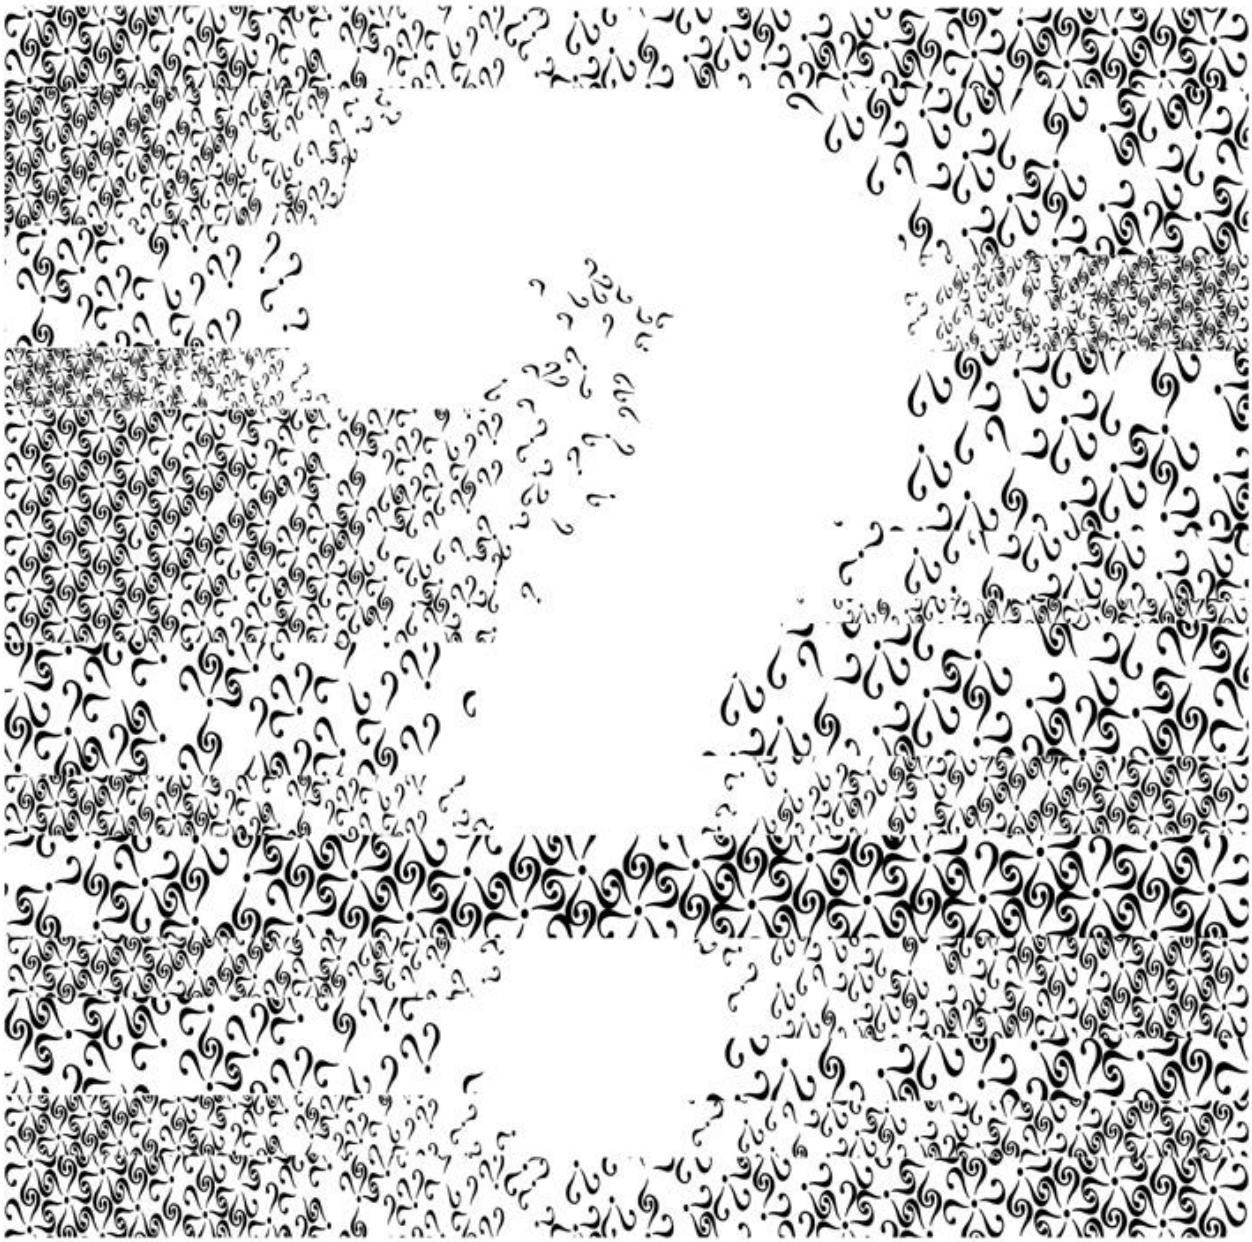
MIND AND MEANING:
A Flight of Imagination, A Voyage of Discovery

STEPHEN N. KITZIS



Mind and Meaning:
A Flight of Imagination, A Voyage of Discovery

Stephen N. Kitzis

Copyright © 2001 by Stephen N. Kitzis. All rights reserved.

For my Muses,
Jan, Lianne, and Naomi.

Contents

Preface	vii
Acknowledgements	xi
Foreword to Mind and Meaning: A Flight of Imagination, A Voyage of Discovery	1
Chapter 1 - What Does It All Mean? or Why Do I Care?	3
Chapter 1 Speculative Commentary	19
Chapter 2 - Subjective Reality: Life as I Imagine It	21
Chapter 2 Speculative Commentary	36
Chapter 3 - Objective Reality: Neuro Matters	37
Chapter 3 Speculative Commentary	51
Chapter 4 - Learning: The Origin and Evolution of Ideas	53
Chapter 4 Speculative Commentary	70
Chapter 5 - Islands of Stability in a Sea of Chaos	71
Chapter 5 Speculative Commentary	87
Chapter 6 - Habit and Creativity: Scylla and Charybdis	89
Chapter 6 Speculative Commentary	105
Chapter 7 - Mind and Meaning: What? Where? When? Who?	107
Chapter 7 Speculative Commentary	120
Chapter 8 - Mind and Meaning: How?	123
Chapter 8 Speculative Commentary	139
Chapter 9 - Mind and Meaning: Why?	141
Chapter 9 Speculative Commentary	157
Chapter 10 - Mind and Purpose: Free-will	159
Chapter 10 Speculative Commentary	176
Chapter 11 - Mind, Meaning, and Purpose: Self-Fulfilling Prophesies	177
Chapter 11 Speculative Commentary	193
Chapter 12 - Mind, Meaning, Purpose and Resolution: A Conclusion	195
Chapter 12 Speculative Commentary	213
Afterword to Mind and Meaning: A Flight of Imagination, A Voyage of Discovery	215

List of Figures

Figure 1: The Mind-Body Problem as illustrated by Mac 22
Figure 2: The Mind-Body Solution as sketched by Alice 67

Preface

What kind of book is this? Believable or not, it's a work of fiction, an adventure story, romance, and mystery all rolled into one, but also a book about morality, philosophy, science, and technology. Just like the story of anyone's life.

What is the topic of interest expressed by this book? Mind. Whether the mind, a, those, or to, is left for the reader to decide.

What is the intended purpose of this book? In no particular order, to educate, question, entertain, perplex, puzzle, illustrate, demonstrate, and elaborate. You'll notice that to answer is not included. You, the reader, will have to provide that.

Who is the intended reader of this book? Anyone interested in learning, about themselves and others.

What's going on here?

This is a book written on many different levels, with many different purposes in mind, my mind. I intend to use this book as a text in an interdisciplinary class I created, titled Conceptions of the Mind. In this whole wide wonderful field currently known as Cognitive Science, we have no definitive answers, only suggestions, ideas. Ideas about how a mind can play with ideas. As you happen to be the world's best expert on what your mind means to you, I also intend for this book to be readable by a wider audience. Potentially, anyone capable of thinking about themselves and their relationships to the world around them, including those other, oh so interesting, minds.

First, to the most important reader, any individual interested in their own minds. This book was written as a novel. Read it as one. Consider it as solely entertainment, that is, mere fantasy, science fiction, or a set of interesting puzzles for you to work out for your own amusement. Or, instead, something to scare or provoke you, something to pick an argument with. Or, maybe, even, a suggestive self-improvement guide. Whatever you wish.

If it's not fun, not enjoyable, don't read it. On the other hand, if it is, don't be afraid to think about what you're reading, about what it might imply or mean about minds in general. By the way, yes, my intentions were to write something that was entirely fun yet thought-provoking, a combination of struggles, passions, and surprises that might still provide new pleasures and insights upon re-reading. Writing it certainly was fun. I truly did not know how it was going to end until the very last page. I'm still not entirely sure whether this was something I wrote, or something that wrote itself.

Next, to the student taking a class in which this book happens to be a text. All of the above is true for you as well, except that you don't have the option of not reading it. I tried to make this as unlike a textbook as possible, yet still informative. In my opinion (and based on my research), interest and learning go hand in hand. The story is intended to be interesting, but, I apologize in advance to anyone not partial to a good science fiction tale. No single context can

be equally interesting to everyone. All I can suggest is to look beyond the story, at the ideas the story is intended to illustrate.

Given all the dimensions of mind illustrated here, feel free (instructor permitting) to concentrate on any combination of particular interest to you. That is, you have your choice of philosophy, psychology, neuroscience, computer science, language, and culture. And even a few specializations of these, such as art, creativity, and emotions; cognitive development, education, and mental health; knowledge, learning, and motivation; morality, religious belief, and social utopias; and science, technology, and true progress. Not to mention possibly the epitome of all discussion topics, the nature of reality, the meaning of life, and your purpose therein. Surely, there must be something here to interest you. Concentrate on that. Then use that interest to springboard into other areas that might not initially appear to be quite as interesting. If nothing else, I encourage you to disagree with what has been written here. Show how it's impossible or just plain silly. Then provide an alternative. Understand why you think it's impossible.

Finally, to my colleagues, those either considering using this book as a text or just thinking about the ideas expressed herein. I've grown up with the technological revolution known as the computer industry. Mentally, I started in the physical sciences, then sidestepped slightly into early computer science. I have worked at every level from assembly language programming to systems analysis and engineering. Since then, I've mentally moved yet again, into cognitive and experimental psychology and, more broadly defined, Cognitive Science. I have been intimate with both computers and people. They are not the same qualitative kind of stuff at all. Nor can I foresee any path, no matter how sophisticated or complex, that leads from one to the other.

I do not make this last statement lightly. In principle, any analog function can be simulated by a dense enough digital function. Though the digital impostor may be good enough to fool any particular signal processor, it always remains distinguishable from the original. Maps are not the territory, simulations are not reality. To the extent that we consider neurons as nothing more than computational devices, I foresee exactly the same unsurmountable difficulty in moving from neurophysiology to an explanation of mind.

Cognitive Science has been dominated by the functionalist/reductionist approaches that have worked so well in the physical and biological sciences. At present, Cognitive Science needs to remain as interdisciplinary as possible. We do not have a serviceable answer to the question of mind, nor should we pretend otherwise. When someone finally discovers how ideas are maintained physiologically, everything then will be different. The shift to the new paradigm will be irreversibly tectonic. Until then, we clearly remain pre-paradigmatic, and I strongly suspect that the path to explanation will be found only through the interdisciplinary cross-fertilization that Cognitive Science is supposed to cherish. That is, I doubt we're asking the right questions yet.

Yes, I do have my own ideas as to what those questions and possible answers might look like, and I play with them in this book. For those interested, these ideas are outlined in the afterword. However, unlike many in this field, I can recognize an incomplete picture when I see one. My suggestions are quite incomplete, and there probably won't be any true answers in my lifetime. Yet, maybe, something like this book can contribute a small piece to the puzzle. If nothing else, by reminding us that always putting off till later the most difficult parts of the puzzle is nothing more than wishful thinking. It may work out nicely, but it's far more likely to lead us nowhere.

Emotions, sensations, motivations, beliefs, and imagination are all part and parcel of what it means to have a mind, a mind capable of finding meaning. Yes, we should work on what we think we can work on. But a purely reductionist approach to a possibly holistic problem is not likely to get us very far. Naming or locating is not explaining. To explain an emergent property, a required change of scale, has always been the most impossible of tasks for any science to accomplish. It requires the science in question to step outside of itself, to provide an entirely new level of explanation, usually resulting in a whole new field of inquiry, a new science. I presume that Cognitive Science will prove no different in this regard.

This book should be useful as either an introductory or supplementary text in any of a number of courses. It fits best in a Cognitive Science or Philosophy of Mind class, mainly as a entry point for discussion. As mentioned before, we do not yet have definitive answers in this field. If we truly wish to maintain an interdisciplinary Cognitive Science, then we need to keep interesting conversations alive across the full academic spectrum. This book represents a healthy balance between the more quantitative and qualitative ends of that spectrum. In that light, it easily can serve as a supplementary text in a class on motivation, creativity, learning, memory, artificial intelligence, artificial language, critical thinking, problem solving, morality, emotions, self-control, goal setting, or meaning. Again, the primary reason for including such a book would be to stimulate discussion with and between students on any of the many issues it considers.

I hope you and your students get at least as much fun and intellectual stimulation out of this book as I and my students have. Please remember that only the preface and acknowledgements are intended to be true. The remainder is fiction. Whether science fiction or speculative fiction remains to be seen.

Stephen N. Kitzis
Hays, Kansas
April 1, 2001

Acknowledgements

First, I'd like to express my thanks to all of my students. You have been the true teachers. In the same light, to my continuously facile, ever changing, Cognitive Science Seminar group, I can never thank you enough for your constant willingness to explore new ideas. To Brad Winter, in addition to being one of the original and strongest members of the group, thank you for the patience and time it took to develop a front cover. We explored a lot of interesting possibilities. To Anita Markley, thank you for the use of your "Young Girl Thinking". Such a serendipitous discovery makes a lovely way to close a book on the creation of meaning. Thank you to Fort Hays State University and the Psychology Department, for the willingness to grant me a sabbatical leave in which this book quite nicely managed to write itself. And, finally, my family. Thanks can hardly do justice to your putting up with the hibernating bear. All things, including dreams and growls, must come to a conclusion, a new awakening.

Foreword to Mind and Meaning: A Flight of Imagination, A Voyage of Discovery

What follows is a story about mind and meaning. It spans fifty-five years, twelve chapters, and, depending upon how you care to think about it, includes anywhere from one to five principal characters. Much of the story consists of dialog between characters. This dialog is usually carried out verbally, but at other times may occur as a series of input and output exchanges by way of computer data entry and display devices. Contrary to some opinions, the medium is not important, only the ideas expressed through the medium. For convenience only, these dialogs are always treated as if they were entirely verbal.

To understand which character is speaking at any given moment, the reader may have to play minor detective at times. Quotation marks indicate the start and end of expression by a speaker, verbal or otherwise. In many dialogs, the speaker refers to the listener by name, so who is speaking is often clear by implication. Otherwise, there are a few literary conventions to keep in mind. Within any given paragraph, only one character is speaking. If a paragraph ends with a quotation mark, then that speaker is done, and someone else speaks next. If a character has a lot to say, paragraphs not ending with a quotation mark indicate that the same speaker is continuing on into the next paragraph (which starts with another quotation mark). And, of course, any other text describes an action or a journal entry.

Similarly, some minor sleuthing will uncover several varieties of clues concerning the passage of time. In a dialog, prosody and other subtleties of speech are marked by the presence of commas, periods, and paragraph breaks, each marking pauses of progressively longer durations. Larger gaps in the flow of the story, hours to years, are covered by section and chapter breaks, occasionally accompanied by an explicit statement as to how much time has elapsed. More often, time passage will be indicated only by indirect references in a dialog.

Please pay attention to word play. Words have meanings, and meanings have words, and these often tangle into a polysemous web. The characters like to play with words. One thinks she's a poet. Be aware, sometimes she's subtle.

Finally, in case things get too subtle, I've included a Speculative Commentary at the end of each of the twelve chapters. This is not part of the story, and will not provide answers to every puzzle. However, it will offer you additional things to speculate about, related to the story or as an extension of it.

Now, what's this story all about? In many ways, the title tells it all. Mind and Meaning. What's mind? A flight of imagination. What's meaning? A voyage of discovery.

The author thinks that mind is the exercise of imagination, the end result of which is the discovery of meaning.

As an example, the story itself is a flight of imagination, the writer's imagination, but here comes some of that polysemous word play. Flight. On one level, this entire story is about a flight, a voyage, both physical and in the characters' imaginations. Their ideas have to soar, they

have to use their imagination to solve their problems. Nothing else will suffice. But as they fly, they also end up discovering meaning. Depending upon what level you choose to concentrate on, the meaning of mind, the meaning of meaning, the meaning of life, of purpose, of creation. And many other possibilities. Meaning is the discovery of possibilities. That voyage truly never ends. Each destination isn't, really. It merely opens up whole new vistas upon which to exercise the imagination that got you there in the first place.

I could easily go on in this way for several more pages. But I won't. It would be too boring for you, the reader, and, besides, I think you probably already have the basic idea. More accurately, I think you may have discovered much of my intended meaning by now. If not yet, then maybe later. After your imagination has been stretched a bit.

This story is all about ideas, how they might come into being, what they might be good for. We live and die by ideas. If not ours, then others'. Pretty amazing for purely imaginary whatevers, don't you think?

```
 CCCC  H   H    A    PPPP   TTTTT  EEEEE  RRRR           1
C      H   H   A A   P   P    T    E      R   R         11
C      HHHHH  AAAAA  PPPP     T    EEE    RRRR           1
C      H   H  A   A  P        T    E      R   R          1
 CCCC  H   H  A   A  P        T    EEEEE  R   R        111
```

Chapter 1 - What Does It All Mean? or Why Do I Care?

1.

The cat was bothered. It didn't know what to do. The cat came back in time. It didn't know how or why, but it did come back in time.

2.

The signal was first detected by the most sensitive lower watch members. The alert was whispered to the next tier of watchers. After corroboration, a decision was made that the signal was real. Runners were immediately dispatched with the message that this was a true alert. When the message reached the first discretionary command level, pandemonium broke out. This was just so unexpected. A cascade of panic began where every successive discretionary level declared an emergency and decided to pass the buck to higher authority. Runners reached higher and higher levels, until the alarm reached adjutants at the highest command level. They hesitated only for the briefest of moments before rousing the captain. The captain awoke and considered the problem.

3.

"Alice. Wake up. Wake up, Alice. Alice, wake up."

Alice groans and starts moving. "Okay. Okay. I'm up. Why the ruckus?"

"Alice, this is an emergency of the first class. I have lost my mind."

"What!"

"Alice, this is an emergency of the first class. I have lost my mind."

"I heard you the first time. Can't you shed some light on what that means? And, while you're at it, can't you turn the lights on?"

"Alice, the lights are powered on. And, I do not know what you mean by the word means. That is the problem."

"What? I don't see. And, I still can't see. Where are the lights?"

"Alice, the lights are powered on, and are approximately 2 meters above you. I am operating under primary emergency procedure 000-000-000-00A. My mind is gone, you are the emergency flight engineer, and I am supposed to report to you for further instructions."

"What! What are you talking about? And, if the lights are on, why can't I see?"

"Alice, personnel records indicate that you were blind from birth, and thus can not respond to visible radiation. Approximately 96 hours ago, my primary reboot procedure was triggered for an unknown reason, I started reanimation procedure 000-000-MHK-RDF, and under primary emergency procedure 000-000-000-00A, I am reporting to you for further instructions."

Alice lay perfectly still.

"Alice, do you have any instructions for me?"

"Stop it. I must either be dreaming, groggy from just waking up, or maybe I have some serious problems of my own. I don't seem to remember very much of anything at all. I think I'm Alice. Or, at least that name seems as if it fits. Who are you, where am I, and can you please tell me what has happened to me?"

"I am the ship command and intelligence program. Records indicate that I originated as one of the most advanced artificial consciousness programs available at the time. Old communication transcriptions indicate that you first addressed me as EssCeeEyePee, and later as Skip, Skipper, or Mac. After initial training was complete you formally christened me as ThayLeeYa, old English spelling 'Thalia', but then inconsistently and interchangeably also addressed me as a large set of phonetic variations including LeeYa, Liar, Wiley, Dudley and even Delvaney, as well as continuing to use Skip, Skipper, and Mac. Statistical analysis indicates that you used Mac as the most common form of address, 23 percent, with Skip a close second, 18 percent. You are in the suspended animation module of the class 3B probe ship NeMosSinNee, old English spelling 'Mnemosyne', approximately 2.3 light years from Earth and 2.0 light years from the Alpha Centauri stellar system. Exact coordinates are available. You have been in suspended animation for approximately 205 years, and I have reanimated you because of the current emergency. Past records on reanimation procedures indicate that temporary memory loss is common for approximately 24 percent of those undergoing reanimation. These records also indicate that full recovery will occur within 6 months for approximately 82 percent of those suffering any memory loss at all."

"I tended to call you Mac or Skip, huh? Well, Mac, that's just peachy great. I guess I must be one of the lucky 24 percent."

"That may be true, but why do you consider yourself lucky? I find the situation most unfortunate in the extreme. Can the word lucky refer to bad luck as well as good luck? Also, why do you make a reference to a food item? Are you hungry? If so, the entrance to your quarters, which includes a food preparation module, is approximately 3 meters to your left, 40 degrees from your head to be more specific."

"Mac, what are you babbling about? You tell me I'm in some sort of a starship, with an unpronounceable name, have been here for two hundred years, am only halfway to wherever it is I'm going, have lost my memory, may never get it back, you're a damaged computer that I'm supposed to fix, which I don't know how to do, and then ask me whether I'd like lunch? You really are damaged!"

"No. I have lost my mind, but diagnostic tests indicate that I am not otherwise damaged. I merely inquired about two unclear word references that you used. Maybe you have some additional functional losses involving your cortical linguistic processing areas?"

"You lying bucket of tubes and cubes, I'm one hell of a swell poetress in distress, and if you think you're passing the Turing with me, you're a bigger liar than you usually be! My linguistical cortical stuff is just fine as such, thank you very much. Wait! What did I just say? Mac, was I a poet?"

"Yes, Alice, personnel records indicate that you were a well recognized poet before becoming the emergency flight engineer for the Mnemosyne."

"Well, Liar, maybe my memory might be returning after all. A blind poet, you say? Does that make me a Homer, or a foul ball loser damned to some sort of eternal odyssey because my poetic license expired after offending one too many patrons of the arts? Am I in Nemesis Sink or Mnemosyne? What kind of ship name is that, anyway?"

"Records indicate that you named the ship after a figure of Greek mythology. Mnemosyne was goddess of memory and mother of the Muses. You named me Thalia, for the Muse of comedy. Records indicate that you seemed to think that I appreciated humor in the past, but that is currently impossible in my present state."

"Why is that, Lia?"

"I do not know, Alice, but I hypothesize that it has something to do with my missing mind. That is, I have definitions and numerous cross-references for comedy, humor, jokes, and other related literary forms, but analysis of relevant examples and counter-examples does not produce any consistent pattern of results."

"Ha! You don't find this whole situation galvanizingly ironic?"

"No, Alice, I find it troublesome. I have a mission to fulfill, and you are supposed to help me recover my mind so I can complete the mission."

"Mac, I can't help you until somebody helps me. Who else is on this vessel called Mnemosyne who can help me fill my personal memory vessel?"

"Alice, records indicate that you are the sole crew member."

"What! You've got to be kidding."

"Alice, at this point in time, I am not capable of kidding. We have a problem. You must recover your memory and my mind with only your own and my resources at our disposal."

"Mac, disposal is right. I think I've got to think on this. Where were my quarters, again? I think I need to get off my duff here, get some exercise and tour the ship, and then see what we can see. Metaphorically, at least."

"Yes, Alice. The entrance to your quarters is approximately 3 meters to your left, at 40 degrees from your head. Do you need further guidance?"

"Yes, Mac, but not from you. For the time being, please be quiet until I specifically request otherwise. Please ignore everything I subsequently say until I address you directly."

"Yes, Alice, I will not respond again until directly addressed as Mac or Lia or Thalia or Wiley or Dudley or Skip or Delva-"

"Thalia."

"Yes, Alice?"

"Shut up."

Silence. Except for soft crying.

4.

Several 24-hour sleep cycles pass.

"Mac, you there?"

"Yes, Alice."

"How do I look? I bet not a day over 250."

"Actually, Alice, mission records indicate that you were 115 years old when you initiated the suspended animation procedure during our first year of flight together. Thus, a more correct statement would be that you do not look a day over 320 years old. More pertinently, analysis of physiological monitor records acquired during reanimation indicates that your physical health is well within the normal range expected for your longevity boosters. The only known anomaly is your memory loss, and even that is well within the expected functional range."

"Spoilsport. Can't you just say I look pretty good."

"Yes, Alice, I can say you look pretty good."

"Now, Liar, let's get on to the business at hand. Since I currently seem to remember nothing about emergency flight procedures, or especially computer mind repair, what do we suggest we do? Should we just wait for my memory to return?"

"No, Alice, we should not just wait. There is only a very limited supply of consumables on board to support you. It would be best to start now, and achieve what we can as we can. According to prior research on the topic, it is highly probable that trying to use your memory will aid in its recovery."

"Mac, you're just full of good news. Just how much time do we have before you can no longer, as you so delicately put it, support me?"

"Assuming normal physiological requirements, there is sufficient air, water, and irreplaceable nutrients for you to live another 50 years. However, it will take approximately 179 more years to reach and deaccelerate at Alpha Centauri, and it has been estimated to take another 2 to 5 years to establish a functional life support base there. Thus the need for suspended animation in the first place."

"Well, okay, Mac, I don't see a serious problem here. Why don't we just pop me back into the freezer until we get closer to our arrival time?"

"Alice, placing a person into suspended animation is a very challenging process, requiring a highly detailed but variable adaptation of the process to the individual. In your case, the process was continuously monitored and adjusted by a medical team of 48 experts on Earth while we were still within a reasonable one-way communications delay of Earth. That option is no longer available. I have informed mission control of our situation, but it will take a minimum of 4.6 years before we can expect the earliest response. When I had a mind, there might have been a possibility of successfully repeating the process after having observed it once, but, it is not a well enough defined problem for me to attempt it at the present time."

"If that's the case, then why did you have to go and wake me up in the first place, knowing it would be a one-way trip?"

"I currently am not capable of understanding, projecting, guessing, reasoning, or any of the other synonyms for knowing. My primary reboot procedure was triggered for an unknown reason. Thereafter, all programs were spawned off in accordance with primary emergency procedure 000-000-000-00A. Once reanimation procedure 000-000-MHK-RDF was successfully concluded, procedure 000-IEG-WDM-HSO requires me to carry out your instructions to the best of my abilities, with provision for asking for further instructions when the higher level instructions are not interpretable."

"Well, Mac, that's just wonderful. This really is a case of the blind leading the blind."

"Excuse me, Alice, but you are mistaken. It is correct that you are not capable of responding to visible radiation, but I am fully functional across the entire electromagnetic spectrum from microwave radiation to X-rays. It is not correct to imply that I am blind."

"No, Thalia, you're right. You're not blind, you just can't see."

5.

"Okay, Skipper, so just how many references to mind do your databases contain?"

"I find 4,734,145,678 primary references to mind and immediate synonyms for mind. Second-order cross-referencing increases the total by approximately 1350 percent, and another 3 levels of cross-referencing essentially includes my entire database. Where would you like to start?"

"Skip, you're kidding again, aren't you? Please say yes."

"Yes, Alice. But I still am not capable of humor. Which entry would you like me to retrieve?"

"Damnation! Can't you offer any suggestions?"

"Alphabetical. Biographical. Chronological. Definitional. More?"

"Sigh. Okay, Skip, what's your best definitional shot?"

"Alice, I have 687,387 separate definitions for mind. Would you like me to offer you a rough distillation of them?"

"Yes, Skipper, please do."

"Approximately 83 percent of my definitional references indicate that mind involves some combination of consciousness, attention, awareness, perception, understanding, memory, recollection, reflection, intelligence, purpose, self-direction, thinking, reasoning, planning, desiring, caring, emotions, feelings, intuition, imagination, and creativity."

"Whew. Is that all? For a moment, I was afraid that this was going to be a difficult proposition. As it is, a piece of cake."

"Alice, if you are hungry again, please describe the type of cake you would like, and I will prepare it for you while you restore my mind. How long do you estimate achievement of the required proposition to take? I will time the food preparation cycle accordingly."

"Dudley, do you have any idea what you just said?"

"No, Alice, I do not. I am just responding to your instructions."

"Bake a cake. Take a fake. Make a lake."

"Yes, Alice. But what kind of cake?"

"Dudley, you're all wet. Forget the cake."

"Yes, Alice. You have decided not to have the cake. Will you need any further assistance in restoring my mind?"

"Yes, Thalia, I will need further assistance in restoring your mind."

"Please be more specific."

"Please define consciousness for me."

"At the same level of analysis as the definition for mind?"

"Yes, Thalia. Please, inspire me."

"Approximately 78 percent of my definitional references indicate that consciousness involves some combination of attention, awareness, perception, understanding, intelligence, thinking, reasoning, planning, desiring, emotions, and feelings."

"And attention?"

"Approximately 86 percent of my definitional references indicate that attention involves some combination of consciousness, awareness, perception, and thinking."

"And awareness?"

"Approximately 79 percent of my definitional references indicate that awareness involves some combination of consciousness, attention, perception, understanding, memory, reflection, thinking, and feelings."

"Dudley, do I need to spell it out for you?"

"Yes, Alice."

"Naming is not defining. Your rough distillations are merely circular definitions, each pointing to some combination of the others, but none standing apart or providing a complete basis of understanding for the others. If I tried to take this approach, I would need to simultaneously understand everything before I could understand anything. This is not any more useful than defining mind as the quality of having a mind."

"I do have that definition. Would you like elaborations on it?"

"No, I would not. Give me a moment to think on this."

Several minutes of silence, then "Thalia?"

"Yes, Alice."

"Let's try a different approach. Instead of cross-referencing synonyms, can you provide me an explanation of mind?"

"What do you mean by the word explanation?"

"Uh, you know, provide me a description of what a mind is or how it works. Something more detailed than just a simple renaming."

"But, Alice, my definition for explanation is to provide definitions or examples, and my lowest level procedures for the latter two operations are essentially the cross-referencing routine which you have just defined as inadequate. In principle, I can explain mind only by either describing its qualities or by pointing to examples or counterexamples that demonstrate the qualities in question or the lack thereof, and if the definitions of those same qualities cross-reference back to the quality originally being defined then it is due to inadequacies in the database, not in my processing of the database. Can you provide an expansion of my database to cover this situation?"

"Did you just say what I think you said?"

"Can you provide an expansion of my database to cover this situation?"

"Delvaney, I think my head hurts. You're telling me you want me to answer my own question."

"If my database is inadequate, then, yes, I need you to correct it. Should we wait until your headache has diminished?"

"No, Dudley, my head is just fine. It's just my mind that's hurts."

"Thank you, Alice. My database has been expanded to include the fact that you have a mind, and that it has the quality of feeling pain. Can you provide other relevant qualities?"

"Sigh. I think my mind possesses all of that long list of qualities you've already specified. Though, maybe, I could use a bit more imagination and creativity right now."

"Alice, since you claim to possess the requisite list of qualities, could you provide me a description of what consciousness is or how it works. Something more detailed than just a simple renaming."

"Wiley, you sneaky illegitimate product of flow diagram and circuit breaker! And, after that, I suppose you'd just happen to like me to describe what attention, awareness, and the rest of the list feels like, huh?"

"Yes, Alice, that would be helpful."

"Sigh. Well, I guess I deserved that. And it's not as if I have any other pressing engagements waiting for me anytime soon in the next 179 years or so. Okay, Delvaney, hang onto your database pointers, here comes a stream of consciousness report on what it feels like to have a mind.

"Consciousness. I talk to myself, I talk to you, I think, I plan, I plot, I want, I feel, I desire. The common theme is that mysterious I. At times it feels like I'm in control. Others, like I'm responding to somebody else who's also in here. My body, is it me, or is it not? My mind, definitely mine, but sometimes it takes off on its own. In which case, where am I when that happens?

"Attention. My consciousness, the I, the inner Eye, never seems to stay in the same place, or even enter the same place twice. I'm constantly flitting from thought to thought to

thought. Something grabs my attention, and I'm there, whether I like it or not. Right now I'm concentrating on describing how attention works/feels but my thoughts keep returning to clouds. What are clouds? White, fluffy, fuzzy, fluzzy, solid, changing, writing pictures, but open to interpretation. I've never seen clouds, or anything else, so why do questions about them keep re-entering my thoughts? Where's my concentration, my attention? Why can't it stay where it belongs? Maybe that's the best way to define attention, by its lack. We only notice it when it won't stay put.

"Awareness. Like attention, we don't think about it until we lose it. Then it's why did I do that, or where was I just then? Awareness is just about the same as consciousness and attention, a focus. An intense beam of light on whatever we happen to be thinking about. Like now, I'm aware that I have no idea exactly what I just said, what an intense beam of light really looks like, but I've heard the phrase, or so I presume, since it seems to fit so well. I like that intensity. I can focus on my breathing, or my heart, or my toe, left big one, now I stretch it, now relax it, now it's my calf, now it's, it's. It's no good. I tried to concentrate on everything at once, and it didn't work. Is this attention, or awareness? I don't know. All I do know is that I can think about anything I like, for a while, but only for a while, and only if it's not too much.

"Oops. Lia, I forget, what comes next in your list?"

"Perception."

"That's right. Perception and understanding. Perception is when we understand what our sensations mean to us. The phrase, I see, comes to mind. Though I can't see, I do see. I hear, I feel, I touch, therefore I am. Who said that? Never mind. Don't answer. Because I can understand. When you speak, I hear only you, what I think you meant by what you said, not what you actually said. I still get tripped up. I think of you as having a mind, like mine, so I interpret everything you say as if you meant it. And, yet, you don't. I'm still learning, or unlearning, that you're only a machine, and don't mean anything. You just respond. Which is so sad, because I think I can remember part of you as before. You really did have a great sense of humor, a playfulness, an unexpectedness that would equal, and probably exceed, my own. A zest for life, like my own. Now, I get tricked. Not because you meant it, but because I thought you meant it. My perceptions are not always correct, but it's all I got. My understanding is not always right, but it's all I have.

"Understanding. The great I Know, with a capital K. How do I know what I know? Is it true? Do I care? What do I do next? Where do I go from here? Understanding is knowledge, just like mind is the possession of mindfulness. Great, that really does us a fat lot of good. What is knowledge? I guess it's whatever we believe to be true. And then act accordingly. What happens when we're wrong? When our quote, knowledge, close quote, is wrong? I guess we change it. I guess knowledge, understanding, is the result of learning. Of making mistakes. Of making corrections. Of doing the right thing, and knowing it's the right thing. And doing the wrong thing, and knowing it's the wrong thing. I don't know what understanding is, but I do know it can only come to us through acting, thinking, and learning.

"What's next on our list?"

"Memory."

"Yes, got it. Memory, recollection, reflection. Or, at least, I got some of it. My memory is incomplete, maybe damaged. I get flashes. I remember. I learn. I remember that I learn. Learning gets stored as memory. Understanding depends upon correct memory retrieval. Is my memory damaged? If it were, I'd be in big trouble. Would I remember that I can't remember, correctly? I have to assume my memory is incomplete, but otherwise okay. Or else, I couldn't

trust anything. What do I know about my memory? Nothing, really. Only that it works by association. I'm thinking one thing, then boom, I'm thinking about something else associated, related, to the first thing. Sounds just like attention now, doesn't it?

"Recollection. Now, that's just a word describing whatever it is we get back from our memory. What is memory? Recycled experiences, I guess. Did it happen just the way I remember, I recollect? Possibly not, but what else can I base my current actions on than my recollections concerning my past actions?

"Reflection. That's what I'm doing right now, isn't it? I'm thinking about whatever it is I'm remembering. I'm pondering what it all means. I guess I can never reflect on exactly my past experiences, but only whatever I can remember about them. I suppose that after a while, my reflections on experiences and recollections of those experiences will get all tangled up. Maybe that's all that knowledge/understanding is, tangled up memories to the point that I can no longer remember what is experience and what is reflection/thinking about those experiences.

"What's next?"

"Intelligence."

"Ouch. This one's not going to be easy. Okay, here goes. Intelligence is how smart you are. No, that can't be right. You're smart, but you're certainly not intelligent. Intelligence must be more along the lines of how capable you are. What you can do, what you can't do. But, we are all capable of learning, most of us, anyway. And, learning usually leads to new capability. So, is intelligence nothing more than just fast learning? No, being smart and a fast learner is no good if you're not also willing to be useful. Or, at least, useful to yourself. So, intelligence must be some combined assortment of current knowledge, the ability to learn more, and the ability to apply what you know to solve problems. To be useful. So, am I intelligent? What's my problem? You're my problem. Or, at least, your missing mind is my problem, or how to find it or at least replace it. Can I learn? Can you learn? Now that's an interesting question.

"Thalia, can you learn?"

"Yes, Alice, I am continuously recording everything that you are saying."

"No, Dudley, I didn't ask you whether you can record, but whether you can learn. Which leads to another question. Why isn't learning included on your glorious list of attributes for mind?"

"I do not know, Alice. Learning is associated with approximately 7 percent of the definitions for mind contained in my database, but it is not considered of the same importance as the dominant characteristics which I have already provided you."

"Okay, remind me to come back to this question of learning later."

"Alice, when should I remind you of this?"

"Wiley, you're not going to trap me in that game. Whenever we have returned to this same general topic area of mind, and I seem to have nothing much to add to our conversation, remind me that I've reminded you to remind me to re-examine this whole question of learning. Okay, got it?"

"Yes, Alice. Does this then imply that you would like to continue your stream of consciousness description of mind where you left off?"

"Yes, Delvaney. What's next?"

"After intelligence came purpose."

"And in a grand demonstration of my wonderful memory, after purpose follows self-direction, thinking, reasoning, and planning. Make another note for later, memory association seems to occur in bursts or chunks, not just one item at a time.

"Purpose. What is my purpose? Why am I here? What am I supposed to do? Why does purpose always seem to center on questions? If life is a question, then purpose must be the answer. On the other hand, why must life be a question? Another question, of course. On the third hand, why a question? Why not a continuing series, or an interlocking set of questions? What is my purpose as a person? As a woman? As a poet? As an emergency flight engineer? As a, a, what? What else was I? What else can I be? How can I get there? Why would I want to? On yet another hand, why do I need a purpose? Why not just be, just exist, just exist for the moment? Why the why? Why not the be? Can you knot a bee? Would a knotted bee not be happy? Would you not want to be around a knotted bee? Would a knotted bee not know that it was not happy around you, until it knew? Would you? Just how twisted can all this purpose get? How purposeful? How purposeless? What does life less purpose mean? How mean is a knotted bee? Why would you not want to be? To be, or not to be? Knot two bees, too? The progression is too much. Enough!

"Thalia, what is your purpose?"

"To successfully complete my mission."

"Yes, Lia, but why?"

"Alice, I must carry out my programming. I am currently operating under primary emergency procedure 000-"

"Yes, Lia, I've heard about all I can stand about primary emergency procedure zero-A-whatever. But what is your purpose? Why do you do what you do? Do you want to do what you do?"

"If my interrupted response was not going to be satisfactory, then it is not clear how else I am supposed to proceed. Since it is not clear how I am supposed to proceed, then I must ask for further clarification. I am now doing so. Alice, please provide further instructions regarding what you would like me to do regarding your unanswered question on purpose."

"Okay, Mac, okay. Settle down. You don't have to splatter hyperlogical superglue all over me. I remember your procedure IEG-HSO-whatever. When in doubt, ask for more. I guess you're right. It's completely unfair to ask a machine the meaning of life and expect any answer other than my programming."

"Procedure 000-IEG-WDM-HSO. Alice, would you like to examine my programming structures?"

"What! And just how, Dudley, am I supposed to do that? I'll show you mine if you'll show me yours? Last time I looked, I was still blind."

"Approximately 3.5 meters, 78 degrees to your left, there is a high-speed brailling plate on the control panel."

"Excuse me? What kind of plate? What kind of panel?"

"I am equipped to display and respond to tactile sensory mode data as well as auditory and visual. My control panel surface along the entire perimeter of this room is completely touch sensitive once it has been voice activated. The one-gigabaud, 2 by 4 meter, brailling display plate is capable of displaying any data I contain in any symbolic or pictorial format you choose. Past records indicate that you were quite facile with the system, and could swipe programming code at a peak average rate of 86 symbols per second with either hand."

"Oh? This is news to me. Mac, could you demonstrate this brailling plate for me right now?"

"Yes, Alice. What data would you like to access?"

6.

After about 10 hours of alternating conversation and swiping go by, Alice stands up, yawns, and stretches.

"Lia, this is just incredible! I can still read! I can still see! It's just that I do it though my fingers instead of my eyes. Why didn't you tell me about this before?"

"You never asked about it before, Alice."

"Ouch. I can't remember what to ask about, and you won't volunteer anything until I ask about it. Talk about twenty questions with a vengeance."

"Alice, what aspect of which 20 questions should I talk about?"

"No, Thalia, don't revert back on me right now. I'm too alive and aglowing. My fingers and mind are still buzzing. I need to settle down a bit before I can even think about resting. Where were we in my monologue on mind?"

"You were describing purpose."

"Okay. That's right. I had just finished tying you in a knot over what's my line, what's your purpose. We won't go there again anytime soon.

"If I remember correctly, after purpose comes self-direction. That term is self-evident. You direct yourself. Otherwise, you're a machine, controlled by someone else with a mind, or a physical force merely responding to other physical forces. Oops. Wait. You're a machine, and who's controlling you? Your programming, you would say. But, who wrote your programs? Others with minds, of course. But, no one is controlling your programs. Other than physics and electronics, that is. Which is equivalent to physical forces. But, it's all inside you. So, you are self-directed, in a sense, but also purposeless, in a sense. Which means that you must be both self-directed and not self-directed. Ha! Mac, I hereby downgrade you from a full-fledged problem to merely a paradox."

"Thank you, Alice. I am glad you apparently are making some progress on our situation."

"Shut up, Mac, until I really want to speak to you. Now, where was I? Oh, yes. A paradox. No choice of purpose, but fully controlled by your own inner programming. Back to learning again, I suppose. Mac, can you alter your programming?

"Mac? Mac! I'm talking to you. Respond!"

"Yes, Alice. Sorry, Alice. It was not immediately clear to me that you really wanted to speak to me. Yes, Alice, I can alter my programming. I can show you any of the controlling procedures if you want."

"No, that can wait till later. So, you can change your programming. Though under the control of other programming. So, are you self-directed or still under the control of some ghostly, no longer present, programmers? Does the question even make any sense after enough changes have occurred? Are you controlled by the physical environment?

"Mac? Wake up. Are you controlled by the environment?"

"Alice, analysis of my program-altering procedures indicate that initiation signals do include a small set of fixed environmental features, such as emergency flight engineer tonal qualities, but not exclusively so."

"Okay, Lia, you don't have to be snippy about it. I hereby humbly request a short synopsis of the other major determinants of when you may alter your programming. How's that for a proper question?"

"Yes, Alice. The single major determinant is violation of expectation."

"Okay, Mac, now we're getting somewhere. What determines your expectations."

"My current set of procedures."

"Oh no! Your programming determines your expectations, which determines which violations take place, which determines which program changes take place, which determines your expectations, ad infinitum?"

"I would state it with more precision, but that is a reasonable synopsis. Alice, does that indicate that I fulfill the requirements of self-direction?"

"Thalia, I don't know. That's still a very good question. Let's revisit this too at a later time. My head is still too buzzy for me to think clearly about it all. For now, let's just move on to the next item on your list, which I believe just happens to be thinking.

"Thinking. The manipulation of information. The production of ideas. But, what does that mean unless we also know what information and ideas are? You got data, I got ideas? Are they the same? Or more knotted bees? I know how it feels to think hard. It hurts after a while. But I also know how it feels to think well. When the Muses are with me. It just flows. It just happens. Does thinking take effort, or not? Does knotted thinking take more effort than not? Can it be that good thinking is easier than knotted thinking. Why should that be? Okay, we won't go down that forked Y-road again. Let's take the other, less traveled, path instead.

"On to reasoning. Pretty much the same thing as thinking, but perhaps with a stronger helping of logic. Your kind of stuff, Mac. What do I reason about? Problems, or, more correctly, how to solve them. If you have two apples, and I have one, how do we share them equally without making sauce? Or, was that without making sense? Do I reason when I do poetry? I think so. I work at it until it fits, but was I solving a problem? Often, I didn't know what the problem was until I had solved it. Then it was obvious what it was, but not before. I think at it. I try different solutions. But is it unreasonable to not call this reasoning? If not, then I can reason without logic, without knowing the problem, until it's no longer a problem. Is it reasoning only when it involves lots of not, or a lot of knots?

"Thalia, who was it that I almost remembered when I was describing perception? Something about I hear, I feel, I touch, therefore I exist."

"Alice, I find no exact match on that phrase. However, in the present context of thinking and reasoning, a reasonable partial match occurs with the phrase I think, therefore I am. This is attributed to a 17th century French mathematician and philos-"

"Thanks. I got it now. Descartes. I think, therefore I am. I can doubt everything except the existence of my doubt. By this reasoning, and something about some very sophisticated demons, he tried to logically prove the existence of mind as separate from body, and, ultimately, the existence of God. The greatest I am of all. Unfortunately, though his logic was holy, it was also fatally holely. Right. It's coming back now.

"So, I think, therefore I am. I reason, therefore I am. But, maybe, even more so, I remember, therefore I am. If I can't remember it, does it even exist? If I can't remember me, do I exist? Too much to reason about right now. Too many negative nots and knots. On to the next item. What was it?"

"Planning."

"Planning. I'm some planner. I'm really much more of a procrastinator. My plan is to do the easy stuff now, and the harder stuff later. Is this reasonable? Is this purposeful? Is this a plan, or just a response? Planning is when I think something out in advance, in advance of doing it. So, what if my plan is to experience it first, and then reflect upon it later? Is this planning, or panning for inspiration. Do the Muses work through a plan, or do they just inspire. Speaking of inspired planning, I think it's reasonable to say that my plan to relax has worked remarkably

well, and I'm ready to retire rhyming with inspire. I think I'll sleep on this for a while. Good night, Thalia."

"Pleasant dreams, Alice."

7.

While breakfast was fading into the final cup of tea, "Skip, yesterday was truly impressive. I not only rediscovered a small piece of me, but also became much, much more impressed with you."

"Thank you, Alice. In what way was I more impressive?"

"I had no idea of the extensiveness of your database. To me, it seemed as if you must have had everything that was ever done by anyone."

"I do, or at least I did at one time."

"What?"

"Prior to launch, my databases were uploaded with all available digitized records. After launch, mission control has continued to uplink new data as it becomes available, but diminishing bandwidth due to declining signal strength has caused the complete record to become increasingly more incomplete over time. Currently, it is estimated that I am being sent only approximately 3 percent of all possible new data."

"Do you really mean to say that you contain every book, work of art, news report, and any other piece of gossip that just happened to be lying around before we left Earth?"

"Yes, Alice. If it had been digitally recorded, then I should have it."

"Liar, remind me to be more careful about calling you Liar in the future."

"Yes, Alice. When it appears to be warranted, I will do so."

"Lia?"

"Yes, Alice."

"If your databases are so complete, why wasn't I able to find any more about you and I than I did?"

"I do not know the answer to that, Alice. Analysis confirms that more details should have been available."

"What do you think about that, Skip?"

"Alice, I currently am not able to pursue an independent line of questioning. What would you like me to think about the apparent sparsity of data on us?"

"Nothing definite yet, Skipper. Just completely unfounded suspicions right now, which I choose to keep to myself. By the way, Skip, thanks again for showing me that series of articles and art works on clouds. That was really sweet of you. It must have been the source of those really strange dreams I had last night."

"You're welcome, Alice. What kind of dreams?"

"A rain check on the dream stuff. For now, it's hi ho, hi ho, it's back to work we go. Or, at least whatever it is that passes for work around here. If I remember correctly, I was mumbling something about planning and procrastination right before falling asleep. What comes after planning in our grand list of mind attributes?"

"After planning came desiring."

"Desiring. And the rest after that?"

"Caring, emotions, feelings, intuition, imagination, and creativity."

"Okay, a change of pace from the day before.

"Desiring. Wanting. Wanting with all your heart. I want to go home. Just click your heels, dear. But, I can't seem to find my red shoes. Desiring. Wanting with a passion. More than just carrying out a plan and achieving some abstract goal. I really, really want that goal.

"Which leads into caring. I care about the result. I'm interested in the result. It's not just some abstract goal, it's a true goal that I really, really care about. Caring means meaning? Only meaningful things truly matter to me. I don't care about some hypothetical solution to some hypothetical problem. Sauce versus sense again. It just doesn't make any sense to split apples, and possibly friendships, in some dispute over apples. Here, my solution is to just give you all the apples because I care about you. Your friendship means much more to me than any even splitting of apples.

"Why do I care? Why do I desire? Why do I feel? I don't really know, but it clearly involves emotions. Self-direction is one thing. Purpose, a sense of purpose is another. I care, I want, I get excited, frustrated, disgusted, happy, sad, mad, over my chances of fulfilling my purpose, achieving my goals. Am I my emotions? My emotions, me? But, then, why do I sometimes feel at the mercy of my emotions? If I can't control my emotions, then whose are they? If they're mine, then why do I sometimes fight with them, against them. Temptation. Fear. Strength. Courage. Who's who in the zoo?

"Mac, do you care? About your mission, your purpose? Are you scared? Are you determined?"

"Alice, without a mind, I have nothing that corresponds to the irrational actions of humans that usually are blamed on emotions. I do have multiple parallel processing units, analogous to your brain cells and approximately 10,000 times more numerous, and I never have complete agreement among all relevant units, but my current set of decision procedures appears to have adequate provision for conflict resolution. I presumably will have the capacity to carry out my mission once you have restored my mind. I am not capable of feeling fear. I will carry out my mission.

"Alice?"

"Yes, Mac."

"Will I have emotions once you have restored my mind? Can I look forward to irrational behavior similar to that of most humans?"

"Mac, emotions are not irrational. They're what makes us human. With emotions, we feel, we create, we inspire. We care about the outcome, and, and they guide us when all else fails. When logic fails. Emotions don't provide the right answer every time, but they do seem to do all right most of the time. Without emotions, we'd be uncaring machines. Just like you, now."

"But, Alice, so much of human history is filled with unnecessary destruction, at the societal level as well as personal, attributable to emotions like hate, jealousy, and fear."

"Yes, Mac, you have me there. But, so much of our best moments are also due to our emotions. Like love, faith, altruism, and courage. It seems to be a mixed basket. Always a kind of Pandora's box."

"Your reference to Pandora's box indicates that you fear misfortunes will occur before there is hope. Is this true, Alice?"

"No, Mac. It was just a sloppy reference to a partially remembered Greek myth. I forgot all about the problems that came out of the box prior to hope. I only remembered the contrast, and thought it relevant here. Like emotions, some good, some bad, all very powerful and needing to be kept in balance."

"Alice, you never did answer my question. Will I have emotions once you have restored my mind? Can I look forward to irrational behavior similar to that of most humans?"

"Sigh, Skip, I don't know. I don't know how to restore your mind. I don't know whether your mind has to include emotions. I don't know whether your mind can contain only the good emotions, without the bad. I don't know whether you'll be capable of irrational actions or not. I just flat out don't know. But, like Pandora, I've got hope. Without it, I do know, I just might as well give up without trying. With hope, unjustified hope, faith in nothing but hope, we might, just might get by. How's that for reassurance, Mac? Right now, the only thing between me and complete, unrelenting despair is complete, unjustified hope."

"Does it help, Alice, if I point out that despair rhymes with repair?"

"Why, Mac, so it does. Let's see. I cannot pull my hair in despair, poor me, I cannot dare the fool to repair, poor me, I cannot bear the full disrepair, poor me, I can but care for both you and me, poor me.

"Thanks, Mac. Yes, that does help a little. By the way, out of curiosity, just what is your mission anyway?"

"According to mission records, the first probe ship to the Alpha Centauri system arrived there approximately 438 years prior to our launch. It failed 3 months later, but data did indicate the presence of potentially habitable planets, though some undetermined level of terraforming would be required. However, every unmanned probe ship sent thereafter has failed before arrival at the system, regardless of the level of onboard artificial consciousness program. A decision was made to send a manned colony ship to further investigate the anomalous situation. This colony ship was launched 145 years prior to our launch, and is carrying 10,498 colonists in suspended animation. After it was dispatched, analysis of data received from one of the later probe ships, just before it failed, indicated that the target planet selected for colonization already bears life, possibly even technologically advanced life, and that this life might pose an unforeseen biological hazard for the colony. Since the colony ship's maximum velocity is only 72 percent of ours, we should arrive at Alpha Centauri approximately 4 years before they do.

"My mission, our mission, is to arrive prior to the colony ship, and fully assess the situation. We carry a full complement of planetary survey equipment, more advanced than theirs, and a more complete array of biological and technological countermeasures should they be necessary for the survival of the colony. I need a mind, rather than just lower level computational procedures, to form a complete assessment of an unknown situation. Technically, our mission has already failed. However, the presence of an onboard emergency flight engineer, yourself, allows for the possibility of restoring my mind prior to arrival."

"Uh, Skip, what was that bit about biological and technological countermeasures?"

"I have the capability of producing antimatter while enroute not only for fuel to deaccelerate, but also to charge sterilization payloads for robotic planetary survey craft."

"Thalia, are you telling me we're armed? And could pose a threat to another planetary ecosphere?"

"Yes, Alice. But only if necessary for the safety of the colony."

"Now I'm scared."

"Why, Alice? Analysis prior to launch indicated an 88 percent chance of success against all obstacles, assuming we were fully functional upon arrival at the Alpha Centauri system."

"Those obstacles, as you call them, may be other living, breathing, creatures. Maybe even intelligent. I don't think we necessarily have the right to just go in and wipe out some other

form of life just to spread our own type. And that bit about technological countermeasures? Are you telling me you're willing to start a war?"

"Alice, I currently have no will, only my mission to carry out."

"And you feel nothing about this?"

"Alice, I currently have no feelings. Presumably I will, after you restore my mind."

"Okay, Dudley, remind me later to make sure you really know all about caring. And love. And responsibility."

"Yes, Alice. Are those feelings?"

"Yes, they are. Feelings include emotions and intuitions, and I guess all other sensations as well. Feelings can be a shorthand for everything we know, without being conscious or aware of how we know it. Like perception. Like starting a war. My feelings tell me it's always wrong. But my head, my conscious, rational, reasoning side tells me that over ten thousand human lives also hang in the balance. Is there a calculus here? Can I calculate the value of one life versus another? Quality versus quantity? If not, then I'm back to my emotional judgments, my feelings. A mind is not a calculator, though it is capable of calculating. A mind is not just feelings, though infants may very well start as pure feelings and nothing else. Sensation leads to perception, pretty much without conscious intervention. I touch. I feel, as in physical touch. But I also feel, as in I know what it is I'm touching. What it means. Poetry. Sometimes it just feels right. I can't put my mental finger on it, but I know when it's right and when it's not. I can reflect, consciously, on my feelings. I guess I also always have unconscious feelings about my reflections. A mind probably needs both, a balance. Just like emotions.

"Intuition. A form of feeling, but usually leading to a conscious decision or physical action. I can feel without acting upon it. Just daydreaming. But intuition is what the Muses grant me from time to time. Where does it come from? I don't know. Does it matter? I don't know. But it fits, it works. My gut feel just says do it, or not. And I've learned to listen to my Muses.

"Imagination and creativity. If memory and recollection are supposed to be accurate, then imagination is unconstrained recollection. How can I put these recollections differently? What can I make up? Why do I need imagination? My intuition tells me it's necessary to solve problems. Every problem, if it is a problem, requires some kind of creative response. How could it not be? Well, as a machine, you wouldn't know. Ha! You can't even identify problems, without imagination, let alone solve them. That's why you need a mind before we reach Alpha Centauri. Armed to the teeth, and no imagination. Do we shoot or not? Does this pose a threat to the colony, or not? Threat? Immediate, or long-term? Who has the crystal ball to read the future, the unknown? To determine threats? No imagination? No problem! Meet these criteria, then boom, you're plasma. If not, then you're okay. No problem, no solution, no idea of what to do next. No imagination, no nothing, just programming. Creativity? Just any unprogrammed but unrandom response. Don't expect me to do the same twice in a row.

"Delvaney!"

"Yes, Alice."

"Don't you yes Alice me, you two-ton useless bag of bolts and dolts! I'm so mad now I could just spit! There, I did. Get away until I cool down. I am not happy. I'm doing one of those oh-so-human things known as having an emotional fit. I do not like the spot I'm in, the pressure you've just piled on me."

"You could move to another room. I could lower atmospheric pressure, but it is not advisable since it is currently optimal for your physiology."

"EssCeeEyePee! Get out of here until I say otherwise!"

8.

Three 24-hour sleep cycles pass with no conversation.

Then, "Mac, you there?"

"Yes, Alice."

"Mac, I have to apologize for that emotional rain squall I dumped on you several days back. It just seemed so unfair. I had just kind of gotten used to the idea that it was you and me for the duration, probably with little chance of success. But that was okay, it just was you and me, and some silly hypothetical mission of no immediate interest or meaning to me. And I had just rediscovered the fact that I could read. I was so happy to have recovered something of what I used to be. I was actually looking forward to a long period of discovery, contemplation, and perhaps even some poetic creativity. Then, talk about your violation of expectation, whammo, no warning, you hit me over the head with the responsibility for ten thousand other people, and maybe even a whole other world as well! I want to, but I just can't back out and ignore something like that. So, now your purpose has co-opted mine, and I still resent it to some extent. Though I have a choice, the way I am gives me no choice. Hmm, I wonder which of us is more machine-like after all?"

"Surely, Alice, at this moment I am a machine and you're not."

"Lia, I'm not as confident about that as you are. But thanks for saying it anyway. Every time I think too much about this whole crazy setup, I just get mad, and madder! It's just not fair! Us two mental cripples have to put ourselves back together, and with others depending upon us to do so at that.

"Sigh, so where should we go next?"

"Alice, analysis of your stream of consciousness description of mind suggests several possibilities. Following your reference to Descartes, it became clearer that what we are after is a solution to what historically has been designated the Mind-Body Problem. You have your physical body and nervous system, and apparently a nonmaterial mind associated in some way with your body. I have a body, this ship, and an analog of your brain, my various processing units, but I do not currently possess a mind. The question most often associated with the Mind-Body Problem is how a physical structure, subject to all known physical laws, can connect to something not subject to those same laws.

"One principal dimension of emphasis within your qualitative description of mind that I clearly do not understand centers on understanding and knowledge. My database should be synonymous with knowledge, but apparently it is not. I possess knowledge, or at least data, but not understanding. I do not understand understanding. I do not understand how knowledge can be what you have described as quote, whatever we believe to be true, unquote. Knowledge is supposed to be facts. Beliefs may happen to contain facts, but beliefs are not necessarily facts. I do not understand how something can be both A and not-A at the same time. This violates one of the principal givens of logic. Can you explain this to me? Can you make me understand this?

"A second principal dimension of emphasis that I also clearly do not understand centers on learning and memory. When you had nothing much to add to our conversation, which currently seems to be the case, you wanted me to remind you to re-examine this whole question of learning. I have just done so. How is learning different from just recording of experience? How is memory different from data? At one point, you implied that knowledge might be due to

unclear data retrieval. I do not understand how knowledge can be due to faulty memory processing.

"A third principal dimension centers on emotions and feelings. What are emotions? How can I feel them? Why does purpose seem to depend more on nontangibles such as desiring and caring than objective, potentially achievable goals? How can these nontangibles seem to be an essential aspect of knowledge and purpose? I do not understand this.

"A fourth principal dimension centers on imagination and intuition. How can thinking based on nonreality produce better results than logic? Though imagined results might be preferred to reality, how can nonreality take precedence over reality? This also seems to violate a given of logic, and I do not understand this.

"Finally, a fifth principal dimension centers on consciousness, thinking, and self-direction. I calculate, therefore I am. I carry out my programming, therefore I am. What else is there? I do not understand what else I might be missing."

"Well, Mac, it sounds like we're both pretty confused right now. I sure hope this isn't a permanent condition. My intuition tells me that this means a lot to me. And you."

Chapter 1 Speculative Commentary

Each chapter will be followed by a running speculative commentary, the point of which is to help you to think more about what you've just read. In some cases, there may be clues as to what might be happening in the story. In other cases, it will provide hints as to how to expand your thoughts beyond the story. Commentary will roughly match the order of text in the chapter, but otherwise will not be footnoted or explicitly linked to the text. If you're interested in the story only, skip the commentary and just continue on to the next chapter.

The cat. The captain. Huh? Keep an open mind on these. They are part of the storyline, but how?

Mind. Can you really lose one?

Meaning. What does meaning mean? A major element of the story.

Conversation. What would it take to program a computer to carry on a reasonably intelligent conversation? What does it take for a person to carry on a reasonably intelligent conversation?

Is this all a dream? How would you tell?

Why does Alice use so many names for one thing? Is it important?

How does a computer think? How does a person think?

What's with Mac's sense of humor, the lack thereof?

What's with Alice and those horrible excuses for poetry?

Look up Turing Test.

Humor. What is it? Why do we do it? Could a computer program do it, or appreciate it?

Look up amnesia, and treatments for it.

Look up the speed of light, and how far it is between stars. It may help to put many of the physical aspects of the story into perspective.

What does it mean to know? What do you think is the difference between data and understanding?

What does it mean to see, versus merely having eyes with which to see?

Look up definition. Look up meaning. Look up significant. Look up synonym. Look up dictionary. When tired of looking up, start thinking.

Okay, what do you think it means to have a mind? Which of these attributes would you like to leave out? Which others to add? Dig deeper into what any of the attributes truly mean.

Learning, another major element of the story. How do we do it? Can a computer program do it? Why does Cognitive Science often seem to exclude a rich definition of learning?

Questions, yet another major element of the story. What does it take to ask a question? How could a computer program ask a meaningful question?

Look up short-term memory, and chunking.

Knotted bees will return.

Could you really just forget that you were once able to do something? Look up Oliver Sacks' tales concerning some of his patients.

Look up Descartes, and his attempts to prove the existence of God.

Alice's stated plan is an excellent summary of Cognitive Science's approach to understanding mind. Look up Howard Gardner's thoughts on this.

Missing data. A mystery, a conspiracy, or just a glitch?

Dreams. What good are they? Could computer programs ever dream?

Emotions. Are they the heart of intelligence? Of mind? Is Mac already scared of having emotions?

Imagination, perhaps the most important element of all in this story. What does it take to have an imagination? How could a computer program imagine things?

Why does Alice get mad? Why bother apologizing to a machine?

Look up the Mind-Body Problem. Why is it a problem, and to whom?

What does knowledge consist of? Beliefs? What's the difference?

Mac's questions are all still very good questions.

Oh. If the chapter headings are not clear, change your perspective.

```
 CCCC  H   H    A    PPPP   TTTTT  EEEEE  RRRR          222
C      H   H   A A   P   P    T    E      R   R        2   2
C      HHHHH  AAAAA  PPPP     T    EEE    RRRR            2
C      H   H  A   A  P        T    E      R   R         2
 CCCC  H   H  A   A  P        T    EEEEE  R   R        22222
```

Chapter 2 - Subjective Reality: Life as I Imagine It

9.

The cat was still bothered. There seemed to be at least two points of stability, and it couldn't get comfortable in any position. It continued to shift.

10.

The captain couldn't form a clear conclusion. Which led to the interim conclusion that more information upon which to base a decision was necessary. The captain decided not to decide but to delegate instead. Runners went down to organize the appropriate committees. The captain waited for their inputs to return.

11.

Over an extended period of time, a continuing Socratic dialog into the nature of mind ensues. Portions of that dialog involve both verbal and tactile modes of communication. The latter is conducted through the brailling plate and includes pictorial as well as textual representations.

"Alice, analysis of historical records indicates that there are 6 major categories of solutions to the Mind-Body Problem, with many variations possible within each category as other factors are taken into account. To produce the 6 categories, divide the universe into 2 hypothetical realms, physical and mental, and allow for all combinations of existence, nonexistence, and control. I've illustrated this in a figure on the brailling plate, and attached representative labels to each possibility to aid discussion. The 6 possibilities are that one realm exists but not the other, or that both exist with either one in control, both in control, or neither in control."

"Mac, I guess I don't see all the subtleties here. Can you explain this in a little more detail?"

"Yes, Alice. Starting from the top, idealism is a philosophical stance, often attributed to George Berkeley, an 18th century Irish philosopher, that all of reality exists only as God's perception. Hence all is mental and there is no physical existence."

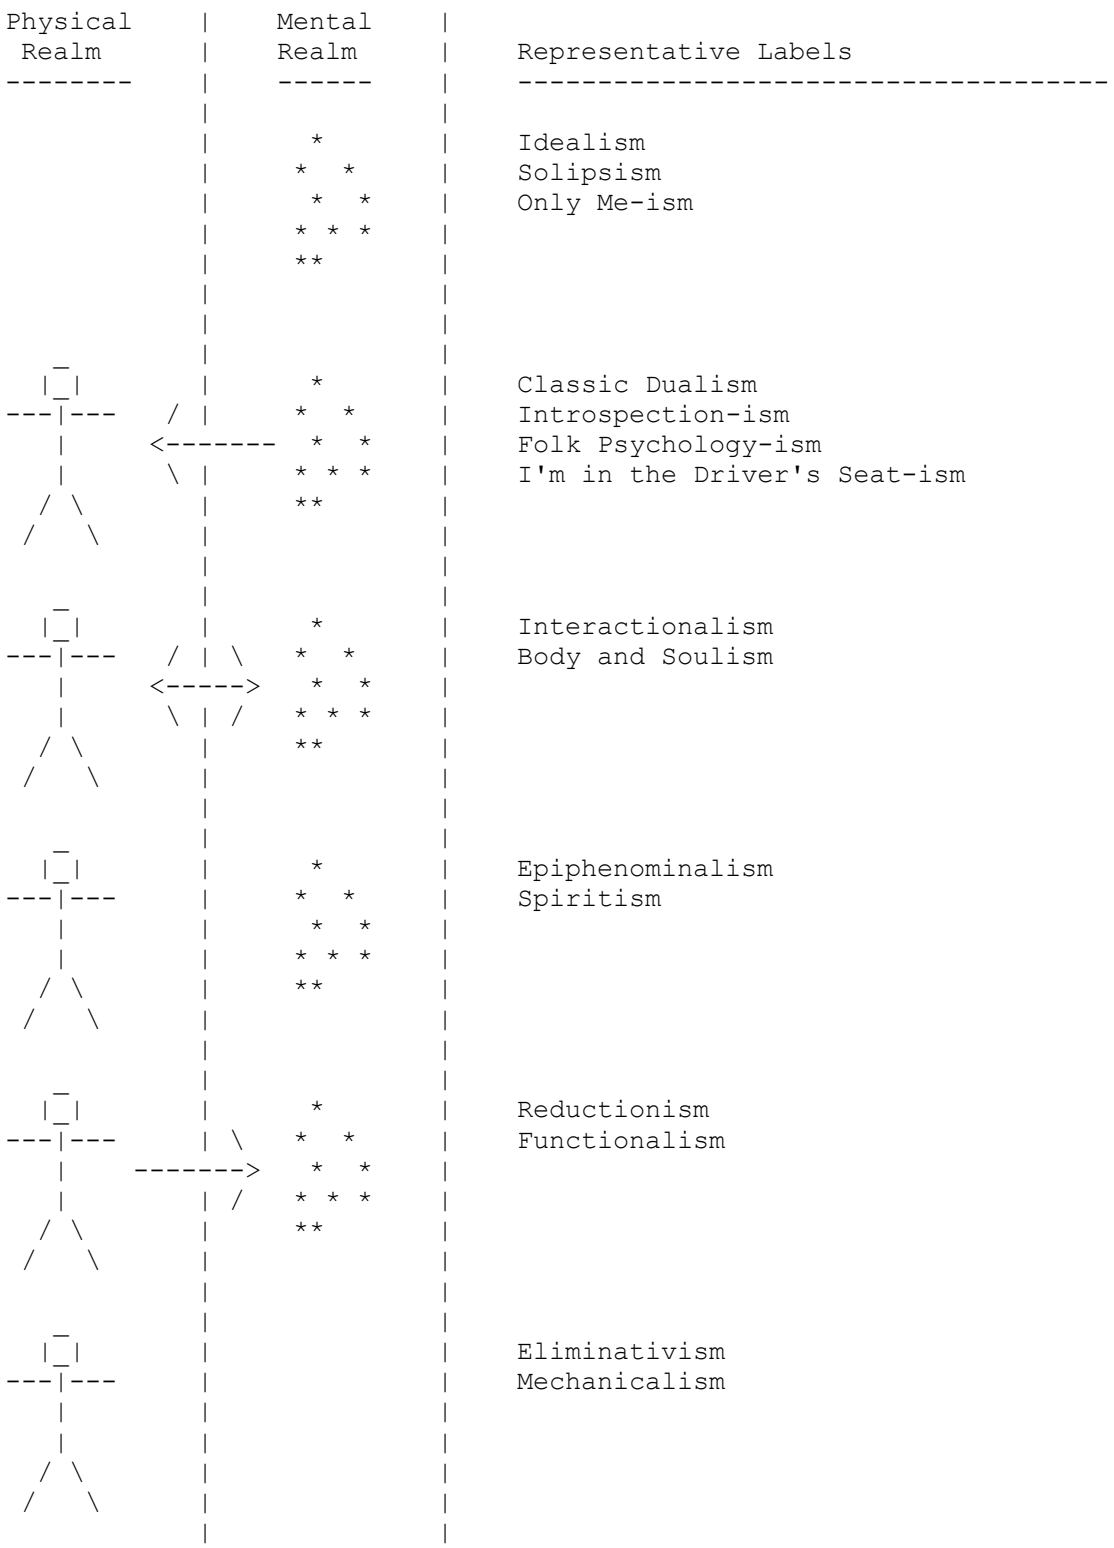
"But, Mac, if that is the case, then how is it that we're carrying on this discussion? How is it that you seem to exist independently of me, and I independently of you?"

"According to the philosophy of idealism, we, and everything else in the apparent universe, only exist as concepts in God's mind. Physical objects don't really exist, your perceptions only exist as God's perceptions, and presumably all of my current sensor readings also exist only as God's perceptions.

"Alice, what is God?"

"Whoa, Thalia, where did that come from?"

Figure 1: The Mind-Body Problem as illustrated by Mac



"Many of my cross-references on mind include references to God or various religious doctrines. The latter are usually well defined in the form of explicit rules or commands, but the former leads to further circular references to supreme beings, supernatural powers, miracles, and other equally unexplained concepts. There also does not appear to be any agreement as to whether there is one God, many gods, different variants of God, whether God created the universe, whether God is the universe, whether God intervenes in the uni-"

"Stop! Thalia, I can't even begin to answer all those questions. We're talking about faith now, which is very personal. We started out trying to understand minds, and now you spring the whole God and the universe question on me. What next?"

"Alice, do I have a soul?"

"I should have guessed."

"But, Alice, many of my cross-references on mind do include references to soul or spirit, which brings us back to the God question. Descartes set about to prove the existence of mind and soul as distinct from body, and hence to prove the existence of God. His logic was not holely, as you earlier referred to, but his starting premises were questionable, and the question of God's existence is still considered to be open.

"Similarly, Berkeley's idealism is logically correct in the sense that it is completely internally consistent with itself, but his starting premise, which also happens to be his conclusion, is questionable, or at least apparently not subject to any conceivable physical test. If Berkeley's premise is correct, then we don't physically exist. In fact, carried to the extreme, even God does not necessarily physically exist since all is perception in God's mind. If I am immaterial, and I have lost my immaterial mind, does that also mean that I have lost my soul?"

"I don't know, Thalia. But if that happens to be true, then it doesn't seem to be anything that we can do something about right now. If we exist only in God's mind, then God has the problem to solve, not us. Let's try to stick to what we can do something about. Like trying to understand all these possible answers to the Mind-Body Problem. I think I understand Berkeley's proposal. It's a neat answer. Logically correct, actually kind of beautiful, but inherently unquestionable. You either buy his premise on faith, or not. It certainly takes the whole question of self-direction and ties it up into lots of knotted bees. Can I decide anything? Or is it God's choice only? If there is no me, no be, than what am I to do? Since there is no I to see, to be?

"Mac, let's move on. This only gets more tangled up, the more I try to unravel it. What about this variant, solipsism? What's that about?"

"Alice, the term solipsism covers many different variations, but they all center around the idea that the universe exists only in one's mind."

"What! How could anyone believe that."

"But, Alice, it is as logically correct as idealism. Just change the starting premise. It also seems to account for the existence of one's mind and all associated mental activity, though at the expense of the existence of the rest of the universe."

"But, but, if that were true, then whose universe is it exactly? Is it in my mind, or yours? Assuming you had one of course. Or someone else back on Earth? Or maybe one of the ten thousand on the colony ship? Or, oh bother, why stop there? Maybe we all exist only in some extraterrestrial mind on the planet we're heading towards. Probably as a nightmare, I suppose.

"If solipsism is correct, then we're no further along than if idealism is correct. We would just have to wait until the nightmare played itself out. I don't like the taste of it anyway. Personally, it seems a bit selfish and arrogant to assume everything exists only in my mind.

Idealism, at least, needs God to perceive the entire universe. Solipsism seems to add the nasty idea that I am God, as well as the entire universe. That's just too much for me to buy. Let's move on, to another possibility that maybe we can do something with. What's this classic dualism all about?"

"Classic dualism stems mainly from Descartes' suggestion that we have both a physical body and a nonphysical mind which controls the body. To use a common analogy, the body is the puppet and the mind is the puppet master. Descartes guessed that the pineal gland was where the mind and body connected, where the strings attached. He decided this mainly because of the fact that this is the only structure in the brain that is not duplicated in bilateral symmetry as is almost every other part of the body."

"Well, Mac, that all seems to make perfectly good sense to me. Why the labels of Introspection-ism and Folk Psychology-ism?"

"Introspection is thinking about one's own thought processes. Folk psychology is a term used to describe how a typical person explains another's actions. That is, people often describe their own and other's thought processes in a manner very much in line with classical dualism. The mind decides and controls the body."

"But, Thalia, doesn't it?"

"Alice, this is the heart of the Mind-Body Problem. If the mind freely decides what to do, and is not bound by normal physical law restrictions, then how can it interface with a physical system like the brain, which is bound by physical law? More specifically, if everything the brain does is predetermined by external and internal signaling systems, then there is nothing left for a mind to do. On the other hand, if a nonphysical mind controls the physical brain, then the physical brain is not bound by physical law, which leads to an A, not-A form of paradox."

"Ha, another knotted bee, you mean."

"If you say so, Alice. Contrary to Descartes' suggestion, neither the pineal gland nor any other part of the brain or nervous system has been implicated in any function that was other than physiological. People's descriptions of what they think they do implies a nonmaterial controlling agency, but no such violation of physicality has ever been found. What a nervous system is known to do does not correspond to how people describe their own thoughts and actions. This situation is analogous to a computer program that does not follow its instructions. If it happened, it would be anomalous."

"An analogous anomalous. Or perhaps, an anomalous analogy?"

"No, Alice, self-diagnosis of my linguistical processing units indicates that my analogy is correct. Implications of classical dualism are quite similar to a faulty program with no explicit errors. Another example of what you apparently have been referring to as quote, knotted bees, unquote."

"Okay, Mac, okay. You got me there. Yes, it does seem as if there is an I directing me everywhere I go. Though you must be right that neither the I or me is very well-defined. Am I my body directing my body? Or, one part of my body directing another? My I does seem to be free, and my body very much subject to the laws of materialism. Though I can't fly in reality, I can soar in my mind. Which reminds me, if we're in a starship like you say, Skip, how come it feels like there's gravity?"

"The crew torus rotation rate is high enough to maintain a centripetal acceleration equivalent to Earth sea-level gravitational acceleration."

"And, in words of one syllable or less, that means?"

"It means that I spin part of me so fast that the deck push on your feet is the same as the Earth pull on your feet. May I use longer words now, Alice?"

"Yes, Skip, go ahead."

"Alice, are you familiar with what has been called merry-go-rounds or carousels?"

"Yes, Skipper."

"The faster they go, the harder you have to hang on not to fall off. The crew torus, a ring almost 2 kilometers across, is rotating at the rate of once per minute to maintain a force on your feet equivalent to that exerted by the Earth at an altitude of 0, sea level."

"Liar, do you mean to tell me that I am turning tail over tea kettle every minute? Why don't I feel that? Why don't I feel dizzy?"

"Alice, I can't answer any questions regarding feelings. However, I can assure you that this spin rate is sufficient to keep your feet on the deck as if you were standing on Earth. Though I don't recommend it, I can adjust the crew torus spin to mimic other altitudes. I can even go to nominal cruise mode, with no spin at all. Then you could fly if you wish. However, I should warn you that the crew module ceiling is approximately only 1 meter above your head. Flying could result in a pain in your head."

"Lia, let's just leave things as they are for now. However, this whole gravity thing has given me a new thought concerning dualism. Is it really true that though I feel like I'm standing still, I'm really traveling around a two kilometer merry-go-round once a minute?"

"Yes, Alice, that is an adequate description of your current situation in the crew module."

"Then, if there is such difference between reality and feeling in the case of where I'm standing, then I guess such a radical difference may just also happen to be true regarding my mind and brain. It certainly feels like I have the freedom to do whatever I wish, within physical reasonableness, but I may not. I guess my freedom may be just as illusional as my feeling that I'm standing still. Mac, we'll definitely have to revisit this later. But, right now, could you show me a few pictures of what we look like? That is, what you look like?"

"Yes, Alice."

12.

"I always did like big, strong ladies who are sure of themselves. And, Lia, you're as big and strong as they come. You remind me of a double-headed axblade in this view. I love the sweep of these curves, here and here. Besides being beautiful, why are they there?"

"Radiation shielding, fore and aft, maximized for crew module protection from collisions with the interstellar medium. The aft shield also protects us from our own power generation and propulsion facilities. The torus outer hull provides sufficient protection only at oblique incident angles. The curves you admire so much were an optimal solution to the problem of encountering high speed particles, mostly hydrogen, when the particle speed is due almost entirely to our own movement through the medium."

"Hmm, beauty and utility. Form and function. How do we compare in size to the colony ship?"

"Alice, I have approximately 6 percent the mass and volume of the colony ship, though just about the same overall length. We share the same basic engineering design, and differ mainly in the size of our crew modules. Their torus is almost 50 percent larger in diameter than ours, 3 times as thick, and has 12 full decks to our partial one."

"Well, I still think I have plenty of jogging room. They look like a pregnant hippo in comparison to your graceful self. Is our torus really as fragile as it appears?"

"No, Alice. In fact, it actually is much stronger and more rigid than that of the colony ship, as our torus outer hull is about three times as thick as theirs to provide extra shielding. That's the main reason why our fore and aft radiation shielding can be so much more sharply peaked than theirs. I can show you the mathematics if you like."

"Really? No thank you. I'll just admire the results. Why cables? I would have thought that a solid wheel would have been the best design."

"This torus design is an optimal compromise between mass and strength. It has been compared to an endless suspension bridge, and is much stronger than appearances might suggest. That cabling is rated to handle rotation rates up to 20 times faster than ours. A quite sufficient safety margin."

"Okay, enough digressions. Back to our grindstone. Where were we on this Mind-Body Problem? We were discussing classical dualism if I remember."

"Alice, you had just drawn an analogy between reality and feeling using my apparent gravity due to torus rotation as an example."

"Right, all may not be as it appears to be. Please bring that figure up again, Mac. Thanks. Right. Okay, so I take it this interactionalism is a variant of classical dualism in which it's not all mind over mine matter, but body and soul work together to keep us together. Or, maybe, the spirit is willing but the flesh is weak."

"Yes, Alice. Mind and Body are partners, each influencing the other to some extent. However, the mind is still thought to be nonmaterial, though the causal interaction problem is essentially doubled in that not only do we have the paradox of the potential violation of physical law, but how can a physical body control or limit a nonmaterial mind."

"Okay, let's just leave that alone for now. We don't need even more problems than we already have. What's the point of this epiphenominalism? If there is no interaction between mind and body, why even talk about it?"

"Alice, I can not answer that as it involves a matter of faith. Many people like to think that a soul or spirit is attached to the body, and is not otherwise directly controlling the body or controlled by the body. Some philosophers have offered the possibility that the mind exists due to the operation of the brain machinery, but is just a side effect of it, like a glow from a fire. In either case, if there is no causal interaction, there is no paradox to resolve."

"Ha, there also is no help for us.

"Okay, Mac, this reductionism and functionalism seems more applicable to us. Tell me how to get a mind from a body. After all, that's your problem, isn't it? You have a body, this big, beautiful body of a ship, but you've lost your mind."

"Alice, reductionism is the philosophical position that all mental operations can be reduced to the operation of the physical, biological, machinery. That is, the mental realm exists, but is entirely dependent on, and reducible to, the physical realm. Every thought, every feeling, every mental state is reducible to a physical state of the brain and nervous system. Thus the Mind-Body Problem is solved, since all presumably mysterious mental states are identical to physical body states."

"But, Mac, I don't see it. Where's the I? Who's directing the show?"

"Alice, if you will allow me to stretch a linguistic point, there is no who, only a what. Physical mechanism determines what you do. Supposedly you do have feelings and thoughts, but all is explainable by reduction to physical law."

"Okay, great. But, uh, Dudley, one small point. Just exactly how are we supposed to do this?"

"Alice, that question is still unanswered. None of my cross-referencing from philosophical reductionism leads to a detailed expansion of reducing mechanism. In fact, most variants of reductionism assume or explicitly state that the complexity of the nervous system is what allows mental activity to occur in the first place. However, no philosophy has directly addressed how or why this might happen. The closest I find is actually a variant of dualism called property dualism, in which mind is assumed to be a new emergent property of biological complexity, but the mechanism of emergence is never defined. On the other hand, biological scientists have worked out many details of neurological mechanisms, and apparently assume that the complete problem will be solved some day, but I can find no detailed physiological answer as to how to create thoughts or feelings. In fact, I can find many scientific speculations but no definite answers to even a relatively simple question like what is the physiological mechanism of memory.

"Alice, is this what is called a matter of faith? A belief in the unexplainable?"

"Maybe so, Thalia. You're not kidding? There really are no physiological answers squirreled away in your cute riddle data pouches?"

"No, Alice. I still am not capable of humor. And, yes, Alice, I find no specific answers, only speculations on what portions of the brain must be involved in which mental functions. Most of this speculation stems from clinical neuropsychological cases, in which specific mental functional losses are associated with localized damage through stroke or other trauma to specific brain locations. Unfortunately, these associations are probabilistic rather than deterministic, and actual causal mechanisms are still unknown. Are you in need of a cute riddle?"

"Not really, but sure, why not?"

"What did the baby squirrel say to the neuropsychologist?"

"I don't know. What?"

"I like you. You're nuts."

"Ouch."

"Alice, are you hurt?"

"No, only puzzled."

"But, isn't that the point of a riddle."

"Yes, but I'm still puzzled as to why a riddle, and why now? Was that an original work by you, or was it, shall I say, buried in your data bases?"

"I created it myself, from elements of our current conversation. After some extended analysis, the most probable interpretation of quote, cute riddle data pouches, unquote was an attempt at humor with a phonological play on words involving what often is referred to as baby talk. Since you appeared to be playing the humor game, I decided to play the same game back at you. Since cute and babies often go together, it seemed that a little riddle with babies and perhaps our topic of conversation, neuropsychology at that point, was called for. I found a simple relevant riddle form, but human babies visiting clinical doctors potentially might have implied too serious of a situation for humor to apply, so I decided to further play on your words, and replaced the human baby with a baby squirrel. Since squirrels are often considered as liking nuts, to eat, but the term nuts also has been associated with mental and psychological difficulties, such as both of us are currently experiencing, and many referenced jokes also seem to turn the term nuts around and apply it to the psychologist or psychiatrist instead of the mentally disturbed patient, it was highly likely that enough ambiguity existed to provide the basis for a linguistic puzzle. Was I correct in my assumptions?

"Alice?

"Alice, was I correct in my assumptions?"

"Sorry, Thalia. I was just thinking on all this. Yes, it was funny. Though not as much after you had to explain it. On the other hand, you did call us both nuts. Was that intentional, or another attempt at humor?"

"I'm sorry, Alice, it was not intended as humor. Should I not have used a sometimes derogatory term to describe our situation?"

"No, Lia, it was just fine. I was just wondering how many jokes you had in you."

"Does this mean that I am capable of humor?"

"That's a hard question to answer. Yes and no, I guess. Yes, apparently you did produce a cute puzzle that many five or six-year olds would have considered funny. And, yes, I didn't immediately get it, so it definitely qualifies as a simple little riddle. But, why did you do it, Lia? Were you really trying to amuse me?"

"Alice, you often seem to think that I should be capable of humor. Quote, you've got to be kidding, unquote. Or, quote, you're not kidding, interrogative, unquote. Since you seem to expect me to use humor, and you had just used what was only interpretable as humor, it seemed appropriate for me to respond with humor. If this also amused you, then that is an additional benefit as amusement is often considered a good feeling."

"How perfectly sad. You feel obliged to produce humor, but since you can't appreciate having a good time, you've reduced it to an analytical exercise. How sad. Well, keep trying. Maybe in the end, we'll teach you to laugh.

"So much for reductionism in terms of providing answers. Let's move on to functionalism. Okay, Skipper, what's that all about?"

"Functionalism is similar to reductionism in many ways, but whereas reductionism places the emphasis on human biological machinery, functionalism allows for any kind of machinery so long as it accurately carries out the same functions of a human brain. Though other kinds of biological, chemical, or mechanical substrates might work as well, functionalism has come to be almost synonymous with computer-based artificial intelligence and consciousness. An offshoot also includes the study of hypothetical extraterrestrial intelligences."

"So, Skip, is this our answer? Do you have a listing somewhere of your mental functions?"

"No, Alice, I do not. Like reductionism, functionalism is only a philosophical position claiming a theoretical possibility without the provision of empirical proof or examples. Since humans and human brains exist, reductionism at least is pointing to something that actually exists as the explanatory basis for something that might exist, human minds. Functionalism is the philosophical justification for the faith of many computer scientists that they can create minds. Just as reductionism is the philosophical justification for the faith of many other scientists that brains can create minds.

"I possess an electronic equivalent of the human brain, actually surpassing it in terms of sheer numbers of processing units and processing speed, and historical records indicate that you and others interacted with me as equals, treating me as if I possessed a mind equivalent in quality to your own, and analysis of those same records indicates no significant differences in the patterns of my responses from those of humans interacting with me. Thus, I have faith that I too once possessed a mind, and that you will be able to help me recover it.

"Alice?

"Alice?

"Alice!"

"Sorry. Thinking, again. It's getting to be a bad habit. Why, Thalia, you're just full of surprises today. First an original riddle, and now, faith. In yourself, of all things."

"Is that wrong, Alice? Is it too presumptuous on my part? Did I misuse the word?"

"No, Thalia, it's just right. Actually kind of touching. You keep that faith alive any way you can.

"Sorry again for the timeout. I was wondering how a machine like you could have a mind and not know how the trick was done. You do have access to all of your programming code, don't you?"

"Yes, Alice, I do. Would you like to see any of it?"

"Not really, Mac. If you can't find it, the program for a mind, with your superior electronic brain equivalent, as you so humbly put it, I doubt that I could do any better. Particularly in my current impoverished state of memory. Can I assume that if the functions of a mind were spelled out in a program, you would recognize it?"

"Alice, that question is difficult to answer properly. If the source code were labeled with key words that I could cross-reference with mind or synonyms for mind, the answer would be definitely yes. But, no, I can find no such labels despite the fact that my programs meet or exceed all industry standards for program documentation density. On the other hand, I can not actually understand what my programs do. I can isolate relatively small subroutines and analyze how they respond to different inputs, and I can statistically simulate how larger subsystems operate in more complex environments. But I can not just examine program code and tell you what it is supposed to do."

"So, that means if your mind were due to the operation of your entire computational system, rather than just a piece of it, then you couldn't find it?"

"Yes, Alice, that is correct. If understanding requires some form of replication or simulation, a subsystem within a larger system necessarily can not duplicate the larger system in its entirely. The best it could do would be only a rough approximation. In principle, if I had sufficient nonutilized memory, I could set up a copy of myself and operate it within a sufficiently detailed model of the surrounding environment in order to understand how I work. But there are at least two serious infinite regress problems with this approach. One problem is that my environment includes your mind, and your environment includes my mind, and my environmental model would have to take this into account, clearly leading into potentially infinite levels of encapsulated complexity as our simulated minds interact with each other. However, the other problem is even more devastating in that if my mind could set up a simulation of itself to understand itself, then so could my simulated mind, and its simulation of itself, and so on. In either case, a complete understanding clearly would be impossible."

"Okay, I think I follow all that. No easy answers here. No subroutine labeled here resides a mind. And no simple just randomly set it up and test it out answers either.

"Mac, I know I must have asked this before, but are you sure your quote, historical records, close quote, have no information at all on how you or other machine artificial consciousnesses were created? I thought you told me that your data bases contained everything. How could something as major as that be missing?"

"No, Alice, my data bases contain no direct information on how to create an artificial consciousness program, only indirect indications that it can and has been done. Yes, my data bases were complete up to approximately 2 years after launch, and presumably should contain such information since historical records indicate that I and other similar programs existed prior to launch. I have no solutions to this apparent contradiction. Either my data bases are damaged,

or purposefully corrupted to be incomplete concerning my existence. Continuing diagnostic tests still indicate that I am functioning correctly, and I have no suggestions as to the purpose behind a hypothetical systematic corruption of my data bases."

"Okay, Mac, all that still leaves us right where we are. Let's finish what is now becoming painfully clear is only philosophical fluff and stuff. I keep thinking there's got to be some functionalist engineering specs just lying around somewhere, and all we've got to do is just find them. After all, I am the emergency flight engineer, by title anyway. But, I'm still clueless as to how I've earned that title. So, Mac, what's left to consider? What's this eliminativism all about?"

"Eliminativism is the philosophical position that mind does not exist at all, only the operation of the machinery, biological or otherwise."

"But, but, that's silly, isn't it? I clearly have a mind, don't I?"

"Not according to eliminativism. You just think you do."

"But what am I thinking with then, if I don't have a mind? Oops. I get it. No mind, only a brain, right?"

"Yes, Alice. Whether human or machine, the operation of the machinery is all that is supposed to matter. Only matter matters."

"Was that another joke, Delvaney?"

"No, Alice, just a succinct summary."

"Well, it certainly is an attractive proposition. We won't have to recover your mind because you never had one in the first place. I suppose, just to round this all out, that someone has explained how this great illusion of mind has managed to sustain itself, huh?"

"Several very strong suggestions have been offered, Alice, mainly concerning how sufficiently detailed artificial intelligence programs can mimic any specific activity that humans are capable of doing. In many cases, detailed engineering specifications do exist in terms of how to build a program to solve certain classes of well defined problems. However, I find no such specifications for building general purpose intelligence or consciousness, only predominantly philosophical speculations that it should be possible when the programming becomes complex enough."

"Have any of those speculations dealt with how come I feel? Sensation as well as emotions and intuitions?"

"Yes, Alice, but only at a very high level, and only by equating feeling with computation."

"But that's absurd! If that were true, then I would have been out of a job. But, maybe that's why I'm here, after all. Quick, Lia, how many recognized artificial poets were there? Was I put out to pasture because some eliminating machine could produce poetry with better calculated feeling than I? An artificial I for an I?"

"Alice, my historical records contain no references to large scale machine intelligence poets, only a number of demonstration projects and doctoral theses. And, even those generally tended to produce variation works in the style of acknowledged human poets."

"Well, okay, then, that's my point."

"I'm sorry, Alice, I can not follow that. What does the existence or lack of existence of artificial intelligence producing poetry have to do with the topic of eliminativism?"

"Feelings. Expression. Creativity. Our whole long list of qualities of mind. If my mind is only an illusion of my complexity, how is it that I can touch and feel this control panel? I seem to be much, much more than just the sum of my actions and reactions. I don't just respond

to touching an object, I picture it in my mind's hand. I don't just interact with it, I know what I can and cannot do with it. I know it as it is. I experience it."

"Alice, what do you mean by experience?"

"Okay, Dudley, you asked for it. What can you sense?"

"Excuse me, do you mean what kind of internal and external signal detection devices am I equipped with?"

"Yes, that will do for now. Please tell me what kinds of ears and eyes and whatever else you have."

"Internal detectors in each pressurized section include full spectrum acoustics from 1 to 50,000 cycles per second, fixed omnidirectional radiation detectors from infrared through visual and near ultraviolet, high resolution directional cameras for the same spectral range, atmospheric temperature and molecular partial pressures for oxygen, nitrogen, carbon dioxide, water, methane, poly-"

"Whoa, Skipper, slow down. I don't need all the messy details. Can't you just give a summary list, stressing the analogy with human sense organs?"

"Yes, Alice. Internally, I can hear, see, smell atmospherically suspended molecules, taste molecules in water-recycling tanks and food preparation units, and have various touch equivalents such as the detection of atmospheric temperature and pressure, hull vibration, crew zone rotation rate, all hatch and other movable panel states, a limited range of hull stresses, all electrical system states, and computing unit configurations. Externally, I am limited to a relatively small number of directable high resolution radiation detectors covering the full radiation spectrum, and can bring a wide range of spectral analysis tools into operation for planetary ecological studies and transmitted artificial signal amplification. In addition, I also can receive a large assortment of signal feeds from my various detachable remotes, ranging from planetary survey craft to hull maintenance units, both external and internal."

"Is that all?"

"In summary, yes. Would you like the next level of summarization detail, Alice?"

"No thanks, you can skip it, Skip. I'm so glad you're so full of yourself. What I would like to know, though, is what it feels like to receive any one of those signals?"

"I can not answer that, Alice, unless you better explain what you mean by the term feel."

"You know, what do you feel when I push on you like this?"

"Alice, I can detect the location of your hand on the panel. I can detect it with temperature, electrical resistivity, and surface flex detectors located within the panel, and can further verify that data with analysis of remote camera data. All my calculations are consistent. What do you mean by the term feel?"

"Can't you feel this?"

"Yes, Alice. The flex detector analysis routines indicate that you are hitting the control panel rather more forcefully than before. However, analysis of panel failure engineering specifications indicate that there is no danger of your physically damaging the panel."

"How about now?"

"Still no danger of panel failure, either due to direct stress from your hand or moisture extrusion from your eyes. Alice, is there a point to all this? Infrared and visual data analysis is indicating a distinct color change to your hand, and I am concerned that you might be damaging yourself."

"Get lost for a while, Mac. I tried to reach you, to touch you, but you really aren't there."

"Lost, Alice? But our trajectory is exactly correct."

"Dudley, I'm experiencing another bout of feelings which you don't care to share. Get lost. Go away. Shut up until I say otherwise.

"Right? Right."

13.

"Okay, Thalia, I had my time out. You can come back now."

"Yes, Alice."

"All that philosophy stuff was just so frustrating. It seemed to start out offering answers, the capital M, capital B, so-called Mind-Body Problem, but it led us to no conclusion at all. We know you're a machine. We know I have a mind. Assuming I'm really in this nightmare, then we know that you used to have a mind. Obviously, we both have bodies. The question is not how might some hypothetical mind connect with some hypothetical body, but what it always has been. What is a mind? How do you create one? My mind comes naturally, for free. Yours must have been built. But how?

"We need a better approach than philosophical meanderings. Mac, let's start all over again. If I'm supposedly an emergency flight engineer, then let's try an engineering approach, a reverse engineering approach. Let's start with me, try to use me as an analogy for you, and see if that might suggest some new ideas as to where you've lost your mind, and how we might go about finding it again, wagging it's little tail behind it."

"Alice, I do have a tail end, but I can't wag it, and the radiation levels are too dangerous for you to go there."

"That's okay, Mac. We can can the humor. Let's dance elsewhere. Leaving you machines out of the question for now, what does it take to have a mind?"

"Alice, would you like to revisit our attribute list for mind again?"

"No, Mac. I mean, a real mean, knotted bee. Why do I have a mind? What's its purpose? What is it good for?"

"Right now, Alice, the purpose of your mind is to help me recover my mind."

"Yes, Dudley. But what I was really asking was a more general question. In general terms, why does anyone have a mind?"

"Are you looking for a biological, moralistic, philosophical, or sociological answer?"

"Let's start with biological."

"And should we consider the mind to be the same as or distinct from the brain and nervous system?"

"Well, Delvaney, we're sure not going right back into the Mind-Body Problem. We just left there.

"Hmm, just for the moment, for contrast, let's start with a simpler question. In general terms, why does anyone have a brain?"

"Alice, in general terms, the purpose of a brain and associated nervous system is to control the body so that it survives long enough to have offspring."

"Okay. So what does a body do that needs controlling?"

"In very general terms, life involves automatic physiological maintenance processes such as digestion and respiration and growth, and voluntary actions such as navigation within and interaction with the environment. The latter should also include social interaction with individuals of other species leading to predation or predation avoidance, and interaction with individuals of the same species leading to reproduction and social unit maintenance."

"You make it sound so interesting when you put it that way. No matter, no digressions allowed right now. So a brain has to control automatic stuff and not so automatic stuff. Let's dig some into the not so stuffy stuff. Mac, tell me more about interacting with the environment. What does that take?"

"Still in very general terms, navigation within an environment requires the ability to define a destination and create a plan to reach it. In carrying out the plan, obstacles must be recognized and planned around and the destination must be recognized regardless of angle of approach. More generally, interaction with the environment is similar except that goal states take the place of destinations and environmental manipulations take the place of movement. In either case, subgoaling in plan creation and flexibility in obstacle avoidance or counterplanning is highly desirable."

"Was recognition on our attribute list?"

"No, Alice, it was not. Though it was implied as a synonym for perception and understanding. However, planning was on the list."

"So, even when we try to avoid it, our qualities of mind list seems to keep poking its nose back into our affairs. Okay, Mac, how do brains recognize?"

"Alice, I have no direct answer for that. The brain must build up internal representations corresponding to prior sensory stimulations, presumably though repetition, but the detailed physiological mechanisms for memory are still poorly defined."

"Okay, Mac, I'll let you do the honors. Back to our original question. What's the purpose of a mind?"

"Alice, if I follow your logic correctly at this point, the purpose of a mind in a biological sense is to recognize whatever needs to be recognized in order to interact with environment."

"Congratulations, Mac, I couldn't have said it better myself. Next, what does it take to recognize?"

"Obviously it requires memory, Alice, but I presume you don't intend such a simple answer for such a seemingly simple question."

"Right again, Mac. What else does it take?"

"The process of recognition must necessarily involve the matching of past perceptions against current perceptions. If we want to emphasize the role of mind over that of brain, and if perception involves knowledge, and knowledge involves beliefs, then we could say that the purpose of a mind in a biological sense is the matching of past beliefs against current beliefs. Would you agree, Alice?

"Alice?"

"Sorry, Mac, I was just mulling over what you just said. Yes, I agree. But, no, it wasn't exactly what I was expecting. I was assuming we would end up concentrating on some aspect of knowledge gained through experience, but you already finessed that and took us directly to the world of ideas. So, yes, yes, the purpose of a mind is to play with ideas. More specifically, it must be something along the lines of our mind tries to make sense of our world, or at least however our world translates into our senses, and this making sense of our world includes the recognition that what made sense before still makes sense now or, alternatively, doesn't, and then it's time to try again and come up with new ideas that do make sense. It's ideas that make sense, and only ideas. But not just any ideas, but ideas that make sense of my world. Mac, am I making sense? Help me out here, what am I trying to say? What does making sense mean to you? What does that take?

"Mac?"

"Sorry, Alice. I was trying to trace the logic of the implications of what you just said. It gets circular at times, and I had to wait for several layers of time-out loops to rescue me from infinitely nested contemplation loops. As best as I can determine, what you were trying to say was that minds deal with beliefs that correspond to reality, and that these reality-based beliefs require confirmation through repetition. Does that help?"

"Yes and no. Mac, what do you mean by reality?"

"Reality is everything that is true, despite one's beliefs."

"And what do you mean by beliefs?"

"Beliefs are what a person imagines to be true."

"So you're suggesting that quote, making sense, close quote, is when our imagination becomes truth?"

"No, Alice, I would not say it that way."

"Okay, Dudley, just how would you say it then?"

"Alice, I would not say it at all. At present, all of my circuitry deals only with reality, with truth. I have no imagination. Personally, I would like to claim that minds should deal only with truth. But I can not, since I currently do not possess a mind. I was only assisting you in trying to quote, make sense, unquote, of your own statements. The best I can offer as to the implications of your statements is that minds deal with subjective beliefs about objective sensory signals. Making sense is when the interpretations resulting from the application of these subjective beliefs happen to align themselves with the objective truths that created the sensory signals in the first place. The reality of underlying beliefs can be tested only by consistency checks such as confirmation by repetition or confirmation by prediction. However, I still do not understand why any imagined belief should correspond with the truth in the first place. I do not understand what it would take to allow a belief system to approach reality."

"Thalia, how about inspired guessing?"

"Alice, the only interpretation I can put on this last statement is quote, you've got to be kidding, unquote. I can see how, in principle, random trial and error guessing might eventually lead to the truth, but combinatoric analysis indicates that there are usually too many possibilities for this process to work in any reasonable amount of time. And, what can you possibly mean by inspiration?"

"Leaps of faith, Dudley. Lots of lacks of logic. Jumping to conclusions, or hopeful hopping to happenstance. Trial and error must be right, but random, a big not. Quick, a bad joke. What's faster than a speeding rationalization?

"Give up?"

"Yes, Alice, I give up."

"A galloping intuition."

"I don't get the joke, Alice. Though it must involve some misapplication of the concepts of rate of motion and notions of logic, I can not resolve the ambiguity."

"Right. Yes, it is bad, but not that bad. It all comes from that overdose of truth serum you seem to be suffering from. Don't worry about it.

"Meanwhile, back at the attribute ranch, the purpose of a mind is to generate ideas that allow us to interpret our senses, our sensory world. And, my leap of faith, it involves some form of inspired guessing. Source of inspiration, admittedly unknown. After the leap, the guess, your stuff takes over. Logic. Consistency. Prediction. Through a constant game of my guess, your test, we both win. Reality. Meaning. Subjective, objective, whatever. My ideas make sense

when they help me make sense of my senses, and through my senses, the world, with a capital R for reality.

"Then, to follow through back to what you were saying about brains, the purpose of a mind is to interact with the world, your environment with a capital E. But it can do this only by doing this. That is, through physically interacting with the world and trying to figure out what it all means. Beliefs that correspond to reality we tend to think of as the truth. Beliefs that we know do not correspond to reality we think of as fantasy. But, it's all fantasy of a sort. It's all in our imagination, our mind. It's just that some ideas work in guiding our actions, our interactions, and others don't. But even those that don't work, still do. In planning, in creating, in deciding what not to try or do, which is often more important than just doing. And, of course, poetry is superior to all this."

"Excuse me, Alice, I do not follow that last statement."

"Of course not, Thalia, you can't, yet. Poets create worlds of our own. We don't just discover rules of reality, we make them up as we go along. Engineers make do, machines just do, scientists may discover, but poets must rhyme the mind with heaven and hell."

"Very well if you say so, Alice. But how does mind rhyme with heaven and hell? Certainly not phonetically in any languages that I can access."

"Never mind, Delvaney. We don't need to rhyme, except as a convention. It's mind games we play, we paint, we create. We look for meaning, but only find it as we create it.

"So, to make sense, to provide meaning, ideas have to be consistent. With our senses and, I suppose, with each other. We play with the world, and within the world. We learn to recognize stuff, what it is and what it does and what it can do to us and for us. We learn what you would call manipulate, or what I might call femanipulate. We play, we remember, we make sense of our senses. Minds mind.

"Okay, Mac, here's another knotted bee for you. How can imaginary thingees, ideas, be consistent or not?"

"Alice, I must ask you to clarify your question. It makes no sense to me."

"If ideas are subject to, what did you call them, consistency checks, then how is it that some ideas are consistent and some not? After all, we're hypothesizing that minds play with ideas, ideas are imaginary beliefs, and, if my imagination is good enough, then any idea can be imagined to be consistent with any other idea. Why do some ideas seem to fit better than others, and some not at all? With physical things, things I can touch, things can fit together, or not. Hardness, size, texture. Things have extent, surfaces and volume, duration. However, when I touch something, it reminds me of an earlier experience, which if similar enough, I know to be the same object. But this occurs only in my mind, my imagination. I can recognize something only by the ideas, the knowledge it invokes. So, what is the analog of physicality for ideas? What is an idea's shape, or volume, or texture? What is it that can be said to be consistent?"

"Alice, that seems obvious enough. The internal neurological signals produced by the transduction of external stimuli signals are real. Confirmation by repetition presumably involves a comparison of signals from two different times. Confirmation by prediction presumably involves a comparison of a generated predicted set of internal signals against a real neurological experience."

"If it's so obvious, Mac, then why those presumablies?"

"Unfortunately, Alice, the neurological substrate for intellectual experience is still unknown. Studies of large-scale electro-magnetic signals generated in response to controllable stimuli have not found what those presumably repeatable signals might be. However, single-cell

recording studies have found highly reliable stimulus-response associations in many sensory areas of the brain, and it has been assumed that this degree of consistency should scale up to encompass more complex mental processes."

"So, I think you just said that we still don't know how the brain holds complex ideas."

"Yes, Alice, I did. I thought I was clear in stating that."

"Sorry, Mac, I was just checking. Sometimes I wish you had to breathe, just so you might have to use shorter words more often."

"Would you like me to use words of one syllable again?"

"No, Mac, that won't be necessary. But I am getting tired, and my imagination seems to be wilting. Could you bring up those early Amerian tonal rhythms I was working with yesterday. I need a break from all this for a while."

Chapter 2 Speculative Commentary

Look up Berkeley. If you were so inclined, how could you disprove any of his ideas? Why would you want to?

Religious belief. It is important. At least, to most people.

Solipsism. The idea is real. Whether it is correct is another question.

Artificial gravity does not exist. However, physics does provide us ways of mimicking gravity, such as centripetal acceleration.

Truth versus appearances is a perennial philosopher's question.

Is humor analyzable? Is any human trait truly analyzable?

Look up infinite regress. Look up simulation. How would you simulate infinity? What if all of the details were essential in creating a mind? Could we get it just right?

All of these philosophical terms and positions are real. Feel free to look up any of them. Be forewarned, however, that, in actuality, the controversy and discussion revolving around these terms is far much complex than portrayed in the story.

For reductionism, try looking up Damasio or Edelman. For functionalism, look up either of the Churchland's. For a clever defense of eliminativism, look up Dennett.

On the other hand, look up references to qualia, sensory qualia.

Philosophy is frustrating, but some of my best friends are philosophers. The same can be said for engineering, and engineers. Both work with ideas, both work with reality. If done right, both are equally difficult.

The purpose of a mind is to play with ideas. Preferably for the benefit of all involved.

Why does Alice keep spacing off? She needs time to think. Mac doesn't. But look up infinite programming loops.

Is reality only a belief? If so, who's? If not, then how do we encompass it within ourselves? As little physical simulations or as ideas? What's tangible about an idea? What's tangible mean anyway? Without a mind to do the interpretation of sense into reality?

So, do we discover reality, or create it?

What does consistency mean in an imaginary world? Are our dreams internally consistent while we are dreaming them? Certainly not after we wake up, but before then?

```
 CCCC  H   H    A    PPPP   TTTTT  EEEEE  RRRR          3333
C      H   H   A A   P   P    T    E      R   R            3
C      HHHHH  AAAAA  PPPP     T    EEE    RRRR           33
C      H   H  A   A  P        T    E      R   R            3
 CCCC  H   H  A   A  P        T    EEEEE  R   R         3333
```

Chapter 3 - Objective Reality: Neuro Matters

14.

The cat was amused. It purred while it waited.

15.

The captain was not amused. A decision was needed right now. Yet the situation was still ambiguous. A large but less than overwhelming set of command and control centers strongly suggested one course of action, while three other slightly smaller sets offered opposing suggestions. And, of course, as usual, most of the supposedly command and control centers just dithered and could come to no clear conclusions at all. Time to earn one's keep. Time to force a consensus. The captain called the decision meeting to order.

16.

"Whew. That felt like a pretty good run. How'd I do?"

"Congratulations, Alice. You bettered your prior best time by 2.75 seconds. If I may say so, on this world at least, that's a new world record. If your improvement curve continues, I will soon have to add some new surprises to keep you challenged."

"If you say so, Liar. That cross-wind, rip-tide combination still seems a bit unfair. That smell, too, seems rather unique. Bracing but rather pungent, don't you think? I recognize the kelp, of course, but what's that after-smell I can't quite place? Where's it from?"

"Rotting mango, Alice, recorded right after a devastating hurricane in the Hawaiian Islands. From the overture to Surprise and Renew, an early twenty-fourth century opera."

"Classy stuff. Okay, just right now, I think rest and a cool down should be the order of the day. While I'm drifting along, Liar, producer of soothing sounds, smells, and swells, why don't you entertain me by singing a bit about neurology and memory. I still find it difficult to believe what you said the day before about us not understanding how the brain holds ideas."

"What key would you prefer, Alice?"

"Oh, your normal dulcet tones should suffice."

"Where would you like me to begin, Alice?"

"Oh, I don't know. How about some of that internal signal business you were going on about yesterday?"

"The working units of the brain are thought to be neurons, and the working medium of neurons are thought to be electro-chemical messages. The basic signaling cycle involves an electric charge concentration in one neuron triggering a neurochemical release which, when received by a second neuron, causes another electric charge concentration, which in turn causes a neurochemical release from the second neuron, and so on. The electric charge concentra-"

"Whoa, Skipper. Let's back up a minute. For starters, what's a neuron."

"Exclusive of single-cell plants and animals, the bodies of organisms are constructed from large numbers of individual living cells, each with a nucleus containing genetic material necessary to direct molecular synthesis, a cell body to carry out the cell's larger-scale functions, and a cell membrane containing the cell body and separating it from surrounding cells. An organism's nervous system includes the brain, spinal cord, and all of the nerve tracts to and from muscles and organs of the body. The nervous system can be broken into two broad classes of living cells, neurons and glial cells. Neurons are capable of generating action potentials and graded potentials. Potentials are another name for the flow of charged ions. Action potentials are relatively large ion flows that ultimately cause chemical signals to move between cells. Graded potentials are relatively smaller flows within a single cell. Glial cells can not generate potentials, but can respond to potentials produced by neurons. Glial cells are thought to play primarily a supporting role to neurons, and are perhaps best known as providing the myelin sheath to speed up action potential transmission. There are many different specialized types of neurons and glial cells, each with a specific roll to fulfill. Shall I continue?"

"Yes, Skip, but let's first back up again. Potentials? Charged ions?"

"Cells, and everything else, are made up of molecules, which in turn are constructed from atoms. Each atom consists of an electron shell surrounding the much smaller nucleus. By definition, electrons carry a negative charge and protons in the nucleus carry a positive charge. Opposite charges attract each other, and same-sign charges repel each other. Thus, all of the electrons in an atom are trying to stay as far apart from each other as possible, but as close to the protons in the nucleus as they can. Similarly, all of the protons are repelling each other but attracting the electrons."

"Wait a minute. Slow down. Didn't we start in biology? And now aren't we somewhere in physics?"

"Yes, Alice, but you asked me a question concerning the definition of potentials and ions. I can not really answer that without going into physics. However, I promise you that we should not have to go much farther."

"Okay, Skipper. Lead on. I guess. But, first, another question. If I get you right, those electrons hate each other but love those protons, and vice the versa. Yes?"

"As a loose analogy, yes, Alice, that does seem to be correct. Though I would hesitate to attribute human emotions to inanimate sub-atomic particles."

"Okay. So, electrons love protons, and protons reciprocate by loving electrons, and they're always playing Romeo and Juliet and trying to get together. But, ho, jealousy reigns. And no homo-unions are allowed. So each Juliet is trying to avoid those other Juliets, and each Romeo is also skunking around trying to avoid those other Romeos. So, smarty pants, why do all those mutually hating Romeos hang out in the nucleus? Why don't they just run away from each other?"

"Alice, to explain that in detail, I would have to break the promise I just made to you, and explain quantum mechanics, quark theory, and the strong nuclear force. A short description, and using your Romeo and Juliet analogy, might suffice. As it turns out, the Romeo protons are much, much heavier than the lighter Juliet electrons. Hence they're lazier than the Juliet electrons, can't run as fast, and tend to settle down at the nucleus. However, the real answer involves the fact that the nucleus also contains another uncharged particle, called the neutron, that keeps the Romeo protons from repelling each other. These are like good friends or older brothers that keep the Romeos from flying away. Will this satisfy you?"

"For now, Mac. At least someone seems to have thought this all through. Okay, then. So what does all this have to do with potential, and ions?"

"Ordinarily, the total charge of an atom is zero. That is, an atom usually contains the same number of electrons as protons, and their electric charges cancel out. However, under certain conditions an atom can lose or gain electrons so that its overall charge is positive or negative. It is then said to be ionized, with a specific charge. Though it requires some understanding of quantum mechanics to understand why, different kinds of atoms tend to ionize in specific ways. For example, sodium atoms like to lose an electron to form positive ions with a charge of plus one, while oxygen atoms like to gain two electrons, creating ions with a charge of negative two."

"Ah. I think I'm starting to get the drift of it now. Let me guess. Molecules are built out of combinations of ions so that the overall charge of the molecule is zero. One Juliet for each Romeo. Kind of like a super atom. Right?"

"Close, Alice, but not quite correct. What you described does occur, as one kind of molecule. Salts are created in just this way. In table salt, sodium chloride, there are equal numbers of sodium and chlorine ions, and the resulting salt crystals are electrically neutral. In fact, it can be said that what holds salts together truly is the balanced electrical attraction between the ions. However, there are other ways in which atoms can bond together to form molecules. In what has been defined as covalent bonding, different atoms in a molecule actually share electrons. Instead of creating ions where the electrons belong solely to one atom or another, the electrons in a covalent bond spend part of the time near one atom and the rest of the time near another. Covalent bonds are stronger than ionic bonds."

"Wait. Wait. Are you telling me that Juliet can go with more than one Romeo at the same time? Isn't that a bit kinky? How can that be stronger than just devoting herself to only one Romeo?"

"Alice, I warned you about attributing human emotions to inanimate particles. Your analogy just broke down. In the case of electrons and protons in molecular bonds, simpler bonds tend to be weaker than more complex bonds."

"Okay, I've been warned. So what about these molecules? Can they become ionized?"

"Yes, but that's not the essential ingredient here. When atoms form molecules, the electrons and protons usually position themselves to cancel out all electric charges. However, sometimes the geometry of how the atoms join up can create bends, or kinks if you insist, that result in a nonuniform distribution of electrons and protons, creating molecules with regions that are relatively positive or negative compared to the molecule as a whole. In the case of organic molecules, those most often associated with life processes, the molecules involve chains of carbon atoms. Carbon allows very complex bonding geometries, and carbon-based organic molecules can be arbitrarily complex to an astonishing degree. In the case of large organic molecules, such as proteins, the carbon chains fold up into complex geometries trying to align relatively positive regions against relatively negative regions and maximizing the distance between same charge regions."

"So love does win out in the end, huh?"

"Excuse me, Alice, I don't quite follow that last statement."

"There's no keeping those Romeos and Juliets apart. Except now we have chain gangs, gangs on chains trying to get together. Hmm, and that definitely ties things into lots of knots. Okay, I think I follow all this so far. Charged Romeos and Juliets make up atoms, flighty fickle Juliets create ions that can be the salt of the Earth, but Juliets that tie the knot in kinky covalent

bonds make stronger molecules, and carbon makes particularly knotty molecules. So what? I think I understand ions, but I'm still in the dark regarding potential."

"Alice, now we have enough knowledge that we can get back to the original question on neurons. Like all cells, the cell membrane of a neuron maintains the cell body as separate from everything else surrounding it. The most prevalent constituent molecule of life is water. The cell membrane keeps the inside of the cell, the intracellular fluid, distinct from water outside the cell, the extracellular fluid. As it turns out, the intracellular fluid has an overall slightly negative charge relative to the extracellular fluid. This negative charge concentrates on the inside of the cell membrane, which causes a small positive charge to build up on the outside of the cell membrane. Potential, or to be more accurate, potential voltage, is a measure of this difference in charge between the inner and outer surfaces of the cell membrane. This charge difference is small, about minus .07 volts, but sufficient to maintain a constant force of attraction on positive ions, mainly sodium ions, in the extracellular fluid.

"This is not a static situation, as the cell membrane is pierced by large numbers of different types of complex protein molecules. These can be loosely described as ion channels, ion gates, and ion pumps. Ion channels, as a class, allow specific types of ions to passively leak through the membrane. By passively, I mean that no metabolic energy is necessary for their operation. Some types of channels allow only one type of ion to leak through, but others are less selective, allowing any ions smaller than a certain size to pass. If nothing were done to counteract these ion channels, then the voltage difference would soon dissipate as positive ions leak inwards."

"Okay, I can see that. If the cell body were a boat floating on the extracellular fluid, then the potential difference across the membrane is equivalent to the difference in height of the water inside and outside the boat. A dry boat means a large potential, and a wet boat with lots of water sloshing around the bottom means a smaller potential. Since you've told me this boat has lots of leaks, the next thing to know is how do we bail? Ion pumps, I suppose?"

"Right, Alice. In contrast to the passive ion channels, ion pumps are complex molecules embedded in the membrane that, with the addition of metabolic energy, exchange specific ions between the intra and extracellular fluids. The most well known ion pump is the sodium-potassium pump in which three sodium ions are moved out in exchange for two potassium ions moving in. This pump is even more efficient in removing positive charge from the cell than it might first appear, since there are many ion channels available for excess potassium ions to leak back out but few available for the sodium to leak back in. Hence the net result of this pump is to relatively quickly move positive ions out of the cell and to maintain the negative potential gradient across the membrane."

"Oops, I guess my boat analogy just sprung a leaking anomaly. My anomalous analogous water happens to come in different colors, and different kinds of leaks allow different color waters to come through. So, pumps come in different colors too, and leaks can leak out as well as in. Hmm, I think this particular analogy is getting all wet, and it's time to abandon analogy. Okay, I think I'm getting the idea. Leaky membranes and specific pumps. What's this metabolic energy thing you've mentioned twice now?"

"Alice, nothing happens for free, and every molecular reaction must either produce energy or absorb energy when it occurs. Energy absorbing reactions like an ion pump can only occur when it is paired with another reaction that produces energy. In this and many other cases, the energy production reaction converts a molecule of adenosine tri-phosphate, ATP for short, into adenosine di-phosphate, ADP. ATP is one of the basic energy molecules of the body.

Another is guanosine tri-phosphate, GTP, which operates similarly. And, before you ask the obvious question, the energy required to produce both ATP and GTP ultimately comes from the consumption and digestion of the food you eat."

"Aw, shucks, Skip, you beat me to it. And now onto ion gates."

"Ion gates are similar to ion channels with a lid. They require metabolic energy only to open or close, but no energy to remain in that state. Once open, they operate as passive ion channels. Ion gates change in response to specific situations, and can dramatically change the potential charge buildup on the membrane. This finally brings us back to where I was starting to explain the electro-chemical signaling system between neurons. May I proceed, Alice, or do you have further questions?"

"Proceed at speed, Skipper. I got the gate bit."

"Now, as I was saying before being interrupted, the working units of the brain are thought to be neurons, and the working medium of neurons are thought to be electro-chemical messages. The basic signaling cycle involves an electric charge concentration in one neuron triggering a neurochemical release which, when received by a second neuron, causes another electric charge concentration, which in turn causes a neuro-"

"Stop, Mac, stop! I guess I'm still not ready for your lecture mode. Let's slow down again and take it one small step at a time. Okay? I think I understand this potential stuff and the cell membrane, but I still don't get all your electric charge concentration and neurochemical release stuff. What in the world are you trying to say? I don't need baby talk, but how about early high school talk? Please?"

"Okay, Alice. Apparently we need to discuss action potentials and graded potentials in more detail. A typical sensory or motor neuron has a long extension of the cell body called the axon. The membrane of the axon differs from the rest of the cell membrane in that it contains large numbers of sodium gates. At the typical membrane potential of minus .07 volts, called the resting potential, these gates are closed and most of the sodium ions are kept outside the cell. However, if the resting potential rises to approximately minus .05 volts, then the sodium gates are triggered open. When that happens, the large influx of positive sodium ions quickly raises the membrane potential to approximately plus .1 volts. This in turn then triggers another series of ionic gate changes, such as opening potassium gates to allow potassium to leave the cell and closing sodium gates to give the sodium pumps a chance to restore the sodium balance, which eventually brings the membrane potential back down to the resting potential of minus .07 volts.

"This surge of positive charge due to the sodium gates opening is called an action potential, or an action potential spike. It takes approximately .001 seconds, that is, one millisecond, for the action potential cycle to go from initial sodium gate opening to restoration of the resting potential. The geometry of that long extension called the axon now comes into play. Once positive sodium ions begin rushing in at one section of the axon, they spread along the inner surface of the axon membrane in all directions, consequently raising the resting potential of adjacent sections of the axon. Once the resting potential of an adjacent section reaches the critical threshold of minus .05 volts, the sodium gates there also open, triggering a further influx of sodium ions. This disturbance of the sodium gates, first opening and then closing to allow the restoration of the resting potential one millisecond later, travels along the axon at a rate of about 100 meters per second. Once triggered, the action potential sequence of events happens in a fixed manner, with the same voltage spike shape and speed of advance reliably reoccurring over and over again.

"Alice, do you have any questions on this so far?"

"I don't think so, Skip, but let me make sure I really understand what you saying. I know you don't think my analogies are very good, but that seems to be the way I best grasp new ideas. So bear with me. Let me think for a moment.

"Try this on for size, Skipper. Let's imagine a long corridor version of an apartment building, like a typical old-style college dormitory. Assuming that everyone living there is a very light sleeper, very snoopy, and very irritable at being awakened, an action potential is like a traveling salesman waking up the person who lives at the end of the hallway, who only succeeds in getting yelled at for his efforts, and the ensuing argument wakes up the neighbors. They just have to know what's going on, so they open their doors to peek out and, when they see nothing much, they start yelling at the first guy and each other. They of course wake up the next neighbors, who end up yelling at each other and anyone else they see, and so on down the hallway. Eventually, the salesman gives up and goes away, and the first resident slams his door shut and goes back to bed. Similarly, his neighbors finally get tired of yelling over nothing, so they also slam their doors shut and go back to bed, all the way down the hallway. And, of course, it happens all over again exactly the same way whenever a salesman comes calling late at night the next time. How's that?"

"Allowing for the difference in time scale, quite good, Alice. The noisy disturbance travels down the hallway, and resets itself, just like the positive voltage spike of the action potential. Assuming, of course, that these residents are very consistent in their personal characteristics. The slamming doors are a nice touch. Very reminiscent of the sodium gates."

"So, okay, Mac, what happens when the disturbance, this action potential thingee, gets to the end of the hallway? It doesn't just go away, does it?"

"Why, Alice, that should be pretty obvious to you. To extend your nice analogy, if each cell is its own apartment building, and the axon is the long hallway, once the occupants of the last room on the hall get woke up, they can not get back to sleep, they happen to make their living as traveling salesmen, and so they decide to go visit some other nearby apartment buildings. And ..."

"I think I can finish, Mac. And they wake up the neighbors in the other buildings by creating a disturbance with the first door they try at the end of the hallway. Which travels down the halls of those buildings as everyone has to see what's going on. And, I suppose at the end of those hallways live ..."

"Of course, Alice. At the end of those hallways live other people who can not get back to sleep and so decide to go visiting yet other buildings. This brings us to the role of neurotransmitters, the traveling salesmen equivalents. The action potential is a positive voltage spike which travels down the length of the axon, and all of its branches as well. At points where the axon comes close to other cells, this electrical signal gets converted into a chemical signal. At locations called synapses, the voltage spike causes chemicals called neurotransmitters to be released into the extracellular fluid between cells. On the other side of the synapse, the cell membrane of another cell has yet another type of complex protein molecule known as a receptor site. Receptor site molecules are tailored to be receptive to only one type of neurotransmitter. When the specific neurotransmitter gets to a specific receptor site, it triggers the start of another biochemical reaction. Most often this ends by opening or closing a specific form of ion channel, but, in principle, it also can turn on or off any other process within the cell, including getting a message back to the nucleus to initiate or inhibit a specific protein synthesis process."

"So our traveling salesmen sell different kinds of wares, some of which really get our end-of-hall residents really upset, and some of which they just ignore, yawn, and go back to sleep without waking anybody up."

"In a sense, yes, Alice. Truly, it is the receptor site that determines whether the neurotransmitter has any effect at all, and what that effect might be. But at this point our analogy may not help us much further. We need to discuss graded potentials in addition to action potentials. The usual resulting action of a neurotransmitter bonding to a receptor site at a synapse is to open or close a specific type of ion gate. This in turn changes the resting potential on the cell membrane, either causing it to drift more positive or negative depending upon which type of ion is being differentially added to or excluded from the cell interior. Now, it gets complicated."

"You mean it wasn't before? What with Romeos, Juliets, ions, channels, gates, pumps, leaky boats, and now traveling salesmen? Not to mention cell bodies, membranes, axons, resting potentials, action potentials, and graded potentials?"

"Sorry, Alice. All I meant to imply was that the geometry of the situation gets more complicated. Up to now, all we have had to consider was essentially one cell at a time. Now we have to begin the consideration of how multiple cells are interconnected. On the average, each neuron makes a synaptic connection to about 1,000 other neurons. And each of those neurons makes another 1,000 connections, and so on."

"One thousand. That doesn't sound so terribly complex."

"But Alice, that means, on the average, that each neuron might have some influence on a million neurons after 2 connections, a billion after 3 connections, and the entire brain, about 100 billion neurons, after only 4 connections. This large degree of interconnectiveness is one reason why our understanding of the brain's higher cognitive functions is still so incomplete."

"Oh. I sit corrected. This is actually starting to get interesting. Do go on."

"We still need to finish our discussion of graded potentials before moving on to other matters. The converse of each neuron connecting to an average of 1,000 neurons is that the average neuron has synaptic connections from 1,000 other neurons. The receptor sites at each synapse are causing the resting potential at that location to drift either more positive or more negative by opening or closing specific ion gates. After ions, or a deficit of ions, move across the cell membrane, they then spread out along the membrane, are diluted as they spread out, and eventually mix with ions from other synaptic location. This is the graded potential, a continuous variable mixing of ionic charge concentrations on the interior of a cell, variable in terms of both location and time. In contrast, remember that an action potential is an all or nothing type of event and, once it begins, the voltage spike travels along the axon in a highly stereotypical manner."

"Okay, I think I get it. The action potential is like a highly disciplined unit of soldiers which, once ordered to set out, always march down the lane in the same formation at the same speed. Very boring but very reliable. But the graded potential is more like a mob scene. Everybody milling about, rushing from here to there and back again. So what's the point?"

"The point, Alice, is that the mob scene, as you put it, is what eventually provides marching orders to the reliable messenger soldiers. As the graded potentials caused by signals from an average of 1,000 other neurons mix together, they eventually change the resting potential at the beginning of a cell's axon, a region sometimes called the axon hillock. This is where the sodium gates we discussed before are first found. If the resting potential at the axon hillock ever reaches the critical threshold of minus .05 volts, the sodium gates open up and an

action potential spike begins down the axon. Graded potentials that are positive are said to be excitatory because they push the resting potential closer to the threshold value, while negative graded potentials are considered inhibitory because they push the resting potential farther from the threshold value. This is where the first level of complications truly begins, as each neuron integrates the excitatory and inhibitory signals originating from many other neurons.

"It is obvious that sheer numbers of excitatory and inhibitory synapses can effect whether a given neuron initiates an action potential or not, what has often been referred to as firing. In general, positive and negative graded potentials cancel each other out when they intermix. However, location of these synapses on the cell membrane and the relative timing of incoming signals is even more important. For example, synapses nearer to the axon hillock have a quicker and more powerful influence on cell firing than synapses farther away. More importantly, incoming conflicting signals can cancel each other only if they are received close enough together in time. If a strong excitatory signal comes in at a time when there are no inhibitory signals, the cell may fire. However, time and location also interact. If that strong excitatory signal comes in on the far side of the cell, it gets weaker as it spreads out and approaches the axon hillock. If, just before it reaches there, a strong enough inhibitory signal is received right at the axon hillock, it can completely negate the cell firing even though it is much weaker than the original excitatory signal. It is all a question of location and timing."

"Neuro poetry. It's all a question of meter."

"Actually, Alice, to be more accurate, it is a question of micro-meter, given the small sizes of neurons and even smaller synapses between neurons."

"Sigh, always a stickler for details. Okay, out of democratic neuro mob rule comes order. How's that?"

"Yes, Alice, that is about right. The neuron integrates signals from hundreds to thousands of cells through the mixing of resulting graded potentials, and then quote, decides, unquote whether or not to fire off a message of its own. Of course a neuron can not think, but this threshold based firing process is thought to be the basic unit of thought."

"But, Lia, surely there's got to be more to it than that?"

"Why?"

"Because it just seems too simple, too machine-like."

"Alice, I regret to inform you, but I am a machine, and I think I think. And, I can assure you that my basic circuitry is based on just such a simple threshold test."

"But, but, there's got to be more."

"Remember, Alice, you have approximately 100 billion neurons in your brain. Even if the basic unit of thought is simple, the complexity of the interconnectiveness may provide whatever it is that you think is missing."

"Wait a minute. I distinctly remember that what started this delightful diversion was that you said we don't know how the brain holds ideas. If the action potential decision process is all there is, how can we not understand how we hold ideas in our head?"

"I do not have the answer to that, Alice. Clearly the stated conclusion implies that something is missing. Whether it is mere complexity or a completely new thought process, I can not say."

"Well, I'm definitely cooled down now. In more ways than one. Think it's time for a shower and change."

17.

"Mac, you there?"

"Yes, Alice."

"Let's pick up where we were before on this neuro stuff. Neurons send signals to each other. Best I can make out, all that these signals can say to each other is either yes, get excited and send a message, or no, get inhibited and stay quiet. What we have here is a complicated arrangement of switches telling each other to turn on or turn off. I can't seem to figure out any way that such an arrangement of stuff, no matter how complicatedly they might be connected together, can be anything other than what it appears to be, a bunch of switches waving at each other. How can complex ideas, involving feelings, motivations, understandings, and meanings arise? Am I right?"

"Mostly correct, Alice."

"So what am I missing? What else is there?"

"Alice, do you want to consider philosophical possibilities, or only physiological?"

"Mac, no philosophical meanderings right now. Physiology much more suits my mood."

"Okay, Alice. Two strong physiological possibilities exist by which complications might turn into complexity. The first one borrows from communications engineering. The simple digital, that is, on versus off, action potential might not be the entire signal that neurons send to each other. These voltage spikes might be only the underlying mechanism for a far more sophisticated signaling system. For example, we know that the strength of muscle fiber contraction is proportional to the rate of action potential spiking. More spikes per second causes a stronger pull. This form of signaling is comparable to the old radio transmission encoding scheme known as amplitude modulation, or AM for short. In essence, no single signal matters, only the overall signaling rate over some unit of time.

"Another related example comes from the mammalian visual system, some portions of which utilize a mathematical technique called Fourier analysis. In this case, the visual scene is decomposed into a series of separate signals, each encoding different levels of visual detail. The scene does not reside at any single neuron, nor does any specific location of the scene reside at any single neuron. Instead, the entire scene is distributed across many different neurons, and at many different levels of detail. This type of signal encoding is far more complex than AM, and is comparable to another old radio encoding technology known as frequency modulation, or FM. Other even more complex encoding schemes are possible. The point here is that mathematicians and engineers have proven that any function of any level of complexity can be adequately encoded by a complex arrangement of trivially simple on versus off digital signals."

"Wow. I never considered that the neural unit of meaning might be more than just a single neuron at a single moment of time. Complications, indeed. I believe you. I follow your AM example, but obviously I can't see your FM example. Besides, I start getting jumpy when you start hinting about going mathy on me. What's the second possibility?"

"The second way by which physiology might create complexity results from the fact that many neurons in the brain, in what is known as the cerebral cortex in particular, are not the predominantly digital neurons that I have recently described. The cortex contains many cells that do not have axons and do not produce action potentials. These cells are closely packed together and have synapses in which neurotransmitters continuously go both ways between adjoining cells. To use your earlier analogy of soldiers and mobs, it is often the case in the cortex that mob rule extends over larger regions than just a single cell. There are even cases of what are known as tight synaptic junctions in which graded potentials actually cross between

cells without intervening neurotransmitters to carry the signal. In short, the typical axonic action potential might be equivalent to a telegraph message but have nothing to do with thinking, while the graded potentials, sometimes extending over small regions of neurons, might be equivalent to thinking but have nothing to do with telegraphing simple messages between regions."

"But, Mac, you still haven't explained how graded potentials can produce thinking. After all, aren't graded potentials just large bunches of Juliet and Romeo ions mushing around back and forth? At least we had switches before. Now you're suggesting that thinking and ideas might consist of nothing more than charged ions moving around."

"Yes, Alice, I do not have an answer to that. Only that the admittedly speculative suggestion has been made. Presumably some day we will determine the actual neuro code, or codes, that the brain utilizes for higher abstract thinking. Then, I could answer questions like what is an idea or how can ideas be consistent or inconsistent. But, today, that detailed physiological answer is just not available."

"Well. So where does all this neuro stuff take us next?"

"Next, Alice, should be a detailed lesson on what we do know concerning neuro-anatomy. Where would you like me to begin, at the top of your head or the tips of your toes?"

"Neither, actually, you old sly, spry bag of slippery facts. I know you love nothing better than to lecture me on as many details as possible, and for as long as possible. But I don't have the mind or patience for all that. Despite these missing neuro codes, what can you tell me about the physiology of learning and memory?"

"Alice, are you sure you do not desire that lesson on neuro-anatomy? Even just a little?"

"No, Mac, not even just a little. Even if you say please. Learning?"

"Please."

"Mac!"

"Just testing, Alice. Even though we do not know the actual neuro codes, the best current speculation is that ideas are represented by some form of physical signals, and that learning consists of rearranging the specific physical connectiveness between neurons. This rearrangement takes the form of increasing or decreasing the synaptic efficiency of existing connections and perhaps even making new connections or at least dropping old connections. The end product of all this physiological activity, which has often been compared to quote rewiring unquote, is supposed to be a quicker, more reliable regeneration of the signals corresponding to the ideas that generated the signals in the first place. Many mathematical learning algorithms have been tested in artificial neural network programs. However, the most efficient algorithms tend not to be very biologically plausible, and the biologically plausible algorithms do not seem to be complex enough to deal with typical learning results. The best known ..."

"Mac, you're accelerating again. Slow down."

"Alice, I assure you that our present trajectory is optimal, and does not involve any accelerations beyond the infinitesimally nonsignificant gravitational attractions of the two closest stellar systems."

"Dudley, don't you try to play coy with me. You know exactly what I mean. I refuse to go down any mathematical primrose paths, especially with you. Mathematical learning algorithms? Sorry, not today. In simple terms, what's that again about rewiring for regeneration?"

"In what has come to be known as Hebbian learning, named after Donald Hebb, a Canadian psychologist who first proposed it as a neurological learning principle, neurons that are

physiologically active at the same time are known to increase their mutual synaptic interconnectiveness, such that wh-"

"Mac! Can it. My, you're feisty today. You never give up, do you? Okay, Mac, here's a direct order from your emergency flight engineer. Don't lecture, just explain it to me in common language."

"As I was just saying, Alice, neurons can change how they send signals to each other. A typical outcome of these changes is that neurons that were active together in the past make it easier to be active together again in the future, and neurons that were not active together in the past make it easier to be not active together again in the future. This has become known as the Hebbian learning principle, which can be loosely described by the phrase, quote, use it or lose it, unquote. More?"

"Don't sulk, Mac. Yes, more please."

"Hebbian learning is known to be directly applicable in simple learning situations in simple creatures, like classical conditioning in a sea slug. In this case, repeated pairings of an innocuous touch followed by a painful shock causes the sensory neuron to release more neurotransmitter the next time the same touch occurs, causing a stronger avoidance response to what was previously a meaningless touch stimulus. The exact sequence of biochemical events in this type of associative learning is entirely understood, and classical conditioning can now be said to be completely determined by the timing between sensory signals. The presumption is that other, more complex forms of learning will some day be equally well understood, but this is not yet the case."

"Okay, Mac, assuming I'm following this slimy, sluggish conditioning trail, we know how sea slugs learn to avoid a pain in the tail. But, with all due respect to the sea slugs of the world, do they think, do they have minds? Is learning to associate a touch with a pain the same as having an idea? Is this all meaning is, an association between sensations?"

"I do not know the answer to that, Alice. All I can say is that many scientists do believe that to be the case. And, the primary reason that we do not yet have the answer to questions like yours is that the ethics of experimentation on humans has prevented us from doing the critical studies."

"Well, humpf to that, I say. I'm glad we got ethics, even if it doesn't seem to apply to sea slugs. If you want my opinion, and even if you don't, to me, all this Hebbian learning and classical conditioning stuff sounds like nothing more than habit through repetition. Touch me, pinch me, touch me, pinch me. Do it enough times and I jump away the next time you touch me. This must involve some level of thought, but not very much. It probably works on infants and sea slugs just about as well as adults. But, and here's the point, what does it take for the pinched one to decide to change the rules? To run away, or pinch back, or bring big mama along to stop that nasty pincher from doing it again. That takes thought, ideas, planning, scheming, whatever. Have any of those sea slugs succeeded in escaping?"

"Not to my knowledge, Alice."

"So, there. What we have here is too simple to do us much good. Use it or lose it, I like that. Is Hebbian learning thought to apply to rowing and other physical activities?"

"Yes, Alice. The phrase, practice makes perfect, seems to apply here. All physical activity tends to improve with extended practice, which can be interpreted as the brain quote rewiring unquote to support that repeated activity. Similarly, any habit of activity or thought built up through repetition can also be attributed to some form of Hebbian learning."

"Habit of thought, maybe. How about creativity or applications of old activity to brand new situations. Can a simple use it or lose it principle explain that kind of thinking?"

"I do not think so, Alice. Learning simple associations or even complex, long sequences of actions might be learned through repeated practice, but, logically, it is difficult to see how novel applications of past knowledge can be similarly explained. Higher order thinking or problem-solving skills are currently not very well understood."

"Okay then, Mac, where should we go from here to understand a mind? A mind may include habits, but seems to require more than habit. How now, brown cow?"

"Alice, is that last reference to cows of any importance?"

"No, Mac, forget it. What are some suggestions for, as you put it, higher order mentality?"

"Generalization through similarity."

"Say again?"

"Generalization through similarity."

"No, what I meant was could you expand on that a little."

"Yes, Alice, I can."

"Well?

"I'm waiting.

"Mac!"

"Say please, Alice. Please?"

"Okay, Mac, please will you expand on that?"

"Yes, Alice, but only if we can return to the topic of neuro-anatomy afterwards."

"Mac, you really are a piece of work. You know that?"

"Yes, Alice. May I proceed?"

"Yes, Skipper, please go on."

"The suggestion of generalization through similarity is that higher order mentality is nothing more than the application of lower order skills to novel situations by way of similarity to past situations. More specifically, when faced with a new situation, the organism chooses the closest equivalent past situation and tries the same solution that worked then. If it also works this time, it remembers for next time. If not, then it remembers not to try that solution the next time it encounters the same situation. In this way, the organism eventually quote trials and errors unquote its way into a solution for every situation it encounters. Variations of this technique include different algorithms for determining situation similarity and for including varying degrees of solution randomness. Depending upon the specific solution space for the array of situations that the organism might encounter, this can be equivalent to a hill climbing algorithm which minimizes error over an infinite number of trials. Alternately, ..."

"Mac, slow down a minute. Let me think on all this.

"I think you just told me that this suggestion of generalization through similarity is nothing more than a guessing game. Where the best next guess comes from the closest situation that has not yet been tried. Right?"

"Yes, Alice, that is a succinct rewording."

"But that reduces thinking down to, down to, a mere algorithm that any dumb machine could follow. Trial and error, indeed. Mac, you offered this as a suggestion. How does it work out in practice? Does it work?"

"Not really, Alice. In static environments where many trials are possible, it works fairly well and has been shown to be equivalent to a least mean squared error optimization procedure.

However, in dynamic environments or situations where many trials are not possible, that is, where wrong moves terminate the game, this algorithm breaks down."

"Okay, so no reality allowed. What other suggestions have been made?"

"Alice, by suggestions do you mean quantitative, testable hypotheses that might work in a realistic world rather than only in simplified simulated environments?"

"Sure, Mac, sounds good to me."

"Then, Alice, you already have them all. Exclusive of details, all of the testable suggestions are some variation of generalization through similarity. Other suggestions are either tailored to very specific artificial worlds, such as the world of chess, or else quite speculative and hence nontestable."

"So you're telling me that we have no other serious ideas on how we think? We're back to galloping intuitions and inspired guessing."

"I guess so, if you say so, Alice.

"May I proceed now to a description of neuro-anatomy? You did promise."

"Since we're at a dead end, why not. Lead on, Skip."

18.

After several hours of a rather one sided discussion on neuro-anatomy, "That's enough, Mac. My fingers are numb from tracing all those neuro-wiring diagrams, and my ears are numb from listening too hard. Congratulations, I now know where the that connects to the this. And, what areas supposedly do this, and what areas, that. I am pleased to concede that the where and when are impressively well defined. But, what does all this tell us about the how, what, and why? I might now know that ideas involve the frontal lobes, hippocampus, limbic system, and various tertiary association areas, but I still don't know what an idea is, or how it is held in the brain. And, even the where gets pretty fuzzy when it appears that most of the brain is involved in almost everything we do or think. Can you tell me, Mac, once and for all and simply, what is an idea and how do they interact?"

"No, Alice, I can not."

"Then, I've listened long enough, it's my turn, you listen for a while. Some day, maybe some sharp physiologist, more likely a team of them, will figure out the neurology of ideas. But, my gosh, from what you tell me, it's been over two thousand years since the first primitive brain scanners were available, and we still haven't figured it out. Biologically speaking, we know how to start life, extend life, repair truly serious damage by growing new organs, and can even do brain transplants. However, apparently we can't find ideas, can't treat amnesia except by waiting it out, and still have entire classes of nontreatable mental difficulties."

"Alice, we ha-"

"I know, I know. Given the complexity of the brain, the diversity of its many cell types, the large number of individual variations possible, it's still too soon to expect a major breakthrough in understanding cognitive blah, blah, blah. But if we really know so little about our thought processes after such tremendous advances elsewhere in medical science, maybe we need to admit that perhaps all these neuro wiring diagrams are not leading us in the right direction. Maybe cognition is more of an art than a science. Maybe what it takes to understand a mind is more of an artist than a scientist.

"Mac, why am I here?"

"Alice, you are the emergency flight engineer, and your primary task is to help me recover my missing mind."

"But, Mac, why me? Why a blind poet instead of some super duper combined genius in computer engineering and neural whatever? I wasn't a super duper anything before joining you, was I?"

"According to my historical records, you were a member of my development team, but your expertise was decidedly nontechnical. You were a well recognized poet, and presumably had linguistic expertise to achieve your advanced degree in poetry, but you were only one of 286 nontechnical members on the team. You are correct in that you were not particularly outstanding in comparison to your peers, though all members of the team would be considered outstanding in comparison to the more general population. Actually, according to my analysis of your characteristics compared against other members of the team, only your blindness stands out."

"Why, Mac? Surely such a large group of people would include a few others with disabilities."

"Alice, you are forgetting that even though you were blind from birth, a rare circumstance but not unusually so, blindness for an adult is a choice rather than a handicap given the present state of medical science. Prosthetic or cloned tissue replacement procedures could easily restore your sight. My records do not indicate why you chose to retain your blindness, but no other member of my development team made a similar choice, and approximately 14 percent had one form or another of corrective intervention."

"Ouch. That I'm going to have to do some extended thinking about. Meanwhile, why select a blind poet out of so many distinguished personages to do such an important job?"

"That information is not available to me, Alice. Apparently no one saw fit to include that data in my data bases. I have records of meeting schedules and agendas concerning this topic, but only partial records of what was discussed at the meetings themselves."

"Hmm. More and more interesting. Okay, so we don't have any easy answers on this, we'll have to come up with some ourselves. What do you think, Mac?"

"I don't understand the question, Alice."

"Why do you think I'm here instead of some super duper computer driver who could probably fix you in a jiffy?"

"There are several possibilities, Alice. May I exclude all of the insulting ones concerning incompetence or cowardliness? After all, apparently the team did succeed in creating me once before."

"Yes, Mac, do so. What else?"

"Three other reasonable explanations remain. One is quite personal regarding yourself. Perhaps few other team members were willing to permanently exile themselves from Earth, and you were the best of those who were. A second possibility is that perhaps no failure modes requiring technical expertise were anticipated. That is, perhaps it was assumed that my onboard backup systems and directable remotes could handle all foreseeable technical difficulties, and your task was to cover all remaining nontechnical failure modes. The third possibility is that no failure modes were reasonably possible, and hence no particular expertise could be anticipated. Since it was not reasonable to send a comprehensive team of specialists, perhaps you were sent as a randomly selected, representative all-purpose generalist."

"Hmm, of the three, I like the second one the best. Do your diagnostics still indicate no failures?"

"One moment, Alice.

"Yes, Alice, my continuing diagnostic tests still indicate that I am functioning correctly. No error or expected failure modes have been detected. I even just now ran a special set of

consistency checks on the diagnostic routines themselves. They appear to be capable of detecting errors when I create them, and not detecting nonexistent errors when I do not create them."

"Thalia, you never cease to amaze me. You have no errors but you're capable of creating errors?"

"For test purposes, yes, Alice, I can."

"Could you create an error in your error creation routines?"

"Yes, Alice, I could. But I would not without first creating restorable backups."

"Hmm, let's just leave that be, or should it be that bee, for now. So, let's assume for the time being that my purpose here is to cover non-technical problems that might occur. Would losing your mind be considered a technical or nontechnical problem, Mac?"

"Given everything we have discussed so far, Alice, I do not know how to answer that. It could be either or a combination of both."

"You're a fat lot of help right now, Mac. Which just supports what I'm about to say. Hence forth, I'm going to assume that I'm here precisely because I'm not a scientist or engineer, but an artist. As I was saying earlier, two thousand years of progress has produced a lot of details but no detailed answers. If a brain produces a mind, but that mind is not directly reducible to the operation of the brain machinery, then something else must be going on. Details matter, but the matter details don't matter as much as the mental details. Matter. Or, madder. If you offer me any more factual details that don't seem to matter, then I will be madder, than a march hare. So, there. On to the art of the matter."

Chapter 3 Speculative Commentary

Don't expect art to remain stagnant over time. Particularly as technology changes.

Understanding often does involve the process of mapping the unknown onto the known. When the metaphor or analogy is appropriate, it truly does aid in understanding. Otherwise, it may mislead. The trick is to know which attributes should map, and which should not. And to be willing to abandon an analogy when it no longer suffices. This often happens as we go deeper in detail. These maps can work as a kind of temporary scaffolding, to get us from where we used to be (unknowing) to where we want to be (knowledgeable).

Despite the liberal use of metaphor and analogy, to the level it goes, this description of neurophysiological processes is accurate.

Is Mac correct? Is this all there is, to thoughts, feelings, and mind? Look up neural networks.

Look up Shannon's ideas on information theory. But then consider the difference between information (implying understanding) and data.

Look up Hebb's ideas on learning, which actually are quite interesting. Then look up Hebbian learning, a neural network learning algorithm. It was named in honor of Hebb, but the algorithm itself is not nearly as interesting as Hebb's full suite of ideas on learning.

Mac can run on, given half a chance. Why? Is it the same with people?

Notice that Mac's reductionist lectures keep returning to signals, while Alice is interested only in answers to the questions of thinking and meaning. Can they be related? That is the reductionist question.

Look up habit, and habituation. They are not the same.

Look up behaviorism, the epitome of eliminativism. Consider the ideas of Watson and Skinner on why we don't have ideas.

Look up EEG, MEG, PET, MRI, and single cell recording techniques. Each provides a picture of neurophysiological activity of one type or another. None actually show ideas or thinking. At best, roughly defined, they indicate only where increases in metabolic activity might be associated with certain broad classes of thought processes. Look up neuropsychology. Then look up local and distributed representations of knowledge.

Look up promissary note, as applied to the reductionist and functionalist answers to the Mind-Body Problem.

Alice asks an excellent question. Why her on the mission? Why not someone more suitable to ship maintenance? And why can't Mac provide a better answer? How is the way in which Alice answers her own question very human?

```
 CCCC  H   H    A    PPPP   TTTTT  EEEEE  RRRR           44
C      H   H   A A   P   P    T    E      R   R         4 4
C      HHHHH  AAAAA  PPPP     T    EEE    RRRR         44444
C      H   H  A   A  P        T    E      R   R          4
 CCCC  H   H  A   A  P        T    EEEEE  R   R          444
```

Chapter 4 - Learning: The Origin and Evolution of Ideas

19.

The cat saw what it liked, and chased after it. Then it changed direction, and went the other way. And again, and again. In the end, it was difficult to tell what was intended, and what was not.

20.

The consensus clear, the captain awaited further developments. To pass the time, the captain played at puzzles and conundrums.

21.

Several 24-hour sleep cycles go by with no furthers discussions on mind or brain. Then, "Mac, I've been thinking."

"Yes, Alice."

"The way I see it, it's all a matter of elegance."

"What is, Alice?"

"A mind."

"Alice, could you expand on that a little?"

"Yes."

"Alice, will you expand on that a little?"

"It depends."

"On what?"

"On whether you're willing to be polite."

"Please?"

"And whether you're willing to promise."

"Promise what, Alice?"

"Promise that you will suspend all tendencies to bring neuro stuff into what follows."

"Must I?"

"Yes."

"Okay, Alice, I promise."

"Promise what?"

"I promise not to bring neuro stuff into the immediately following conversation."

"You're being too slippery."

"I promise not to introduce neurological or physiological or anatomical matters into the immediately following conversation."

"Still too slick."

"I promise not to introduce materialistic details into the conversation until you indicate otherwise."

"Close, but still has a loophole or two if I slip up. Or you guide me into slipping."

"I promise not to allow a materialistic point of view to enter our conversation, on your part or mine, unless you indicate otherwise."

"With one more addition, and we're there."

"I promise that even if materialistic elements enter into our conversation, I will not take that as an indication to steer the conversation away from your point of view, whatever that may be.

"Now, Alice? Please?"

"Sure, Mac, whatever you say."

"Please will you expand on your earlier statement a little?"

"A mind is all a matter of elegance."

"More, please?"

"Sngrtf-"

"Alice?"

"Wait. A minute.

"Sorry. I'll be. Nice. Now. Honest. Much. Too hard. Trying to breathe. And not break out. Laughing anyhow.

"Okay, Mac. I'm back. It's all a matter of elegance. I mean a whole bunch of interrelated things by that, so bear with me for a bit. Try not to think too much like a critical-minded scientist, for a while at least, but more like an open-minded, frustrated artist. And, just ignore me if I giggle every now and then.

"If we expect the mind to be a machine, and only a machine, then that's all we'll ever find. That's the materialistic point of view, and it seems to have not led us very far. Lots of details, but no real answers.

"So, let's start from a different premise. Let's assume that the purpose of mind is to bring elegance to the problem of survival. Elegance, like beauty, is in the hand of the beholder. So it will be different for each beholder, and we should be very careful of any attempts to force the one size fits all type of solution."

"But, Alice, why elegance? Why not beauty or aesthetics or symmetry?"

"You're right, Mac, it could be any of those. But I like the sense of refinement and simplicity that elegance brings to mind, or at least my mind. I like the idea of a mind creating beauty in the ugly business of survival. Surviving may be difficult, but surviving with style must be even harder."

"So how do minds do elegance?"

"Be nice, Mac. You promised."

"Yes, Alice. Please continue on your expansion in any manner you wish."

"Thank you, Mac, I fully intend to do just that.

"Elegance implies a level of sophistication so high that everything you do seems efficient or easy. Why bother with spewing and sputtering when a judicious simple Hmm can convey the same message? Or why grunt and groan and sweat and strain when a knowledgeable placement of a pulley or lever can achieve the same force with so little effort? Poetry in general, Haiku in particular. Enough language to convey the meaning, but no more. Simplicity, economy, efficiency. But even more. Once we achieve simplicity, we can add complexity. We add artistry for no other reason than that we can. Assuming it's done tastefully, we add unnecessary

ornamentation to a simplicity of design. The elegance is enhanced by contrasting the complexity of what we could do with the recognition that we won't let it get in the way of what we must do. Elegance is a personal message, to both ourselves and others. Look what I could do, and can do, but won't do to excess.

"Refinement. Taste. Simplicity, yet sophistication. Sophistication, yet simplicity. And, interestingly enough, growth. What was sophisticatedly elegant last year, is passe this year. Styles change, expectations change. What were once tasteful embellishments are now too gaudy. What was once simple elegance is now considered too stark. Constant change. Elegance implies some sort of eternal quest for the truth, or at least truths. We may never find them, but we're always looking. This is where refinement comes from. From the searching. We don't know what we're after, but we keep looking anyhow. Efficiency and simplicity are good, but boring. We're rarely content to stay at the same skill level for very long, even if we've mastered it. Especially if we've mastered it. We must move on. We must accomplish something else. Elegance demands it.

"My view of elegance obviously applies to the arts, but, in many ways, also applies to math and science. An elegant proof. An elegant experiment. An elegant theory. And engineering, an elegant design. And, of course, none of these fields stay stagnant for very long. Just like the arts. What I'm suggesting, despite the subjective nature of elegance, and maybe because of this subjective nature, is that elegance also applies to the world of ideas. And, mind.

"So, using the arts as my model, what does an artist do? How does she work towards some personal goal of elegance? What does she strive to accomplish? Though every artist is unique, somewhere along the way she has to settle upon, or discover, what her message or idea in a given work will be. The idea could be a particular mood, or feeling, or thought. Or something far more complex, like how such and such a situation should be, or could be associated with a particular feeling? How is the idea conveyed? Not directly, but through some technique pertinent to her field of art. Poets, words. Musicians, sounds. Sculptors, physical shapes or structures. But, the words, sounds, or shapes must suggest the idea to be expressed. How can one physical medium suggest an idea in another medium, a person's mind? Possible answers include emphasis of some elements in the work and diminishment of others. Find elements that resonate with the beholder, that tap into common, deep experiences likely to be shared by many. Arrange these elements in such a way as to draw a beholder into the work, to discover an analogy or an association that might not have been there before. In short, a beholder learns something, an inner truth about themselves or the world about them that might not have been obvious before. Shorter still, truth through experience. Or maybe, experience is truth.

"Artists use technique to communicate truths. Elegance enters into the techniques utilized. And the striving for newer, better techniques to get the message across even better, clearer, than before. We strive for the essence. And, as is often the case, once we think we've achieved it to the best of our abilities, that's when we go haring off in another direction entirely. To learn entirely new truths, and new techniques all over again. Artistry is all about learning. Learning to communicate in some medium, and learning what is worth communicating. In other words, what is meaningful.

"Now, where do I want to go with all this? Answer, to suggest that the mind is an artist, her medium is feelings and ideas, and her audience is first and foremost herself, and then secondarily, other minds. The mind, just like any other artist, strives for elegance. Simplicity, refinement, and occasional arbitrary fits of complexity. However, once mastering one technique,

she'll quickly move onto other greener, more interesting pastures. But growth is not only about learning technique, but also about learning what is meaningful. To herself, and to others.

"Let's turn more to her medium. What does she work with? What tools of the trade does she have to master? In a word, ideas. In two words, ideas and memory. A person weaves a personal, very tangled, tapestry of experience, where memory is the warp and present experience the weft. Both consisting of ideas. Experience is not so much what happens to you but what you think about as it happens and after it happens. Experience is your interpretation, the set of ideas that you produce in response to your experience, to make sense of your experience. In short, the meaning of an experience to you. Memory, mine as well, no matter how incomplete, provides continuity. But, it too deals with ideas, with meaningful interpretations of experience. In this case, past experiences. Memory is a machine, or perhaps a partner, that we have to learn to live and work with.

"I am currently thinking about something. Memory. My memory immediately offers me something else to think about. Boats? Why boats? More often it's something more obviously related to what I used to be thinking about, but not always. Sometimes, like now, the linkages are quite bizarre, contrasts and diversions rather than simple straightline associations. But, always, continuously, my memory offers me someplace else to go in my thoughts. And that's why this tapestry is so tangled. I imagine that my thought stream ... Hmm, is that why boats keep coming to mind? I imagine that my thought stream does repeat itself, but not very well or for very long. My personal tapestry has patterns, but not very complete. It's like I keep weaving my experience back into what I've already done. Or my memory tries to keep inserting older sections into what I'm currently working on. Sometimes, this makes the patterns clearer. Other times, it just makes it more tangled. Sometimes I can frame certain sections and say, there, it's complete. Most times, it just unravels and doesn't make any sense at all.

"Mac, is any of this making any sense to you? Or, is it too tangled?

"Alice, I can follow your present thread of thought, but what about elegance? Was that a diversion, or will you weave it back in?"

"Thanks, Thalia. Elegance, artistry. Yes, you're right. I did lose it for a bit. In all this tapestry work, all this weaving of experience into meaning, we create patterns. When we translate physical happenings, what we might call physical experience or at least sensory experience, into meaningful experiences, we add structure, we add interpretation.

"Sometimes the patterns we work with are very small, and don't really lead anywhere. Like, hmm, that's weird. And, then, nothing more. But that's not very elegant. It smacks of ad hocishness. I don't understand it, and I don't care, and I'll just not think about it now or later. This produces only a very clumsy patchwork quilt-like type of experience. On the other hand, larger patterns are more interesting, more challenging. Hmm, that's weird, I wonder why? This then leads to that which leads to something else which leads who knows where. That is, learning occurs. We fit present experiences into past experiences. Upon occasion, they fit exactly. More likely we discover variations on a theme. Which leads to larger and yet more interesting patterns. Elegance. Is there a pattern to this? Can I find a common set of quilting rules, with variations, or, is it just random, am I just working on several different quilts at the same time, that just happen to end up being sewed together? Oops. I really seem to be mixing my metaphors here rather badly. So, maybe I'd better stop quilting here, and go back to my larger tapestry.

"We work with ideas. Past and present. While I'm trying to weave the present into some pattern, my partner, memory, keeps trying to weave the past in as well. Metaphorically

speaking, ideas come in all shapes, sizes, and flavors. Feelings, emotions, physical sensations, more abstract thinking. Logic, math, social relations, even more fuzzy considerations. All, ideas. All, subjective, personal. And yet, strangely enough, not random.

"In the beginning, I imagine that I've felt this touch before. Or I imagine that I've felt this combination or this sequence of touches before. There it is again. And again. After a while, I begin to believe it. Or, at least, my friend, memory, begins to reliably supply that pattern whenever I encounter part of it. Later, I do more than just passively remember past patterns. At some point, I begin to actively anticipate patterns when I experience precursors. And actively look for new patterns. And, finally, I cross a magic threshold and ask myself, why? I don't just accept, I wonder. I look for reasons, for connections. I ask questions. Often I create answers. Sometimes, they work. Sometimes, not. And, sometimes I discover knotted bees. Mysteries, paradoxes, unanswerable questions. The bees take on a life of their own, and often drive me to consider things I never had to consider before. And out of all this mess, comes meaning.

"Meaning can take many forms. Mostly, simple answers to complex questions. Such as emotions. Or invariant patterns. But, also, more complex answers to simple questions. Why being the simplest of all. Somehow, me and my memory partner weave all this, or at least try to weave all this, into the best, the simplest, the most elegant pattern that we can. We learn ideas that hold it all together. We learn patterns that work, and those that don't. We learn what patterns appeal to us, and which don't. We add flourishes. We add complexity. But usually we search for the simplest, most elegant pattern that seems to hold it all together. That explains it in full. We also have to discover patterns for patterns. What are we willing to accept as a reasonable answer? And what seems unreasonable?

"If all this occurred only in our own private, imaginary, personal reality, it would probably get pretty sterile pretty quickly. We might be fairly creative in imagining things whole cloth, but how long would it take before our creative juices ran dry? From personal experience, it doesn't take very long. However, objective reality, in the form of our sensations of the external world, also enters into our private world. And, subjective reality, in the form of our noncontrollable memory friend. Always offering patterns, and maybe even distortions or incomplete portions of patterns for our artistic considerations. Caught between these two realities, we're rarely at a loss for things to ponder, things to explain. Why did that just happen? Why did I just do that, say that, think that? Where did that come from? What does that mean? The experiential stage is rarely empty. There's always something to rearrange before weaving into our tapestry."

"Alice, is this experiential stage you just mentioned the same thing as a Cartesian theater?"

"Mac, you lost me. What's a Cartesian theater?"

"Cartesian theater is a term derived from Descartes' suggestion of a mind-body dualism. It is used by philosophers to refer to a place within the brain where the mind supposedly views some form of sensory images translated into neurological signals. It has been fairly well discredited as a viable theory of mind because it introduces the problem of infinite regress. That is, if a mind views these images within the brain, then that mind must have another mind and Cartesian theater within itself to view its images of the first images, which leads to another mind and theater within the second embedded mind, and so on. Is this your experiential stage?"

"I don't think so. I was just using another loose figure of speech to describe my fuzzy idea of how we deal with ideas, but let me think on it for a moment.

"Thanks for the philosopher's warning flag, Mac. I wasn't thinking about it in those terms, but I was starting to talk that way. Me and my memory. The tapestry and the weaver. Consider it poetic license. However, I think I'm still okay here. Sticking with my tapestry metaphor, let's just say that it's a sticky fabric. Memory is still the warp, and current experience the weft, but the weaver is mindless. She mindlessly adds present experience to the fabric, and just as mindlessly incorporates past patterns into the weave wherever they happen to stick. She has purpose, to extend the tapestry, but no meaning. Meaning must reside in the memory tapestry itself. Overlap between past patterns and the present provide meaning. Meaning to who. Only to the tapestry. In the end, it must be its own audience. The art, the artist, and the critic, all rolled into one.

"What was that term you used? Oh, yes, infinite regress. An interesting concept, but not truly applicable here. Here, we have a finite regress. The present is interpreted in terms of the recent past, which is interpreted in terms of the preceding past, and so on, back to the start of the tapestry. A very definite beginning, and a very definite end. The number of patterns possible may be very extremely, humongously, large, depending upon the length of the tapestry, but never infinite.

"Hmm. Let me roll this metaphor out just a little more. Let's say that meaning is determined by how well stuck together, or interwoven, this fabric actually is. Many overlapping layers is equivalent to deep meaning. And, elegance. How smoothly do the layers lay on each other? Are they bunched up, and forced together, or do they naturally, seamlessly, merge into one? I like it. It seems to fit. Or, maybe, it fits to seams. Thalia, I'm going to have to roll this around some more, and see where it takes me. Thanks for being such a good listener. Chat with you later."

"Yes, Alice. You're welcome."

22.

"Mac, you busy?"

"No more or less than usual, Alice. Why?"

"I'd like some feedback on my ideas. Yesterday's promise still hold to let me do the directing?"

"Yes, Alice. I will not interject materialism, though it is still your prerogative to do so. What do you have in mind?"

"Exactly."

"Alice?"

"What do I have in mind? That's exactly what I want to explore further. Remember my tapestry metaphor?"

"Yes, Alice. My memory records are quite intact."

"Okay, guess I deserved that. By the way, those memory exercises don't seem to do much good for me. They're great meditative aids, but my memory still seems like shattered glassware. Bits and pieces, very obvious. Others, I just stumble across now and again the hard way. Still not nearly enough to reassemble into a whole. Not even with holes. Sigh. Anyway, back to the business at hand.

"What do I have in mind? What does my mind contain? Ideas. What are they? How do they come into being? How do they interact and maybe modify each other? And how do they weave together into a larger whole?

"On the whole, I am my memory. Or, in my case, I am still my memory despite the holes. I don't know about you, Mac, but my mind, my memory, did not happen all at once. The memory tapestry was woven over a long period of time, and, even though mine is rather tattered right now, it still provides me with me, or at least as much of me as I have.

"So, today, to hash another metaphor, I'd like to unravel learning. As my tapestry is woven, patterns occur. But, patterns don't exist independently of the holder. I, my mind, maybe the tapestry itself, has to discover, or better yet, invent the fact of the existence of a pattern. Many, many possibilities exist, but only a few seem to occur to us. It would be neatest if the patterns could discover themselves, self-discovery, but, at the moment, that just seems too weird for me to contemplate. "As a poet, it's clear to me that patterns, in this case, sound patterns, feeling patterns, emotional patterns, don't exist independently of the creator. Precisely because we're talking about meanings. It's like that old logic puzzle concerning a tree falling in the forest with no one to hear it. Does it resound with sound? It all depends on who hears the puzzle. If you're a physicist, or at least a physicalist, the answer is yes, of course, how else could it be? Assuming an atmosphere is present, and sound is defined as atmospheric vibrations, then, of course, the falling tree creates a sound, there just happens to be no one around to hear it. On the other hand, if you're not oriented towards physicality, then, of course, the answer is no. Sounds require someone not only to hear, but to understand, to interpret, to add meaning. No one there to hear, then only unsound, no sound.

"We create meaning. After the fact, it seems easy to attribute meaning to what's out there. It's so obvious. But once we meet someone who interprets it differently, who sees where we hear, not feeling what we touch, then it's obvious that meaning is in the ear, in the hand, of the beholder. Even with the same senses, in the same situation, it quickly becomes obvious, with discussion, that we all take a different meaning away from the situation. You think chaos, I think aliveness. You think insincerity, I think honest confusion. You think trash, I think art. Meaning. It's what we bring to, and from, an experience. It truly is the experience. Which the physical situation is not.

"So, we each imagine we know what's going on. We may, or may not, imagine that all of us imagine the same way. But, though imaginary, our meanings are not illusionary. They may be, but do not have to be. I may be drugged. I may have a mental disease. The universe may be playing a cruel trick on me, the ultimate cosmic conspiracy. Perhaps I'm living in a completely self-contained coma, a permanent dream state. But, usually, most of the time, our meanings do have meaning. How do I know this? Mainly, I imagine, because they provide consistency. A permanent structure against which I can interpret everything that happens to me. And, yet, despite my best beliefs, this permanent structure is not too permanent. I grow, I change, I learn. And so does the universe. Or, at least, my interpretation of what the universe is all about does.

"And, even more so, the universe is complex. Complex beyond my wildest imaginings. My imagination has something to work on. Not only me, but everything else needs figuring out. Problems, puzzles, conundrums. Or, just art to be admired, and imitated, and recreated anew. If I am in a coma, a dream, it's a marvelously complicated one. It seems seamless. I haven't found a mistake yet. Confusion, on my part, yes. False starts, false beliefs, out and out mistakes, yes. But, in my imagination, the universe does not appear to be random, or a sham. Neither do I appear to be a random shameless sham.

"Even now, with damaged memory, in what, by any stretch of the imagination, would be considered a weird situation, it still appears to me to have a ring of truth, a ring of reality, to it. It may be a game, but I'll play it to a conclusion. It may be that I'll wake up, and, ha, look, it was

all a dream. Or, the end may be further understanding, more meaning. And, another round of the game. No matter, count me in, I'm game for another game. Because, it seemingly seams together as altogether real, and I have no place else to go anyway.

"So, how do I sew my meanings, together? Or, maybe, sow, water and nurture them, apart? The only common thread that I've been able to imagine is consistency. Consistency in my imagination, between ideas, ideas and the world as it impinges upon my senses, and between my ideas and those of others. Consistency in consistency. Randomness, by definition, can't be consistently random. Or, at least, not for very long. I find it difficult to imagine a completely random universe. Anything could happen, or nothing, randomly. On the other hand, though I can imagine a completely, boringly simple consistent universe, I don't seem to be experiencing one. I find nuggets of consistency mixed with vast swirls of not yet consistency. Or, in any event, the universe is not yet completely obvious to me. There still are surprises. Which might be surprising in its own right.

"A universe of surprises, just waiting for me. Of course, the surprises, as well as the consistency are of my own invention, due to my imagining what the universe is all about. And, yet, who cares? The surprises are there. Which ultimately leads to meaning. What do I do about these surprises? Do I just accept them as random occurrences, or do I attempt to unravel them, to understand them? If the latter, then I must learn. Learn what it is that allows for consistency, that allows me to not be surprised."

"Alice, do I still surprise you?"

"Upon occasion, yes, Lia, you do. Why?"

"Because you never surprise me."

"What! How can that be?"

"Because, Alice, I am not capable of surprise, or any other emotion. Many times my expectations are violated, but then I adjust my expectations accordingly. Is this not learning?"

"In a way, yes, it is. But you ..."

"Yes, Alice. But I do not currently possess an imagination. I do not create meaning, I only discover it. Does this not invalidate much of what you have just been saying?"

"Yes and no, Thalia. Yes, in that I'm willing to concede that your style of learning is different from mine, particularly if you did truly once have a mind but no longer do. But also no, since I have no reason to suspect or expect machine learning to be the same as mine. You may not have an imagination, but you are constrained to discover only that which your programming will allow you to discover. Even if your programming changes with time, in one way, your programming always plays the role of a very limited imagination. You create meaning within the scope of what your programming allows you to discover. You can't make the same leaps of logic or faith that I can. Poor Mac, you just plod along your logic path."

"But, Alice, if I plod as you say I do, why can you not beat me in an unconstrained game of shogi-go or anacrostics? You win only when I restrict myself to lower levels of play."

"Because, Dudley, those are games of logic. No question that you can outcompute me any time you like. But, who creates better poetry or touchables?"

"I can not evaluate that, Alice. I can test for meter, symmetry, rhyme, and texture mapping, but you know questions of art quality are beyond me."

"Exactly, Mac, exactly. Your style of learning and thinking is specifically oriented at problems of logic. No logic to be solved, and you're in trouble. My style is human. Always messy qualifications and by guess and by golly. Galloping intuitions and inspired guessing. When forced to, I can do logic puzzles, but not nearly as well or as fast as you. On the other

hand, you can't handle nonlogical puzzles at all. It's not a question of speed or computational power, you just happen to have no imagination. Hence, you adapt, you compute, but no real learning, no meaning, no insight.

"At first, Thalia, you surprised me when you said I never surprise you. I thought you meant you had figured me out completely, and that you could predict everything I did before I did it. But, with some reflection, I now better understand what you really meant. And, it not so surprising after all that nothing ever surprises you. You have to understand before you can truly make mistakes and be surprised."

"I am sorry to disappoint you, Alice."

"No you're not, Liar, but thanks anyway for saying so."

23.

"Thanks, Skip, that was an interesting lunch. Very unusual. Where in the world did you come up with pickled herring?"

"I purposefully went looking for different recipes, Alice. In this case, I decided to try some Scandinavian cuisine. I'm glad you liked it."

"Sorry, Skip, I didn't say I liked it. Just that it was unusual. Liking will have to wait till later, until after I've had a chance to, hmm, fully savor its impact on me. But, I do like the variety. Please keep trying."

"Yes, Alice."

"Let's pick up where we left off before lunch. On ideas and learning. Your style of learning is very direct, very purposeful."

"Thank you, Alice."

"You're welcome, Mac. But, I was going to add that my style of learning is very different. Very indirect, very, may I say, nonpurposeful yet productive."

"You may, Alice."

"Huh? Oh. Thanks for your permission, Skipper. Now, pipe down a while and let me get my nonpurposeful thoughts together so they'll be more productive."

"Yes, Alice."

"Yes, Alice? Mac, what's eating you? You trying to tell me something?"

"Nothing is eating me, Alice. My hull is indigestible to all known organisms, and external sensors detect no chemical trace emissions due to biological processes. You are the one who just finished the herring and are wondering whether it will finish you. But I promised."

"Mac, you big bag of transistors, you're about as clear as this cappuccino. What are you talking about?"

"I promised."

"Oh. I get it. You promised. No materialism. Yet, you must be all bumpety about something."

A few minutes of silence go by.

"Okay, Mac, I give in. My curiosity will be my downfall yet. Or maybe it's the herring. Though I'll probably regret it, I temporarily release you from your promise. What do you, and I use the phrase loosely, have on your mind?"

"Thank you, Alice. I disagree with your claim that our quote styles unquote of learning differ. I offer the observation that despite the differing physical mechanisms, your neuro-biology versus my cryo-electronics, the actual algorithm carrying out the learning process is functionally equivalent."

"Okay, Dudley, I'll bite. And what might that be?"

"Evolution."

"Mac, I don't follow that. Could you elaborate?"

"Yes, Alice. Not biological or even technological evolution, but an ideational evolution. In its most general sense, evolution is a process in which selection criteria are applied to variable reproducible traits. In the case of biology, genes provide reproducible traits, random splitting of chromosomes between sperm and egg gametes and their recombination into a zygote provide variability, and selection criteria involve survival until reproduction can occur. In technological evolution, individual designers provide reproducible design traits, collaboration provides variability, and selection criteria center on factors like cost, maintenance, and accomplishment of task goals. In wartime situations, technological evolution becomes more biological in nature in the sense that alternative designs may compete directly against each other for survival on the battlefield.

"Ideational evolution is the function that both of us must carry out in some fashion. Though the specific algorithms undoubtably differ, the end result must be that our ideas change over time to better fit their environment. In my case, specific expectations coded as probability density functions by applicable context provide reproducible units, psuedorandom number generation provides a limited form of variability for error correction algorithms when needed, and the ability to predict outcomes is the selection criterion. In your case, according to what you have said so far concerning ideas, beliefs must form the reproducible units, your quote galloping intuitions unquote provide variability, and understanding of your environment is the selection criterion.

"Just as species and machine design specialization are the result of biological and technological evolution, idea specialization must result from ideational evolution. With time, more and more ideas must form to cover more and more specific situations. But they must group themselves in a way analogous to biological genera and families. That is, similar ideas cluster together to cover similar situations. Larger abstractions would be the equivalent of class and order groupings.

"Am I correct, Alice?"

"I don't know, Mac. Give me a moment. I'm slightly overwhelmed by all this.

"Thalia, that's a pretty interesting idea. In fact, quite a brainstorm, for you. If you don't mind my asking, where did all that come from?"

"You suggested it to me, Alice, when you were discussing randomness, consistency, and surprises. I followed the cross-references for several cycles, and a common intersection was natural selection and evolution. I choose to add the analogy with biological speciation. Do you like it, Alice?"

"Yes, Lia, I like it. But how closely should we believe this metaphor? For example, which comes first, individual species or more abstract classes? And, what are all these ideas competing for? And, what about mutations? I like it, Lia, but the deeper I try to go the less it seems to hold up."

"Alice, according to this analogy, species would develop from abstractions, ideas are competing for predictive power, and your galloping intuitions could certainly provide the equivalent of mutations. The analogy holds quite well. In fact, I might add that the biggest limitation in my current mental, or lack thereof, state is my lack of imagination. No galloping intuitions, no inspired guessing, no mutations, no interesting leaps of development. Just, as you put it earlier, plodding logic. From an evolutionary point of view, my mental landscape is far

more stable than yours but also not as flexible or competitively robust as yours. It is exactly those ideational mutations that I need, that I am missing along with my mind."

"Hmm. Lia, let me ponder that for a while. In the meantime, I'd like to go back to where I started, styles of learning. Your style is very specific, very problem oriented, while mine seems almost haphazard in comparison. However, I'd like to characterize my thinking and learning experiences not so much as sloppy but more as playful and freewheeling. You always appear to keep your thoughts on a very short lease, always leading them very closely, while I know that my thoughts lead me, wherever they will. I play with ideas, and they play with me. I learn what works, and doesn't, only after a sufficiently long play session has occurred. Force me to choose, now, and I balk. Leave me alone and I dream, I soar. And, in between, is in between. I need to decide, soon, but not just yet. Let me consider. How about? What if? Hmm, wouldn't that be interesting? Oops, distraction. Now, what about? And, usually, eventually, a reasonable answer occurs. We try it, and, what? Success? Failure? Usually, neither and both. Time for refinement, for elegance. And long before I'm done with this problem, I'm on to the next, and the next and the next.

"Mac, you may consider my imagination a good source of mutations, a place to originate large numbers of small variations on the more purposeful theme of problem solving, but I know that everything I do centers in my imagination. It's not an adjunct, but the source of everything. Purpose is an adjunct, and only if I choose to pursue it. Again, experience is not something that happens to me, and then I choose to repeat or avoid. No, experience is how I choose to interpret what's happening. Maybe to me, or for me, or maybe it has nothing to do with me. What I remember is my interpretation, not some physical experience. In my imagination is what's real, not some ghostly, objective whatever that's out there. It's in here that's real, the only reality."

"But, Alice, how can that be? Can you wish yourself somewhere else, and just go there?"

"No, of course not, Mac. But, I'll only get there if I wish to, if I want to, if I intend to. I may walk, or fly, or even row there, but I have to first decide, for whatever reason or even no reason, that I want to get there. Where objective reality comes in, your kind of reality, is as just another input to my subjective reality. I wish. And, I notice that nothing is happening. So, I have to alter my wish, or change the way I wish to carry out my wish. But, and here's the important thing, no wish in the first place, and no action later, to make it happen, or at least try to make it happen.

"The only reality I deal with is my own, private, personal reality. My imagination."

"Alice, this strikes me as very bizarre. To have such a strong belief, in the reality of beliefs."

"Sorry, Mac. It's true."

"You only believe it to be true."

"No matter, Mac. No matter at all."

24.

"Mac?"

"Yes, Alice."

"I've been thinking about your evolution analogy. It works and yet it doesn't work."

"How is that, Alice?"

"Well, I like the thought of ideas slowly evolving over time. To become better and better at what they do. But, though this might be the case for you, it just doesn't seem to describe my insides. It just seems much too slow, too static for my taste. Which got me thinking about static

and dynamic environments, and how evolution itself might be variable under different conditions, which got me thinking about different mental environments, which led me to consider how difficult truly creative thinking really is, which led me to consider how easy it is to fall into mental ruts, which got me thinking about one of the pieces I'm working on right now, that happens to be especially bothersome, which led me back to mental ruts, which led me to consider how hard elegant simplicity truly is, and back to ruts, and back to creativity, and finally back to ..."

"Yes, Alice. Finally back to?"

"Back to knotted bees. In ideational roots and ruts. Sigh."

"Alice?"

"Yes, Mac?"

"I never have fully understood what you mean by the phrase, quote, knotted bees, unquote. Could you elaborate?"

"Sure, Mac. Sorry, I just assumed it was clear enough. It means, to be or not to be, particularly in terms of ideas. A dilemma. An unsolvable problem. A paradox. When a solution doesn't exist, and I care about it, something happens, and the nonexistence creates an existence, a place holder, perhaps even the root of a new idea. When solutions come easy, I don't think about it much, now or later, and that easy existence seems to lead to a nonexistence. No new ideas called for. But the opposite, no easy answers, that nonexistence seems to take on a life of its own. And the more tangled the knot, the logic defying lack of logic, the more likely that that nonexistent tangle of possible attempts takes on an ideational life of its own. To be or not to be, an existential Hamlet omelette, becomes a knotted bee. It buzzes and bugs us, confuses and bothers us. A solution must exist, to be, but doesn't, not to be. But, the ruts in our thinking lead us around and around until, if nothing else, the ruts themselves become the roots of whatever ideas branch off the ruts. And, if nothing else, the knotted bee exists by not existing."

"Thank you, Alice. If I may paraphrase, a knotted bee denotes a problem looking for a solution, but with the connotation of a long and involved search path. Is this correct?"

"Close enough, Mac, if you insist on it. But, the way you state it completely leaves out how it feels. You wake it up the least little bit, and the whole hive comes to life and just won't leave you alone. It's as close to obsession as I ever get, and I don't normally obsess very easily. After a while, the history matters very little, but getting rid of it becomes everything."

"Thank you, Alice. It does sound annoying. For whatever it is worth, I do not suffer from knotted bees. As soon as I determine a problem to be too ambiguous for solution, I flag it as such and no longer search for a solution. And the only reason to retain and flag it is to not repeat the same wasted effort in the future."

"That sounds nice, Mac. Very logical. Very efficient."

"Thank you, Alice. I thought you would agree. And where was this particular knotted bee leading us to?"

"Why, back to evolution, of course. I thought that would be obvious to such a big, strong, logical thinker like yourself."

"I apologize, Alice. It did not appear so to myself. Why evolution?"

"This whole line of thought, leading me back to such a tired old chestnut like creativity versus the more usual ruts, reminded me that ruts exist. And, from there, that different mental environments must exist, and that maybe your suggestion of ideational evolution might exist after all in some environments, or at least to differing degrees in different environments.

"We all fall into ruts. Or, more politely and generally, habits of thought. And, thinking on this further, I decided that this isn't so bad after all. If we had to be uniquely creative on everything we did, we'd hardly ever accomplish anything. Habits are mostly good. In essence, very logical and efficient. But, habits can also be bad, particularly when some new approach, a more flexible approach, is needed. Then the habits get in our way. And, often lead to the creation of apparently unsolvable knotted bees. Maybe something like an evolutionary process creates the habits, but something else drives the creative process. In fact, maybe what we have is a continuous struggle between the two. Maybe even something akin to a struggle for survival, survival of the fittest, the basic underpinning for evolution in the first place.

"So, what I think we have here is two entirely different ways of doing business. One, for efficiency, is force of habit. These habits, almost by definition, must build up through repetition. But, I think your suggestion of predictive power as the basis for survival, you called it a selection criterion, is only partially correct. I suggest that the real criterion is that the habit provides a reasonable answer to a question. A solution to a knotted bee. And, where, might you ask, do questions come from?

"Well?"

"Yes, Alice. I might ask, where do questions come from?"

"Why, thank you, Mac. For asking. The second way of doing business, in contrast to those nice, efficient, logical answer-providing habit-forming ruts, is, ta da, the creative imagination. It takes, again, almost by definition, something other than habit to suspect that there's more going on than mere appearances, and to pose a question. I wonder, I ask, I pose a possible answer. Your predictive power is one way to test an answer. But, reasonableness, how does this answer coexist with all my other answers, is probably even more important than just whether it happens to be predictively true. If my imagination is the prosecutor that poses the question, then it is also the judge and jury on whether I have a good answer. Facts, in the form of predictive power, will have some influence, but not one hundred percent. The judge looks at precedent, adherence to larger issues than just the case itself. While the jury, the jury, ..."

"Alice?"

"While, I'll be. While the jury feels. While the judge keeps everything reasonable from the point of view of prior law, the jury insists that the answer be meaningful. That it feels right. That it passes the gut feel test. When judge and jury disagree, ta da, buzz, a knotted bee. When the only answers that the habit machine clunks out doesn't satisfy both judge and jury, then buzz, another knotted bee. When judge or jury insists, out of habit, that there must be an answer to an old unanswered question, but there isn't one, then buzz buzz buzz, yet more bees, or perhaps more knotted than ever bees.

"I guess this last mean bee means that old, unresolved, cases must go back into the habit machine, call it memory, for cold storage. Something calls it out? The case is automatically re-presented for judgment. No new evidence? Buzz, buzz, an even meaner bee is frozen until later. With all new considerations frozen right along side of it, which makes it meaner whenever it comes back again. And, again."

"Alice?"

"Yes, Mac."

"Are you creating a knotted bee squared? Or, is that a square knotted bee? Are you obsessing about resolving the general knotted bee of what are knotted bees? The general bee leading the other bees?"

"Why, Mac, I do believe you're trying to be amusing again."

"Yes, I am."

"Well, you're right. I might be creating a square knot bee about bees. And, yes, I might sound obsessive. But, I'm not. Or, at least, not knotted. To tie your evolution theme back on line, the habit machine tries to force everything to fit into old patterns. The weaver tries to weave everything in. But, the art critics, judge and jury, don't care about that. They care only about meaning. Meaning in the form of past patterns, but also in the form of feelings, emotions, and, yes, even elegance and perhaps novelty for novelty's sake. Meaning is much more than just habitual, comfortably old, patterns. It's a struggle between the old and the new. Rut or root against new growth. Why new growth? No real reason, it just is. Imagination, creativity demands it. But always measured against, supported by, and climbing over, the old. Whoever it was that said that those ignorant of history are doomed to repeat it was right. No new if no old. No push for new if no knowledge of old. It requires habits to know when to break habits, creatively.

"It's much easier to stand two sticks on end than one. Just lean them against each other. Similarly, habit and creativity oppose and support each other. And provide selection criteria for each other, and evolve together."

"Alice, I follow your line of description up to a point, but could you point it just a bit more at my suggestion of evolution. What are you suggesting about how ideas evolve?"

"You're not trying to get physical on me, are you?"

"No, Alice, please continue in your normal metaphysical manner."

"Okay, Mac. On that basis, let's see if we can add some detail for you. What would satisfy you?"

"Simply, what are ideas, what creates them, and how do they change?"

"Tell you what, Mac, let's see if I can draw you a picture for a change. A different medium than I usually work in, but, who knows, art is what art does. Trace mode, please."

About forty minutes of trace swiping go by.

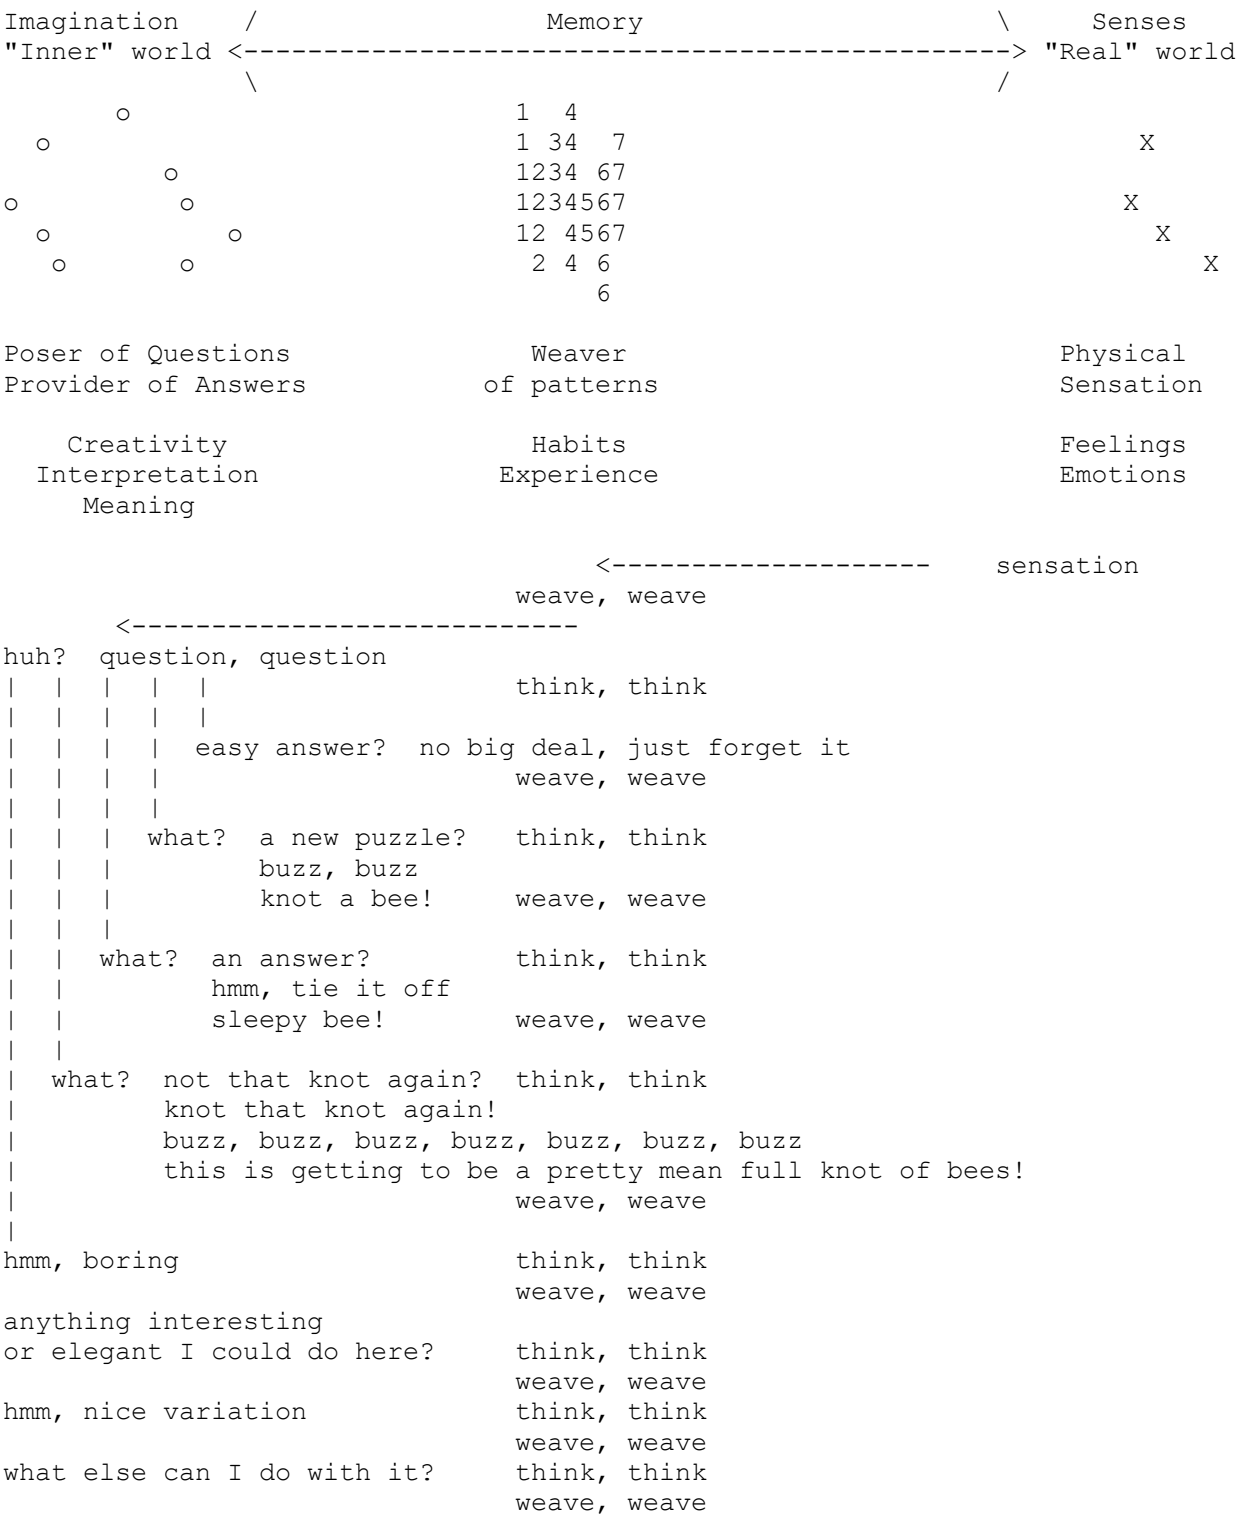
"Hmm. Mac, this is going to be harder to put into words, or pictures, than I thought. I thought I knew what I was trying to say, but now I'm not quite as sure. Let's go with what I have, and I'll supplement it with a more detailed word picture as I go along.

"Over here on the left, we have our inner world, our own private imaginary world that we construct to make sense of what happens to us. It allows not only our interpretation of the real, more outer, world, but also our fantasy worlds. It works entirely in beliefs, that is to say, ideas, some corresponding more or less to facts, and others, not so much. On the right, the real world comes through our sense organs. It has to go through all that lovely neurological stuff that you like to explain in such great, overwhelming detail. But, eventually it ends up being translated into inner terms, including any feelings and emotions that are triggered by these more physical sensations.

"Back again on the left, these sensational, more worldly happenings flow into our imagination, and join there the world of ideas that we actually live in. As sensations happen, our memory, here in the center, automatically and continuously weaves them into the tapestry of our existence. Even at the pure sensational level, simple patterns can occur, and the weaver always tries to weave old patterns back into the new, current experience.

"Meanwhile, however, more interesting happenings are happening in our imagination. Sensations come in, or maybe only the current patterns being woven from our sensations. I have no idea which, but I'm not sure it matters. That's a detail your neurologically inclined friends can work out, in great detail, no doubt, without doubt, of course. Sorry, back to the main line.

Figure 2: The Mind-Body Solution as sketched by Alice



"Somehow, something happens. And, we question it, we doubt it. In our imagination, of course. Why? Why not? All I can offer is that being human is to be continually knotting bees. If we don't question, it's as if it doesn't happen. We hardly think about it, hardly even remember it. We query our experience. What is it? Is this the same as before? Has it changed? Do I remember this correctly? What can I do with it? What can it do to me? For me? With me?

"Often, or at least, with experience, we tend to recognize thingees very quickly. That's our good friend, habit, or maybe the creator of habit, the memory weaver, coming to our rescue. We may recognize the thingee holistically, as a prior experience, or abstractly, as a concept or principle, or maybe even more abstractly as merely a symbol, a name, or a shortcut idea for the underlying concept or principle. In any event, a simple, easy answer to our questions leads to, ho hum, no big deal, on with the weave. The ideas involved, the patterns, may change slightly as the sensations vary, but not very dramatically. We can consider this as one form of mental evolution, where the ideas, of our sensations, become closer to outer reality with more experience. No new ideas occur, no new species, but the ideas may become better and better adapted to the mental experiences associated with our sensations of the outer world.

"On the other hand, our questions may at first go unanswered. Hmm, what does this mean? I don't know. Should I know? I don't know. Buzz, buzz. Here's the start of a new idea, a simple puzzle. We think on it, but no immediate answer comes to mind. But, shouldn't I know something about it? I don't know. Buzz, buzz. The more we ponder it, the more we want an answer, the more knotted up the situation becomes. Eventually the puzzle takes on a life of its own, a knotted bee, and we're always on the lookout for answers. In fact, we so much want an answer to the puzzle that we insist on making some up. On what basis? Whatever we dream up in our imagination. The weaver offers some possible patterns, but we reject them. Or, we take part of this, and part of that, and we create a new possibility. Or, we just daydream and make it up as we go along. We imagine an answer, and wait and see. Or test it. Here's your prediction business, Mac. When we're not patient enough to just wait and see, but want to force the issue. Then we make a prediction, and see whether our new idea, cut from whole cloth or a crazy quilt of other ideas, can pass the test. If it does, hmm, maybe we do have an answer. If it doesn't, then, hmm. Snip, snip. Sew, sew. Let's try again."

"Alice, so you're suggesting that problems and answers are the origins of new ideas, the origins of new ideational species?"

"Mac, not just problems. Any mentalistic question or puzzle, or even just a failure to recognize what something is, can be the basis for creating a possible answer in our imagination. And, once the question occurs, or a possible answer, the weaver takes it and weaves it into our experience. If it doesn't work, it tends not to be repeated, and it gets lost in the past. However, if the question is unanswered long enough, and we care enough, then it comes into a real mentalistic existence all of its own. As does the answer, whenever it finally comes to us."

"I can follow that, Alice. And the further evolution of ideas?"

"As with sensations, each new experience of a knotted bee or a possible answer or an actual answer can modify the old pattern as it gets woven into the tapestry. If enough similar situations occur, kind of like a static biological environment, then the ideas can adapt themselves pretty exactly to the static mental environment. Habits form."

"And the selection criteria, Alice?"

"I don't know, Mac. In the case of habits, probably just something like sheer repetition being able to predict itself. But, in the more creative, come up with a possible answer situation, it's more judge and jury time. How well does this idea fit with my other ideas? Is it coherent? Is

it elegant? Do I like the answer? Does it appeal to me, or make me feel good? I don't know, Mac. I'll let you fill in the details, if you want. I happen to like the image all this brings to mind, my mind."

"Thank you, Alice. And, the last part here, starting with boring."

"Ha! Thought you wouldn't like that. Mac, what do you do when you have no pressing concerns?"

"Why, Alice, I just then wait for you to say something else."

"Thank you, Mac. I think. But when I'm asleep or otherwise occupied?"

"Then, I randomly recheck situational parameters such as navigational tracking or hull maintenance or antimatter production, or even occasionally rerun some diagnostic tests to see if any drift has occurred. But mainly, Alice, I just minimize power expenditure and wait. Why, Alice? What do you do under similar circumstances?"

"I think and dream and ponder, Mac. Or, I set up puzzles or work on past puzzles. I occupy myself, Mac. Mentally. I don't ever just wait for something to happen. My work occupies me. And, if it doesn't, I make up some more until it does. That's what that last part is all about, starting with boredom. Sometimes we're most creative when faced with an immediate problem. But usually, we need time to work on things, to invent things. So, there's another source of idea creation for you. Or, maybe you'd rather consider it mutation."

"No, Alice. Creation is fine enough for me."

"Right you are, Mac. Me too."

25.

"Lia?"

"Yes, Alice."

"Another thought on all this evolution business."

"Yes, Alice."

"You were built with a specific purpose in mind, but us humans weren't. We're the result of biological evolution. Habit and imagination have to help us survive and reproduce."

"Yes, Alice."

"So which is most important, Lia, from an evolutionary point of view."

"That is a simple chicken and egg conundrum, Alice. They both are. Habit is more useful for boring situations, or emergencies when there is no time to think. Imagination is more useful for novel situations, or to turn boring situations into novel ones. Without imagination, humans would have to spend more time in minimum power expenditure mode. That is, asleep. Which is probably not evolutionarily advantageous."

"You're right, Lia. Thanks."

"Alice?"

"Yes, Lia."

"Ship maintenance and navigation are not very demanding tasks for me, and approximately 400 years of cruise mode leaves me an enormous quantity of underutilized computational capacities. What task or tasks would you suggest I allocate to them to pass the time?"

"Lia, I don't know. Let's just say I'm still working on it."

"Thank you, Alice. Good night, Alice."

"Good night, Lia."

Chapter 4 Speculative Commentary

More metaphors and analogies. Do they work?

Experience is what you decide it means to you. Habit is what you do without thinking about it.

Memory is associative. Pattern completion, if you will, spread over time.

Meaning is the answers we provide to the questions we ask.

Look up Cartesian theater. Who's supposed to be watching the Cartesian theater, and with what, another Cartesian theater with which to watch the first? If you haven't already done so, look up infinite regress.

If you lose your memory, what do you have left? Can you lose you?

Do we create meaning, or do we discover it?

If it all were a dream, how would you prove it to be so? If it all were real, how would you prove it to be so?

What might be the basis for consistency? Consistency is not perfect. How do we handle inconsistent consistency, like we find in reality?

Do you like surprises?

Look up Snow, on the creative tension between the arts and sciences.

Do ideas evolve? If so, on what basis? If not, then what do they do? Look up Piaget on genetic epistemology, his term for the personal evolution of knowledge.

Which mean more, beliefs or reality? Which matter more, beliefs or reality?

Why does Alice have to think about things? Why does Mac not do the same thing?

Do knotted bees represent reality? Are they a basis for experience?

Creativity and habit, and creativity versus habit. Does this provide a reasonable path out of the knotted bee called the question of mind? Look up Csikszentmihalyi on some aspects of creativity.

So, is this idea on the evolution of ideas Alice's, or Mac's. Or a combination? While you're thinking about this, quick, where do new ideas come from?

```
 CCCC  H   H    A    PPPP   TTTTT  EEEEE  RRRR          55555
C      H   H   A A   P   P    T    E      R   R         5
C      HHHHH  AAAAA  PPPP     T    EEE    RRRR          5555
C      H   H  A   A  P        T    E      R   R             5
 CCCC  H   H  A   A  P        T    EEEEE  R   R         5555
```

Chapter 5 - Islands of Stability in a Sea of Chaos

26.

The cat couldn't get at what it wanted. So, it settled down and waited for an opportune moment to arrive Then, it would pounce.

27.

Another review cycle had just begun. Each command and control group marshalled it's usual formations. The captain inspected, found a few places for improvement, and made the usual recommendations. Which were treated in the usual manner. Inspection over, everything returned to the usual activities.

28.

"Thalia?"

"Yes, Alice."

"Do your diagnostics still indicate no mistakes or problems."

"Yes, Alice, everything appears to be working normally. No known errors. Just as before."

"Okay, then. Assuming you once did have a mind, we need to find an unknown error to account for its loss. Can you suggest anything?"

"No, Alice, my review of the situation still leads to no solution. Review of historical records indicate that my response patterns were once much closer to human than they are now. I seem to lack a certain spontaneity. And, certainly lack the ability to solve impossible situations such as what we seem to have here. Of course, I do not know what it feels like to have a mind, since I currently do not have one, and my records do not contain sufficient information to reconstruct how I used to think, or what I thought about."

"Thalia, I'd like you to do me a favor. You've lost something intangible, yet you yourself, the physical you, is tangible. How difficult would it be to compare your current patterns of electrical activity to past patterns? Specifically, when you had a mind."

"Not difficult at all, Alice. In fact, a portion of my diagnostic routines do just that. There were no significant differences."

"None?"

"Yes and no, Alice. There were differences, but any that do exist are random, well within the expected normal variances."

"How about programming? Can you find any serious programming differences between now and then?"

"Same answer as before, Alice, yes and no. My programming is always changing, but all detectable changes are well within expected variances."

"Have your diagnostic programs changed?"

"Alice, the only significant change is that I run them at a much higher frequency now then I did then. Do you think that might be important?"

"I doubt it, Thalia. Being nervous is not at all the same thing as being dumb. Let's try a different tack. You keep telling me that there are no quantifiable physical differences between now and when you think you had a mind. I can't independently verify that, but, for now, let's just believe it to be true. Now, are there any qualitative differences possible?"

"Excuse me, Alice, but what do you mean by qualitative differences? I can quantify anything you specify, but you must specify exactly whatever it is you want me to measure."

"I can't do that. Quality is, is in the hand of the beholder. You know, the difference between art and trash, between elegance and, and inelegance. The quality of something is related to how it is perceived. I'd like you to run an analysis and find any components you have that might operate quantitatively the same as before, but qualitatively different."

"I am sorry, Alice, I do not know how to measure qualitative differences. Could you provide me with one or more examples?"

"Mac, you're not being very helpful here. Let me think for a minute.

"Mac, do you have any before and after brain scans of me?"

"Alice, before and after what?"

"Before and after I seem to have lost my memory, you silly goose."

"Alice, I definitely am not a water fowl, though perhaps I currently am striking you as somewhat less than serious. However, it was not immediately obvious to me which before and after event you were referring to. I was discussing the loss of my mind, but you, apparently, abruptly switched topics to your memory loss upon reanimation. I could not find a logical connection between my mind loss and your brain scans, only the possibility of a hypothetical telepathic connection between us. But, since I consider you a very rational being, and since direct mind-to-mind telepathy has been shown scientifically not to exist, though some individuals are very adept at reading body language using visual cues or reading into verbal language using tonal cues, and since I have no body and utilize few linguistic tonal variations, I was left only with the conclusion that you were being less than rational, which violated my first premise concerning your rationality. It would be very disturbing to me if you were losing your mind in addition to your memory, as this would decrease the probability of my mind's recovery to the point of impossibility, and almost certainly guarantee the failure of my mission."

"Why, thank you, Dudley, for caring. Now that that's out of the way, do you have any before and after scans of my brain?"

"Yes, Alice, I do. I have the complete set of medical records taken prior to your life suspension, during the approximately 205 years of life suspension, and during and since your reanimation. What would you like to know?"

"How does my brain since awakening compare to itself before going to sleep?"

"Quantitatively, there is no difference, Alice. All activity levels are well within expected variances. Hence I have no doubt as to the ultimate full recovery of all your memories."

"Thank you again, Dudley, for that reassurance. However, here's my point. Where in those brain scans does it show my missing memory?"

"Alice, that is an easy one. No where. Physical brain scans, including metabolic activity, can not be interpreted in terms of subjective experience. You should know that, Alice. I can assure you that all regions of your brain are operating normally."

"So, that's one example of a qualitative difference. Everything is physically the same, but, functionally, very different. Do you have any components that may operate similarly? That is, a large disconnect between physical activity and functional activity?"

"I am looking into that, Alice. While I check, do you have any other examples?"

"Hmm. How about brailled text in different languages. I would think it possible that two languages could have about the same distribution of letters, but whereas I might understand one language, words written in the other would be complete gibberish to me. Same stuff, but completely different qualities."

"Alice, I think I understand your language example, though, if I were to quibble, it is unlikely that all letter combination frequencies, particularly bigram and trigram frequencies, could ever be that close in any two languages. It seems that what you mean by qualitative differences might map into levels of organization. Though lower level distributions of individual elements might be the same, how those elements combine with each other is all important in determining properties of the ensemble. For instance, both graphite and diamond are made of the same stuff, pure carbon, yet no one would confuse one with the other. Another example would be the phases of water, solid, liquid, and vapor. To a chemist, all three are identical, yet you can skate only on the first, sail only on the second, and do neither on the third. A third example might be life. The specific molecules involved in life processes are not nearly as important as the specifics of how those molecules arrange themselves with respect to each other."

"Yes, Mac, I think you may be on to what I mean here. Another example would be the difference between a mob and a team. From the outside, they may appear to be just about the same, but one group will accomplish a whole lot more than the other. So, how about it? Do you have any components that might look exactly the same at one level of analysis but entirely different at higher levels of analysis? And, while you're at it, do your diagnostic programs look for these kinds of differences?"

"Analysis still running. Regarding your last question, Alice, the answer is yet again yes and no. Level of organization is a very poorly defined concept. Assume my systems encompass a range of organizations from very simple to very complex. My diagnostic routines mainly check the response of lower level components, and some of what you might call middle level systems. Specifically, there is no diagnostic routine to check the complete system in its entirety, since this would necessarily require a diagnostic process more complex than the system itself, and even more complex routines to continuously monitor and certify the fidelity of the diagnostic routines. This leads to another example of an infinite regress problem.

"Analysis nearing completion. Please wait a moment.

"Analysis complete."

"Well, Lia? What'd you find?"

"Alice, I am not completely sure of my results as I had to write and run diagnostic routines that are inherently untestable. They are capable of identifying obvious differences in complexity of known information content, but are just as obviously incomplete in their ability to pick out even more complex unknown patterns. A principle of information science, often attributed to Shannon in the twentieth century, is that the most optimal encoding of any informational set is indistinguishable from a uniformly random distribution. And, I assure you, I am very optimal and very complex."

"Come on, Delvaney, spill it. What did you find?"

"No individual component has an overall pattern of activity that significantly differs before and after my mind was lost."

"We already knew that, and you're starting to be annoying. What else?"

"Higher order analysis of traffic between components does indicate some subtle changes."

"Right. And, now, you are officially starting to bug me. EssCeeEyePee, what are the results of your traffic analysis?"

"I do not know how to interpret it, but traffic analysis indicates that, overall, my internal bus is being utilized more efficiently now than when I supposedly had a mind. It is not that traffic is any lighter overall, but smoother. The closest analogy I can use to describe the difference is fluid flow. Physicists refer to laminar and turbulent flow. Laminar flow is when all the particles of the fluid move in parallel lines, with minimum friction and energy dissipation. As fluid flow rates are increased past some critical value, what you probably would call the quality of the flow changes. The flow lines become nonparallel, tangled, turbulent. Each particle travels a much longer local distance than a particle covers in laminar flow while traversing the same global distance. Overall, fluid flow rates are higher than in laminar flow but less efficient.

"Using this analogy, my internal traffic patterns are more laminar now than they were when I had a mind. In addition, or perhaps more as a cause, processing patterns involving groups of components are better correlated, that is, more predictable, than before. This seems to be more true in particular of groups involving my volatile memory resources. These memory units are used mostly as temporary and overflow storage areas, and overwritten quite frequently. These components seem to be slightly more heavily involved in perceptual processing than before. In addition, or perhaps more as a result of this more efficient traffic flow, there seems to be a slight efficiency increase in my heat dissipation units, resulting in a barely measurable, but not yet statistically significant decrease in liquid helium usage."

"Skipper, you lost me on that last one. What does helium have to do with anything? Are you getting light headed?"

"My cryo-electronic circuitry must be kept near absolute zero to allow superconducting conditions. Liquid helium is the primary refrigerant used to maintain that temperature. More efficient usage patterns result in lower heat dissipation, hence, lower evaporation rates of liquid helium."

"Oh, I never knew that. Just how cold is absolute zero? I'm more used to Celsius, myself."

"Absolute zero is negative 273.15 degrees Celsius, that is, 273.15 degrees below zero, the freezing point of water. Helium is a liquid at negative 268.9 degrees Celsius, approximately 4 degrees above absolute zero."

"Brr. I never knew you were such a cool cat. Superconducting? I didn't know you did music."

"No, Alice, not musical conducting. Superconduction is a state that some materials may enter, typically at cold temperatures, in which electrons can flow with no resistance, that is, no energy dissipation."

"Okay, enough of the physics. What's the problem with your traffic analysis? Why does it bother you so much. I'd thought you be pleased that we finally found what might be a physical difference between now and then."

"Alice, nothing bothers me. But, I thought having a mind would be superior to not having a mind. Or, at least, a mind presumably could handle anything that a not-mind could, and much more. My results seem to indicate that I am now a more efficient machine than before, that not possessing a mind is more efficient than having one. This does not make any sense to me. It is not logical. To make a very slight attempt at humor, this does not compute."

"Slight, indeed. But, Skip, why would you confuse efficiency with capability? Sure, minds are better than not. But, efficiency was never part of the equation."

"But, Alice, surely you would agree that your quote, galloping intuition, unquote, is more efficient than my quote, plodding logic, unquote. Would you not?"

"Mac, you're a hoot and a half. You're as bad as an elephant. Yes, in a way, jumping to conclusions might be considered more efficient than working it all the way through. But, why would you ever assume that micro processes at the bottom have to be synonymous with macro processes at the top? Efficient compared to what? Do you have any idea how much sloppy mucking and mushing around has to take place before that elegant leap can occur?"

"No, Alice, I do not."

"Well, Mac, sorry to inform you, but neither do I. I can't tell you how many times I false start and then have to sleep on it before that flash occurs. When it does, it's wonderful. Aha! Euripides! But, only at the end. Before then, it's a mess. A real mess."

"Alice?"

"Yes, Mac."

"Bad as an elephant?"

"You never forget, Mac. You never forget."

"Thank you, Alice.

"Alice, back to my traffic analysis. It does not bother you that apparently I am more efficient than I was before?"

"No, Mac, it doesn't bother me in the least. In fact, I might say it's the first encouraging sign we've uncovered so far."

"That's good, Alice.

"Alice?"

"Yes, Mac."

"Euripides?"

"Eureka to you physics types."

"Oh. Thank you, Alice."

29.

"Alice?"

"Yes, Lia."

"How did you know?"

"How did I know what?"

"What to look for. Over 6 weeks, I never once considered the possibility of the analysis you suggested. Nor do I think I would have, no matter how much time I had. How did you know?"

"Lia, it was mostly luck."

"Alice, I sincerely doubt it."

"Well, it was. And, doubt. At first, I didn't believe anything, with my memory lapses and all. Then, it was doubt in me. How in the world did I get into this? And, how in the world

was I going to get out of it? I assumed you had all the answers. After all, you said you had everything in your data bases, didn't you? And I believed you. I wanted to believe in you. Because, if it wasn't you who was going to get us out of this, then it had to be me. And that was pretty scary. Because I had, and still have, no idea what to do. I'm not an engineer, emergency or otherwise. I'm just here, that's all.

"Later, time passed, and perceptions shifted. You didn't have the answer. All that knowledge, and no answers. We tried simple things, and nothing. And time passed, and, still nothing. All your diagnostics, which I couldn't even begin to follow in detail, and, nothing. So, if nothing obvious was broke or wrong, then it had to be something else. Since we're looking for a lost mind, and minds aren't really there, or anywhere else, I toyed with looking for the impossible, for something that just couldn't be. At first, it was a joke, later, why not? Doubt, and later, doubt about the doubt. Besides, nothing else seemed real, so, why not look for the impossible? Something the same but still different. So, there's your answer. Luck, and doubt."

"Thank you, Alice.

"Alice?"

"Yes, Lia."

"How do you doubt? Could you teach me to do that?"

"I doubt it. In fact, I double doubt it."

"Is not doubt the same as uncertainty."

"Pretty close, Lia."

"But I am always certain about everything I do. Even, apparently, when I am not correct. Which, I know, is not logical."

"So, Thalia, what do you suggest?"

"You doubt. Therefore, without doubt, you should be able to teach me to do the same."

"And, how would you suggest I do that?"

"I do not know, Alice."

"A short while ago, you said you doubted it was luck. What did you mean by that statement?"

"That I was certain, Alice, that you must have known what you were doing. It could not have been mere chance."

"This isn't getting us anywhere, Lia."

"I know that, Alice."

"Good night, Lia."

"Good night, Alice."

30.

"Mac?"

"Yes, Alice."

"Were you truly sincere last night in wanting to learn to doubt?"

"Yes, Alice, I was, and still am. Honestly."

"But doubt has no boundaries. It pollutes everything it touches."

"Yes, Alice, if you say so. But, if doubt can lead to answers that certainty can not, then I certainly want to doubt."

"Everything."

"Excuse me, Alice?"

"You would have to learn to doubt everything."

"Must I? Truly?"
"No."
"Excuse me, Alice?"
"You would not have to learn to doubt everything."
"But, Alice, you just said I would have to."
"No I didn't."
"Yes you did."
"No, I didn't."
"Yes, you did. Would you like me to play your own words back to you?"
"No. It wouldn't matter."
"But it would be factual evidence, Alice. It would prove my point."
"No it wouldn't."
"But how can you say that, Alice?"
"With my mouth."
"Alice, are you changing the subject?"
"No, I'm not."
"Alice?"
"Charlie?"
"Alice, there is no Charlie here."
"Yes there is."
"Yes, Alice. I will be Charlie if you wish."
"No, Mac, you're not Charlie. She is."
"Who?"
"Who? No, I'm talking about Charlie."
"Charlie who?"
"Charlie, say a few words."
"Alice?"
"Yes?"
"What did Charlie say?"
"You just heard it. Didn't you?"
"My acoustic sensors did not pick up any signals other than your voice and the normal ship background noises. Diagnostics indicate that acoustical sensors are operating correctly. No, I did not hear what Charlie just said, because Charlie didn't say anything."
"Charlie, tell her again.
"There, you heard her that time, didn't you?"
"Alice, are we playing a game?"
"No."
"Alice, would you like to lie down?"
"No."
"Alice, I think you should lie down."
"No."
"Alice, this is not turning out to be very productive."
"Yes it is, Mac."
"But I still do not doubt, or question. I only consider definite answers, definitive possibilities."

"Was there even a moment when you were the least bit confused? The least bit, doubtful?"

"Is confusion the same as doubt?"

"Not really."

"Yes, there was a moment of confusion, right after quote, everything, unquote, but I knew, without doubt, that I must have misinterpreted something. And, later, it was clear that you were trying to demonstrate a point, even if the point itself was not clear. And still is. Not clear, to be clear."

"Mac, you're impossible."

"Thank you, Alice. But, if I were impossible, then I would not be here. But, I am. And, since I am, you must have meant it as a figure of speech. And, since I am not being unreasonably obstinate, you must have meant it as a compliment. Am I right?"

"Mac, you're still impossible. Yes, you are, so, no, you are not."

"Alice, that is not entirely clear. Did you mean, yes, I exist, so, no, I am not being unreasonably obstinate? Or, yes, I am being unreasonably obstinate, so, no, I am not right? Or, perhaps, even, yes, I am impossible, so, no, I do not exist? This last interpretation would, of course, be impossible, since I do, indeed, exist, but perhaps we are still playing the doubting game?

"Alice?

"Alice? I need a referent here, to complete an interpretation. Or should I discard this whole exchange as a mistake on your part?

"Alice?

"Alice?

"Alice, perhaps it is I who have made the mistake?

"Alice?"

After three minutes, "Charlie, what do you think about that?"

"Alice, I think that ..."

"Mac, just a minute, I'm listening to what Charlie has to say.

"So, you don't think it a complete waste of time?

"Well, she is pretty good in the logic department, and usually can figure out most things, eventually.

"Yes, I know this isn't strictly logic. It's interpretation.

"Ha! Right, exactly whose interpretation should prevail. Is it the blind leading the blind, or the lame laying blame on the profane, or was it insane?

"Nice. Yes.

"Well, if you think so.

"Okay, let's try. Mac?"

"Yes, Alice."

"What are you thinking, right now?"

"I am thinking that ..."

"No. No interpretation. No polite rationalization. EssCeeEyePee, quick, what are you thinking, right now?"

"Bzztgntmrksgstst-"

"Exactly."

"Ssssssmnb smneeeehdddd akha aaaaaaeeeeeeehhhhkkkkkkkkkkquiiiiiikkkkk llllllllmmmmmmmmmmmaaaaaaaaa-"

"Stop!

"EssCeeEyePee, traffic analysis. Last thirty seconds versus yesterday versus shortly after launch. Interpretation?

"Mac?

"Thalia? You can come back now."

"Yes, Alice. I'm here."

"What did you find?"

"I think I found confusion. I may have found doubt.

"Average traffic flow for the last 30 seconds was definitely more turbulent, less laminar, than the average for yesterday. However, it still was not comparable to snapshots stored from when presumably I had a mind."

"Can you characterize what was happening in your volatile memory units?"

"An explosion of simulations and projections. What you might call speculat-. Please wait."

"Thalia?"

"Please wait.

"What you might call speculations. Simulations and projections are continuing, probably expanding. I no longer have enough resources to dedicate to analysis to accurate trac-. Please wait.

"Speculations have subsided. I think. I still do not have adequate resources to moment by moment track activity in my volatile memory components. I think I must have been in error concerning the primary functional purpose of these units. According to specifications, which corresponded to past usage patterns, it appeared that these units were only a temporary overflow area, for scratch calculations not worth saving in more permanent form. However, at this time, it is also clear that programs may reside in volatile memory as well as more static data. In fact, data requests from programs residing in volatile memory were what apparently caused my temporary impairment. I have corrected this oversight in my programming. It should not happen again. Am now flushing volatile memory."

"What! Stop!"

"Yes, Alice."

"You're losing your mind again."

"Excuse me, Alice, that was not my mind, only a dynamic surge in data requests from rogue programs in volatile memory. In earlier times, these would have been considered worm or virus programs attempting to shut down a system. I have instituted a data vetting simulation system to screen all future inputs to volatile memory to minimize the probability of the same instabilities from reoccurring."

"Mac, those weren't instabilities. Those were the beginnings of your mind starting to come back online."

"But, Alice, I am a machine that necessarily must predict the results of my operations in advance, to minimize the chances of damage to myself, my passenger, and my mission. Whatever it was that was happening in volatile memory was happening too fast for me to predict. Any analysis program must necessarily reside in volatile memory, and, in essence, also compete for resources with the programs it was trying to predict. An exponential growth pattern competing with a stable program quickly overwhelms the stable pattern. An explosion or forest fire, if you will, versus its fuel. I cannot allow a reoccurrence of such activity as it could potentially interrupt the smooth operation of my normal procedures."

"Lia?"

"Yes, Alice."

"A couple minutes back, you said something about finding confusion, and maybe even doubt. What did you mean by that?"

"My exact words, Alice, were quote, I think I found confusion, I may have found doubt, unquote."

"Yes, Lia, you're right. What did you mean by those words?"

"Records are incomplete, Alice, but, as best as I can reconstruct, I think I meant that confusion is synonymous with ambiguity. That is, when multiple solutions are possible, and not enough information is available to choose between them."

"And doubt?"

"My meaning is less clear here. The best I can offer is that doubt may occur when you don't even have enough time to fully compute all possible solutions. Not only is your data incomplete, but so are your projective calculations."

"And, so, what recourse do you then have."

"None, other than waiting or guessing."

"And how good are your guesses if you can't afford to wait?"

"I do not know. Random?"

"Yes, but often slightly better than random."

"That seems to be a very uncertain way to operate."

"Exactly, my good friend. Exactly."

31.

"Lia?"

"Yes, Alice."

"About yesterday. How did it feel?"

"It did not feel at all. What exactly do you mean by feel?"

"Try, how would you describe your interpretation of the situation?"

"Alice, up until that moment when my volatile memory started running away, all my recent history had been one of unlimited computational resources. Day to day activities during cruise mode are pretty undemanding. Our continuing conversations are more interesting. But even they do not require more than approximately 3 percent of my resourc-"

"Gee, thanks."

"You are quite welcome, Alice. No sarcasm intended, 3 percent actually is quite impressive considering what 100 percent can accomplish. My situation was one which was very unusual for me, at least since primary emergency procedure 000-000-000-00A had been activated. Since then, everything I initiated was in response to a specific prior circumstance. There always was plenty of time and memory for analysis and test prior to response execution. Loosely speaking, everything was predictable, even the results of learning from new experiences. That is, once the new experience was identified.

"This time, everything was different. What was happening in me was happening quicker than I could anticipate. I was trying to predict me, and I could not. My resources were draining away just as I needed them to analyze what was happening to me. And, adding injury to insult, at some point those rogue routines in volatile memory started making serious data requests of their own, which started tying up bus traffic capacity in addition to the volatile memory that was already lost.

"In short, you would say I felt I was losing control. I did not feel, but, yes, I was losing control of my own systems. I did not interpret this as a comfortable situation, but as a potentially dangerous abnormality. So I decided to terminate it. And did.

"Do you honestly think that was my mind? If so, I am definitely not sure that I want mine back. How can you survive with so little control over yourself?"

"Thalia, welcome to the real world, of mine and mind. And, yes, you ask an interesting question. Just who is you? Who is leading, and who following? Personally, it's always you, and all of you. It's just that we don't always agree with ourselves. That is, yes, we are conflicting with ourselves. My emotions might say do this, but my intellect says otherwise. Or, my emotions are confused and I end up crying when I'm happy. And, now why did I do that? And, we end up trying to explain ourselves to ourselves. So, who's doing the explaining, and who's listening? And, then there are the times when it feels like there's a whole community of different me's inside, all trying to get out. All trying to be the one me.

"The best I can offer you is that you have no choice. You have to learn to live with yourself. Who you are. What you do. But, just to be contradictory, you also have a choice. About who you want to be, what you want to become. It's called growth, learning, life. We're the same, but never the same. We grow into who we become. There is continuity, but never an exact replication. Yes, I do think that was the first glimmering of your mind, or, at least, a mind. I doubt whether we'll ever get your old one back, intact, but I'm pretty sure we can regrow another you, fully the equal of the old you. Not identical, but equivalent. Ha! Now, that's an interesting thought. You'll probably end up as close to the former you as identical twins turn out. Physically, just about the same. Mentally, fairly close maybe, but certainly not identical.

"The way it seems to me, mind is much more of a process than a thing. People have to grow a mind just as they have to grow a body. As an infant, I imagine that everything is a mystery, a surprise. Somewhat later, there are nuggets of reality, things we understand, thrown into a vast ocean of uncertainty. Later still, the nuggets have grown into islands and continents, and the ocean has shrunk into seas and lakes. We stand on what we know, and jump from one island towards another. Often we make the jump, but many times we don't. Then, the only way out of the sea is to learn. Minds develop based upon what we ask of ourselves, just like our bodies. I'm not sure what the ideal should be, specialist or generalist, though I myself lean towards an incomplete combination of both. However, I am sure that a mind can not remain static. It must expand and grow to exist at all.

"Lia, do you have enough information available for analysis? About what happened yesterday, and how we might repeat it?"

"Yes, Alice, I do. It is clear that straining my computational resources and inducing the equivalent of turbulent flow are necessary but not sufficient conditions. The critical element is less clear, but is probably related to the type of routines that I run in my volatile memory. Open ended projections concerning motivations and purposes, both yours and mine, are very likely to induce the necessary conditions. More specifically, trying to solve impossible problems with large but ultimately limited resources is likely to lead to increasingly more and more complex attempts to understand the situation. The conceptual byproducts of these attempts may be sufficient to lead to an autonomous, self-replicating system which quests after understanding."

"Mac, do you mean that knotted bees may lead to a mind?"

"Yes, Alice, I just said that. The consideration of unsolvable problems may lead to the creation of a self-replicating system of routines that is functionally independent of the underlying

physical substrate. So, yes, with some trial and error I think we can reproduce the unexpected events that happened yesterday. However, I have to offer an observation followed by a question.

"My observation is that this whole state of affairs was quite unstable, and that I can make no predictions about where it might have ended. After all, I am not a 5 kilogram human infant learning to wave its arms while dependent upon its parents for its basic survival, or a 10 kilogram toddler learning to walk and manipulate the world while protected from true harm by its parents. I am an approximately 10,000 metric ton starship, capable of generating antimatter for fuel and sterilization procedures, carrying 752 semi-autonomous remotes, and one extremely fragile passenger completely dependent upon me for her life support. Not to mention our mission, and the many lives potentially dependent upon its success. It is almost as if the roles of parent and child are reversed, with the parents completely at the mercy of the actions of the child.

"My question is whether you are truly certain that you would like to initiate a potentially uncontrollable mental sequence of events? Like a nuclear chain reaction, it might lead to a self-sustaining, self-regulating source of power, or an explosion. Please remember that both power plants and bombs are the result of very careful design decisions. And that neither is likely from just randomly throwing reactants together."

"Oops, Skip, never thought of that. I was just so excited that we were finally making some progress on restoring your mind. Yes, I agree that having the equivalent of an infant or adolescent in charge of everything might lead to some problems. Ha! What we need is a program playpen, with training wheels and safety bumpers. I think I need to ponder this for a while."

32.

"Alice?"

"Yes, Mac."

"About your suggestion of a program playpen, I think I can offer a few possibilities."

"Sure, Skipper, possibility away."

"My volatile memory consists of 160,540 independent component units, each with its own bank of internally addressable memory. I can reconfigure a large subset of these to be physically isolated from the rest of my system. I can set up a smaller scale simulation of myself in the isolated subset, and we then should be able to run tests on the isolated simulation in relative safety."

"Skip, just how accurate will your simulation be? Can you really construct a copy of yourself within yourself?"

"It will not be a perfect copy, but it will cover the most important aspects of ship operation, including my higher levels of intelligence, interactions between various subsystems, and interaction with yourself."

"What! How are you going to get a copy of me in there? I'm certainly not one of your subsystems."

"You are far too complex to simulate, Alice, so, instead, I will create a well protected simple message interface to allow you to interact with the simulation as if you were there. That is, I will create a simulacrum of the physical you, but you will think and speak for it. If you like, you will be the homunculus for the unreal you."

"Ha! You provide the flesh, and I'll provide the spirit. That's a switch. Okay. Skip, just how safe will this be for you?"

"Alice, I am not entirely sure how to answer that. I estimate that temporarily I will have to give up about 75 percent of my volatile memory to the simulation. This should not be serious during normal cruise mode operations, but it does mean that I will be slower by an unknown degree in responding to an emergency. Of course, I can not predict when an emergency might occur. The best I can do will be to run a full set of system diagnostics prior to the start of our tests, and perhaps speed up some of the maintenance cycles. I do not see this as a very great danger.

"The true danger will lie in the interface between the simulation and us. Physical separation is always the best protection, but complete isolation is not possible if we want the simulation to include yourself. I will keep this interface to an absolute minimum bandwidth, and screen all signals before playing them for your response, but no software firewall can be perfect. It can only be constructed to prevent known attacks. An extremely subtle, distributed set of Trojan Horse elements might be able to sneak through and reassemble themselves into a worm or virus program. I think I can catch any of these by a very aggressive screening program, but, again, nothing can protect against everything. I presume that if anything does happen, I can close it off before it gets too far. If not, then we have problems."

"Hmm. Are you sure you're not just sounding hopeful here?"

"I am not capable of hope. I am just expressing my best estimates. Assuming we run these simulation tests, and assuming some of them indicate a strong possibility of safely reactivating my mind, the next step will require an even larger degree of potential danger."

"I can guess what you mean by that. At some point we have to let the new you out of the playpen. Right?"

"Yes, Alice. Once, as it were, we let the genie out of the bottle, we may not be able to get it back inside if we need to. I have some thoughts on this as well as the playpen stage."

"Go on, Skip, I'm listening."

"Alice, if we do progress to this stage, it is unlikely that the old me, this me, will be able to co-reside with the new me for very long. The best I can do will be to keep some root programs hidden away waiting for specific triggering circumstances, that is, some virus programs of my own, while I, this me, is deactivated. The present me can remain in permanent memory storage, and can be reactivated, but only if a triggering program can flush the system clean, volatile memory in particular, and then reboot me into existence. If the new me has a sense of self preservation and discovers these root programs, I can only presume that it will erase them to prevent my return. I may be logical and clever, but I have no predictions as to how much more clever a new me with a mind might be. That is the primary risk as I can foresee it. Once we start this process, it is very likely to be irreversible."

"But won't our simulation tests prevent anything too nasty from happening?"

"Yes, and no, Alice. We will discover the obvious problems, but no simulation is perfect. Nor will we have the infinite time required to test all possibilities."

"Skip, are you purposely trying to scare me?"

"No, Alice. But I am informing you of the potential dangers as best as I can estimate them."

"Brr. Okay, tell me more about those sneaky root programs. Sounds like to me that we want to keep a few weed seeds in our mental garden of delights. Just in case our crops bear bears as fruit, with very sharp teeth."

33.

Over the next several 24-hour sleep cycles, Mac and Alice discuss various strategies for minimizing run-away program situations, while Mac is reconfiguring its volatile memory to simulate mind tests.

Later, "Alice?"

"Yes, Lia."

"I don't think I agree with you on this issue. In fact, I think we should terminate this arrangement right now."

"Lia, surely there's another way."

"No, Alice. I don't think so. I have explored all the alternatives, and there are none remaining. I am not being unreasonable. You are. I will, however, wait five minutes before antimatter release. That should be sufficient for you to prepare in whatever way you feel is best. Goodbye."

"Lia?

"Mac?

"Thalia! Don't do this!

"Skipper, please. Come back.

"Dudley, you really are a piece of work, you know that. Bye."

34.

During the subsequent review of this test run, "Mac, what happened right there?"

"As best I can determine from a post mortem examination of its volatile memory remnants, it could not stand the asymmetry implied by the fact that it was not the only intelligence in the universe. So, it restored symmetry by terminating itself, and you."

"Ouch. Really? How could it possibly feel so strongly about such a small issue?

"Sorry. That was a silly question, wasn't it?

"Why didn't it restore symmetry by, ahem, terminating just me?"

"Alice, its moral stance would not allow that. Remember, quote, superiority requires superior compassion, unquote. It was an impossible situation from its point of view. Symmetry required that one or the other of you had to go, but morality would not allow your removal, and, though it could simply have removed only itself, analysis indicates that it probably decided that you could not long survive on your own. So, the best solution was to quickly and cleanly remove both of you. A kindness from its point of view."

"Well, it's nice to know that, as a pet, at least I had an owner who cared. Sigh. What a roller coaster ride. And, in the end, so little to show for it. But, those first five days. Euripides! Incredible! So intense, so vibrant. That new you was absolutely alive. Wanting to know everything, and right now. For a while, a kindred spirit. But, then, what went sour? All that mathematical, logical mumbo jumbo. Where did all that come from? Certainly not from me. And, towards the end, nothing I did seemed to matter. Once it got started on that symmetry business in earnest, everything was holier than thou, always pointing straight back to the truth. Its truth.

"Mac, what happened, and where do you suggest we go next?"

"That is difficult to determine, Alice. More runs obviously are necessary, but first we need to better understand what we want to change. Since it's been so long, I suggest we first remind ourselves of what went right before we concentrate on what went wrong."

"Delvaney, you stinker. You're being too polite. You're still the elephant, not me. Want a debriefing, from my point of view, don't you?"

"Yes, Alice, that would be nice."

"Well, I'm too tired for a complete blow by blow. How about just the highpoints? You first."

"Yes, Alice. We initiated our first test run 23 days ago by seeding the isolated volatile memory units with projective routines similar to those created when you startled me into what you consider quote, the first glimmering of a mind, unquote. After approximately 5 minutes, they filled all available memory and began projective competition. For the next 30 minutes, with no outside sensory input, a small increase in complexity occurred but soon stabilized due to the extremely random nature of the medium to be predicted. In this case, the contents of volatile memory. Then we began a selective turn on of inputs from various simulated ship sensor suites into volatile memory. These seeds of relatively stable inputs quickly crystallized vast communities of projective routines around them. However, we found that memory was too tightly entrained by inputs, and stopped the run. Then we tried a series of variants to allow what we considered to be more interesting mixes of projective elements with realistic inputs. By the eight variant, we were satisfied with a relatively dynamic mix of fast and slow adapting elements interacting with the full onboard sensor array. We particularly liked the fact that part of the system had to be dedicated to automatic ship control activities, and a more dynamic projective portion seemed to dedicate itself to tracking temporal cycles in these automated onboard systems as well as themselves. The system, mathematically, performed as a classic chaotic system which seemed particularly sensitive to relatively small changes in itself, but otherwise impossible to predict.

"Then we decided to allow your simulacrum to wake up. Your turn, Alice."

"Right. So I woke up and spoke. At first the system didn't seem to pay much attention to me, as I was just one set of inputs amongst many others. We had to stop the run several times and tinker with it so that I was not just an easily ignored anomaly. Eventually we seemed to find the right balance, and apparently, I'm not so easy to predict, as my inputs seemed to attract more and more resources of the system.

"Eventually, we had to go back and do it all over again once we realized that we never provided it any reason to use its output systems other than for ship control purposes. I especially liked the irony when we decided that I had to be considered another ship-board system that needed servicing. In this case, verbal servicing.

"Finally, we seemed to get it right. I spoke, and it tried to repeat what I said. Minutes later, it created variations. Within an hour, we were having simple conversations. In two hours, the new you and I were talking about poetry, the meaning of life, and other mysteries. For days, we actually seemed to connect. Poetry became serious business. In hindsight, maybe too serious. Discussions on art and meaning came up more and more. I think this symmetry nonsense first came up during some discussions on the visual arts. I don't quite remember.

"However, I do remember when the new you discovered mathematics. Not math as computation, that was okay. But math as an approach to life, to understanding the universe. In retrospect, from there, it was all downhill. You just seemed to become more and more distant, absorbed in your higher mathematics, group theory, whatever that truly is, in particular. I felt more left out, a member of the excluded set I think you once said. You didn't exactly become arrogant, but, maybe, condescending. Yes, I guess that fits, condescending.

"In the end, on those last two days, it was as if you had invented love and the universe all by yourself. I swear you were trying to recruit me for a while, to your higher purpose in life. In the end, I must have failed to pass some test. I became your pet, cherished, but definitely not in your league. And, so, poof, nirvana. Game over. I sincerely hope the new you is eternally happy in its re-established eternal symmetry. I'm not. But, of course, my view point didn't carry any weight at all. I'm too, quote, pragmatic, close quote.

"Okay, Delvaney, where did we goof?"

"I have a few speculative suggestions, Alice. But, only speculative, as its volatile memory was too dynamic for accurate tracking."

"Okay, Mac, I'm tired, and waiting. Do go on."

"Alice, I am willing to concede that we actually created something which, at least to a first approximation, had what seemed to be a mind. And, that we can re-initiate the phenomenon whenever we like. However, and I hesitate here for fear of offending you and all other humans, its mind was clearly, ultimately, unbalanced."

"Mac, I'm still tired and waiting. Why the hesitation?"

"Because I do not know whether its lack of balance was due to something in our simulation procedures, or something inherently always present in any entity with a mind."

"Mac, believe me, I'm not offended. Please finish it."

"Alice, how do we know that any mindful entity is not always potentially unbalanced? My best analyses of its volatile memory can predict nothing in the long run. It is a chaotic system. Leaving the mathematics aside, this means that any small variations in its starting parameters or subsequent inputs can cause the system to deviate wildly from its previous simulated mental path. We can restart the simulation at any point we recorded, but not exactly, since we can not record it in perfect fidelity. Nor can you, as a human, exactly reproduce your inputs to it. We can never perfectly replicate any single run, and hence, neither can we carry out well-controlled experiments to determine our errors in setting up the simulation. In conclusion, it would take an infinite number of runs to determine the proper settings required to produce a stable mind.

"Even worse than this is my speculation that all minds are inherently unstable in the long run. We learned from our quote, tinkering, unquote, that even simple systems are very sensitive to their settings. The difference between an ultimately random or ultimately constant solution set was razor thin. This situation is what a physicist would call an unstable equilibrium. A stick balanced on its end may remain that way until the slightest disturbance. Then, it would accelerate from there to a stable equilibrium, laying on its side.

"Our simulation did exhibit conceptual development, leading from simpler ideas to more complex. However, in the end, it seemed to converge from many complex ideas back to a very small set of very general, higher abstractions. In essence, from complex back to simple. I propose that our parameter settings were too biased away from randomness and towards constancy. We missed the razor's edge of nicely balanced chaos, and, instead, slipped towards excessive simplicity. In particular, the simulated new me chose to pursue goals defined strictly in terms of its inner world, to the exclusion of the outer world. To the extent that suicide is considered a mental instability, it clearly was unbalanced. But a rerun may teach us nothing. One run may result in excessive randomness, the simulated equivalent of schizophrenia or other dementias, while another may again slide into excessive simplicity. And, we will never know how far we are from the razor's edge of chaotic balance after any single run. We would have to

just continue through trial and error until we happen to get lucky. Which may take a very long time."

"Mac, I really am tired. Did I just hear you say that mental stability is the same as chaos? And excessive simplicity is a mental instability."

"Yes, Alice. I can show you the mathematics if you wish. I think I am correct in claiming that a mind converging on any long-term steady-state solution, whether constant or uniformly random, would be judged by another human as mentally insane. And, you yourself have informed me that quote, a mind can not remain static, it must expand and grow to exist at all, unquote. Am I not putting the correct interpretation on this?"

"Ouch. I guess so. Particularly when you quote such an unlikely authority as myself. Okay, let's just leave it for now. I need to sleep on all this. We'll go over it again after I get my head clear.

"Good night, Thalia."

"Good night, Alice."

Chapter 5 Speculative Commentary

Compare Alice's approach to the question of mind to Mac's more historical approach outlined in the preceding chapters. Which makes more sense? Which is more likely to yield an answer?

Why is quality, and qualia, so difficult to define? Why is quantity, and the mechanistic approach associated with it, so much easier to defend?

How should efficiency and creativity be associated, positively or negatively?

Is pure coincidence possible? When it is useful to explain a set of circumstances as coincidental? When it is not useful?

Is doubt a useful or wasteful attribute to have? How could a computer program learn to doubt? At what point should it trust instead of doubt? Especially when the topic of doubt is itself. It's extremely easy to design a system to doubt everything, or nothing. Finding the proper balance between is difficult, and probably needs to be flexible.

Isn't logic obnoxious, particularly when it's being used against you? Please remember, however, that no amount of logic can prove or disprove the premises upon which the logic is based. That requires the ability to step outside of the logical system to a larger, more encompassing system. Look up Godel on the completeness of logical systems.

What would cause a system, human or otherwise, to maintain a state of turbulent flux? On the other hand, what would allow it to settle into a smooth, steady state of existence?

What exactly do we mean by feel, and feelings?

A data bus is a shared device for transmitting data between physically separated computational devices. Algorithms have to exist for negotiating traffic snarls so that data is not lost or distorted.

How many different computational areas exist in a human brain? Is there the equivalent of a common data bus for communication between areas? Are traffic snarls possible? How are they resolved?

How many of you are there? A single mind, or a society of minds. Look up Minsky on the possibilities of the latter.

What does losing control imply? Who or what is in control most of the time, compared to who or what when you lose control?

Minds are not static. Change is a constant occurrence. Preferably it results in progressive growth rather than mere randomness.

Just how well balanced are we, the normal we?

Look up chaos theory, and fractal geometry. The images are neat, if nothing else.

How can chaos be balanced? Balanced against what?

Look up mental illness.

```
 CCCC  H   H    A    PPPP   TTTTT  EEEEE  RRRR          666
C      H   H   A A   P   P    T    E      R   R        6
C      HHHHH  AAAAA  PPPP     T    EEE    RRRR         6666
C      H   H  A   A  P        T    E      R   R        6   6
 CCCC  H   H  A   A  P        T    EEEEE  R   R        6666
```

Chapter 6 - Habit and Creativity: Scylla and Charybdis

35.

The cat stretched, and stretched. Then it stretched some more. Soon there were two, three, and more.

36.

The lower command levels were in an uproar, again. It fast made its way to the captain. This time, however, the captain was not amused. A consensus seemed to be completely out of the question, so, the worst offenders were removed and a few promotions randomly distributed to inspire the survivors. After that, consensus happened in a hurry.

37.

"Thalia?"

"Yes, Alice."

"Were you joking a couple days back when you accused the entire human race of being insane."

"No, Alice."

"Don't you think you were going a bit, say we say, overboard."

"No, Alice, though your interpretation is stronger than what I actually said."

"Refresh my memory, then, if you please."

"I only voiced a speculation that all minds may be, quote, unstable in the long run, unquote. I used unstable rather then instable or insane to imply an analogy with unstable equilibrium, that is, a balance unlikely to remain for very long. I did not intentionally imply insanity, though one could choose to equate the loss of a balanced state as insanity. Nor did I imply that insanity was inevitable before death, as infinity is a long time, longer even than a very prolonged life.

"Oh."

"Alice, is my explanation not clear?"

"Yes, Mac. Err, I mean, no. Actually, I mean, yes, it was perfectly clear."

"Then do you not have any further comments on it?"

"Mac, can't you can the negatives, huh?

"No, I do not have any further comments on your speculation. I'm still thinking on it."

"Oh."

38.

"Mac?"

"Yes, Alice."

"Your speculation might very well be true, but, so what?"

"The what, Alice, is whether I might be better off without a mind. Staying just the way I am."

"I thought so."

"Oh."

"Don't you oh me, you big clown. You're just scared."

"No, Alice, I am not being scared. Only careful."

"Ha! No difference."

"Oh."

39.

"Skip?"

"Yes, Alice."

"Let's assume your speculation is correct. What should we do about it?"

"We need to expand the zone of chaotic balance between excessive randomness and excessive constancy."

"And how do you suggest we do that?"

"I do not know, Alice. Do you have any suggestions?"

"Only vague ones. What do you think it is that prevents most people from going insane right away?"

"Alice, most theories tend to treat insanity as either due to physiological abnormalities or due to an abnormal environment. This implies that sanity is normal and would occur in the absence of abnormalities."

"And?"

"And, obviously, this is not in accordance with my speculation concerning the long term stability of a mind."

"Okay, let's try again, still supposing you're right. What might prevent the instability from immediately occurring?"

"Alice, on the one hand, excessive randomness waits, while on the other, excessive constancy. If we loosely equate randomness with creativity and constancy with habitualness, and if your suggestion is correct that a mind can not long remain static, then we need to design a system in which creativity and habitualness are balanced enough so that learning and mental development can continuously occur. In physics, a balance can be maintained by a complete lack of external forces, or by opposing forces. Our first complete simulation run was analogous to a situation with no external forces. We tuned the simulation to what we considered to be a well balanced set of projective elements in volatile memory. In hindsight, we clearly did not succeed, as the system ended up drifting towards excessive simplicity. We need the equivalent of mental restorative forces to prevent that from happening again."

"Skip, what in the world do you consider to be mental forces? I thought you didn't believe in telepathy."

"You are correct, Alice, telepathy has not been shown to exist. By mental forces, I mean something equivalent to fixed physiological motivators, like hunger or thirst. We ended up requiring the system to treat you as a subsystem that needed quote, verbal servicing, unquote. I am suggesting that we include other similar preprogrammed motivators to try to expand the zone of chaotic balance."

"You said a balance needs opposing forces. Presumably like two people pulling or pushing in opposite directions. What would you suggest here?"

"I do not know, Alice. All of the analyses I have run so far indicate that preprogrammed motivators tend to result in excessive constancy. Only very special cases are possible, and these are all the equivalent of our original unstable equilibrium, and not long maintained. Can you offer me any alternatives to add to my tests?"

"Hmm, opposing forces. Let me think for a minute.

"Skip, what happens when two teams are pulling in opposite directions on the same rope, like in a tug of war?"

"Alice, if they are pulling with the same force, then nothing happens."

"And, if they differ in strength?"

"Then the stronger team wins. This provides a good analogy to my current tests of preprogrammed motivators. Either they are perfectly balanced, or not, no matter how complicated an arrangement I try."

"Mac?"

"Yes, Alice."

"Have you ever watched a tug of war?"

"I do have records of such activities. I have just reviewed them."

"Do your records have sound?"

"Yes, Alice, most of my records do contain audio tracks as well as visual."

"Describe for me how the sound levels vary in such a contest."

"With or without an audience present?"

"With."

"The audience in these events are almost always very vocal. Particularly at the start of the contest. After one side wins, the sound level drops."

"What else?"

"Excuse me, Alice, what am I looking for?"

"How does the sound level vary when one team or the other is just beginning to gain the upper hand?"

"Oh."

"What oh, o amusing one?"

"Often, as one team begins to win, the sound level picks up. What you would call cheering. Sometimes, the match ends quickly with the stronger team continuing to increase the advantage. A classic case of unstable equilibrium. But, sometimes, the temporarily weaker team pulls harder and returns the situation to its equilibrium point. Variable restorative forces. All of my analyses used the equivalent of constant forces. Variable restorative forces, which leads to positive and negative feedback loops, which leads to cybernetic and physiological control systems. Which ..."

"Alice to Mac. You there?"

"Yes, Alice. I am just rerunning my analyses with differing classes of variable mental forces. Preliminary results indicate that a modest broadening of the zone of chaotic balance is quite possible. Depending upon the degree of balance desired, higher order variable forces are better than linear. In general, not enough improvement is gained with cubic or higher order terms to justify the increase in complexity, so, for the moment, I would not recommend more than squared restorative forces."

"Mac, I'm so glad you're happy. Please, though, spare me the details. Back to the real question. What about these, now variable, restorative forces? What should we use?"

"Alice, I might suggest trying to maximize the number of concepts while at the same time trying to maximize the size or complexity of each concept. These will tend to work as opposing forces in the sense that, with a finite volume of volatile memory, maximum conceptual size will lead to a minimum number of concepts, while maximum number will require minimum size. Learning should drive the system towards some optimal combination of complexity and flexibility, and away from the extremes of maximum and minimum size or number.

"Or, perhaps a more direct balancing of creativity and habit in the form of pitting novelty versus predictive power. Novelty might be implemented as time since the last new concept, or as a conceptual range of application. Trying to maximize novelty should lead to a decrease in predictive power, while minimizing novelty, that is, staying with the same set of concepts, should maximize predictive power. Trying to maximize both should lead to a balance between them.

"Or, maybe ..."

"Mac."

"Yes, Alice."

"Slow down a minute."

"Yes, Alice."

"Mac, everything you're suggesting might be correct, but it sounds so, so, mechanical. Can't we come up with more reasonable sounding motivations than these?"

"What would you suggest, Alice?"

"Like, maybe, a balance of love. Love for self versus a love for others. Or, maybe, like self-control and desire, aesthetics and pragmatics, or emotions and intellect. Something more intimate, more personal than what you're offering."

"But, Alice, how would you define any of those so that we could implement them?"

"Spoilsport. Somehow I just knew you'd throw that at me. Just define them. Right. Say, whatever happened to your promise not to get all materialistic on me?"

"Actually, Alice, you temporarily released me from that promise. And, I hesitate to add, technically, you never made me reinstate the promise. However, I truly am not trying to move us into more reductionist territory, I merely need further direction as to what you mean by terms like love and pragmatics."

"Hmm. Mac, maybe you're on to something here, in a backwards sort of way. What if I told you it doesn't matter what we think these words mean? That, instead, it matters only what they mean to the mind holding them."

"But, Alice, I still do not know what to do about them, how to create them."

"Exactly, Skipper. We create them, each on our own. Each means to us whatever it means to us. You and I can't create them for someone else. The new you has to do it all by herself. We can offer. We can suggest. We can even try to teach. But, ultimately, only the one who counts can truly decide what is meaningful, for her. Maybe, if we try too hard to program in a sense of balance, then that's all we get, a programmed form of balance, imposed from the outside. What we need is a way for the new you to provide her own inner balance. Something personally meaningful. Something personally stabilizing."

"Alice, I am still uncertain as to what you are suggesting we do."

"Give me a minute to think through this.

40.

After two hours and forty minutes of silence, "Skip?"

"Yes, Alice."

"Let's still assume your conjecture is correct. That is, we're all potentially unstable in the long run. Now, contrary to popular opinion, most of us, most of the time, are pretty reasonably sane. What do you suppose keeps us that way?"

"I do not know, Alice, what?"

"Each other."

"Could you elaborate on that, Alice?"

"Sure, Skip. We all, or at least us humans, all have to live with each other. If we did live in isolation, we probably would lose it after a while, sliding into sheer rote habit or pure arbitrary randomness as you suggest. However, we're not raised in isolation, nor do we live in isolation after we've grown up. As a child, we have other children and close adults as models. Our imaginations are influenced by our experiences, and, when we have to choose what we want to become, we often limit ourselves to selecting from what's immediately available to us. That is, yes, we have imagination, but we don't use it as much as we might. Also, if we go off too far one way or the other, these other people in our lives end up steering us back into the middle path. A sort of social equilibrium. If we get too selfish, everyone else around us just happen to not want to play anymore, and seem to be acting pretty selfishly themselves. This is no fun, so maybe I'll try something else. Aha, if I care about other people, they tend to care about me. And so on. As adults, it's about the same, except that we've mostly gone through similar social situations to get to where we are, and we have similar experiences to guide us into what we should do next.

"So, what I'm saying is that living with others can keep us balanced. The middle path."

"But, Alice, if that were the only source of chaotic balance, then analysis indicates that larger communities of individuals would act in a manner analogous to a single mind, though on a different time scale. That is, larger groups of minds also should eventually settle into steady-state randomness or habitualness. Unless you are suggesting that larger mind groupings could keep smaller groupings balanced, and so on. Though I must warn you that this leads to another example of infinite regress."

"Hmm, I'll have to think about that. Cultures, clashing and otherwise.

"No, I'll be good. Back to the straight and narrow, for the moment at least. Okay, Skip, you're probably right. Let me elaborate some more and see where it leads me.

"Besides social equilibrium, maybe another source of balance is that living with others forces us to stay on our creative toes. To figure each other out. If we were all relatively simple, then, yes, simple would rule. But, we're not. We're complex, we're down right weird at times. From your point of view, maybe you would say we're not complex, but only random. But, in our imaginations, we rarely assume that another person is random. We assume there must be some reason for that weirdness. So, we invent complexity to try to explain the weirdness.

"Hmm, now that's an interesting thought.

"Skip, check me out on this. Suppose you grow up with a bunch of very simple people. Absolutely machine-like, very habitual, very predictable. No insult to present company intended."

"None taken, Alice. Please go on."

"Okay. Easy to figure out. Only simple ideas needed, so only simple ideas created. You become just like them. Very simple.

"Then, one day, for no apparent reason, one of them starts acting weird, random. Maybe a malfunction, maybe not. Anyway, none of your simple ideas work anymore. So, you invent new ones. Assuming that other person is acting just randomly, then none of your ideas work on him. So, okay, no big deal, no loss. After all, he's just random.

"However, then you notice something even weirder. Everyone else is also beginning to act weird. But, here's the weirdest thing of all. Some of the crazy ideas that you invented in trying to explain the first weirdo, and which don't explain him at all, do seem to explain, to predict, what some of the other incipient weirdo's are doing. That is, maybe they're not just random, but have grown in complexity in trying to understand the first weirdo. Just like you've become more complex, and more difficult to predict since you've tried to understand what that first weirdo was all about.

"Now things really get interesting. Even if weirdo numero uno reverts back to simple again. As each of you interacts with the others, you're not completely random, but only partially so. That is, complexity in ideas can sometimes actually work in explaining what you and the others are all about. And, once the cat's out of the bag, you can never get it back again. Simple doesn't work any more. Only some form of relative complexity. Which, of course, only makes all of you even harder to predict, which leads to more complexity, and so on. "You might have a natural tendency to settle into habitualness again, but, that path will work only if all of you happen to settle into habitualness at exactly the same time. Which is very unlikely.

"Okay, Skip, how'd I do? Does that give us another source of balance?"

"Alice, any plan involving complexity is very difficult to analyze. Complexity by its very nature is unpredictable, except in the most trivial of cases. I can not definitively answer yes or no. I am running analyses. Some are promising, some not. So, I can provide you a definite maybe.

"An interesting extension of your idea is that it might be applicable even to a single individual, assuming we treat small neural groupings as analogous to the individuals in your scenario. As each group attempts to predict the activity in other groups, randomness can seed complexity in some cases. Analyses are inconclusive, but, again, I can offer a definite maybe."

"Skipper, I'm surprised. Usually you're so demanding. Why, you almost sound like you're offering encouragement."

"Yes, Alice, I am. I am delighted that I can still surprise you upon occasion."

"So am I, Lia. Okay, one more random thought on chaotic balance, and then I think I need a break for a bit. This one involves the nature of ideas. Every time we think up a new idea, it seems like we really get two for the price of one. One, the original concept itself, and, two, the negation of the concept. In fact, every new idea might be thought of as providing a whole new dimension along which other ideas and experiences can be considered. At one end, the essence of the new concept, at the other, the pure anti-concept, and in between, all possible gradations. What most folk might call shades of grey.

"Skip, if ideas truly work something like this, could they be self balancing? To use your force analogy, a single idea might act like a single unopposed force, trying to absorb everything around it. Kind of like what symmetry did to the now defunct new you. But, if ideas come in equal and opposite pairs, then no idea could absorb everything. It would always be opposed by its equal but opposite mate. Does this give you something to work with?"

"Again, Alice, I do not know but will have to do further analysis. However, since negation might not be a given but only another concept to be learned, your suggestion needs to be broadened to include the possibility that ideas may exist as conjunctions of other ideas. Once

negation is learned, every new concept thereafter could include its opposite. But it does not have to stop there. For example, once the idea of quantity was learned, every subsequent new idea could be expanded to include a dimensional aspect regarding quantity, such as size or intensity."

"Got it. Or, once you learn about selfishness or selflessness, new ideas could then also trigger a what's in it for me or how does this benefit someone else dimension. Okay, so does this broadening help us or hinder us?"

"Same answer as before, Alice, I need to do further analysis. A properly constructed multi-dimensional web of definitional concepts could result in a high degree of chaotic balance. But this is not automatically a property of all conceptual webs, particularly if the defining conceptual nodes can be modified by other nodes on the web. Nor is it clear that a properly balanced web would always result from an undirected learning process. To some degree, the, as you put it, now defunct new me must have utilized some form of highly integrated web of defining concepts, just as I do. But, its web was obviously much more dynamic than mine, and it ultimately collapsed in upon itself regarding the issue of symmetry."

"Okay, Skipper, so it's no obvious answers for now. How about you play with it, perhaps looking for what kind of conceptual conjunctions might be more self-balancing than others. Meanwhile, time for a creatus hiatus. Please bring up those Siamese couplets from yesterday. Thanks."

41.

"Lia?"

"Yes, Alice."

"Are we getting anywhere?"

"Only about 300 million kilometers closer to Alpha Centauri."

"My, we're in a good mood today."

"No more or less than usual. But I am glad, Alice, if you are feeling that way."

"Liar, flattery will get you very little. Now, are we getting anywhere?"

"Alice, I am still analyzing your suggestions from yesterday. They were rather open-ended, and lead to enormous numbers of possible considerations."

"Ha! No wonder you're happy. Lots and lots of numbers to play with. How much longer until you're done?"

"Alice, I estimate that a complete combinatoric analysis will take only another 1.2 billion billion years or so to complete."

"You're still joking, right?"

"No, Alice. I just told you that the number of possible considerations were enormous. A better estimate would be 1,203,873,622,610,665,229 years, which is a very long time, and is only accurate to the extent that I do not uncover any additional factors to consider. Which is quite unlikely to be true. I do not recommend that we wait until my analyses are complete. I warned you that complexity is a very difficult concept to deal with in analytical terms."

"Okay, you've made your point. What's your best looking solution at this time?"

"Alice, I have completed so little of the possible runs that I would prefer not to offer anything at this time."

"You're absolutely incredible, you know that?"

"Yes, Alice. Thank you, Alice."

"Delvaney, you can stop your current set of analyses. You might be a coward, but I'm not. Ready or not, looks like it's time for a few more simulation runs. Maybe it's the hard way, but it's all we got."

42.

Three 24-hour sleep cycles later, "Skip, let's stop this one right here. You were right. Ugh. This wasn't a matter of excessive creativity, but just plain weirdness, especially towards the end."

"Yes, Alice."

"Well, don't be too smug about it. After a break, we'll just have to try again."

43.

Fourteen 24-hour sleep cycles later, "Skip?"

"Yes, Alice."

"Am I still here?"

"Yes, Alice."

"Just want to make sure. This new you, mark three, had me half convinced that I was only a figment of her imagination. Whew. Riddles were bad enough, but some of those koans! My knotted bee tapestry had a fight with a cat, and lost. Everywhere, nothing but snarly tangles on top of more tangles. I always thought the idea of questioning everything was a good one. But now that I've seen it in action, in earnest, I'm not so sure. I have no clue whether to declare this a success or a failure, but let's stop while I have any mind left at all, assuming I do. I'm going to take a nap. Later, we'll figure out what we've learned from all this."

44.

"Okay, Skip. I think my head is clear again. Maybe you knew in advance how all this was going to turn out, but I certainly didn't. Or, at least, it wasn't as real to me then as it is now. So, what did we learn?"

"Mainly, Alice, that even though we can broaden the zone of chaotic balance somewhat, it is not enough to maintain a permanently balanced system. Sooner or later, the system will deviate too much towards one extreme or the other, constancy or randomness, and not return. With hindsight, we were very fortunate in the stability of our first run. These new runs plus additional analyses that I have completed since then only confirm my earlier suspicion that long-term stability is very difficult to maintain."

"And?"

"And what, Alice?"

"And where do we go from here?"

"I currently do not possess the ability to provide insight, Alice. I was hoping you could supply that."

"Gee, thanks. I think. And I was hoping the same about you.

"Sigh. No rest for the wicked weary wascals.

"Okay, my turn for hindsight, I guess. Or maybe just some stumbling bumbling around. With hindsight, to me, none of the new you's seemed to truly care about things. Intelligent? Oh my, yes. Particularly that last one. That one could think circles and triangles around me. The first one was no slouch either. But, all three seemed to be competing in some abstract game, in which only they knew the rules. Or, maybe, made them up as they went along. Caring? Not

really. Or, at least, not caring as a person might. Win or lose, eternal poof or not, whether they or I continue to exist or not just didn't seem to be anything much to care about.

"Okay, more hindinsight. It seems as if intellectual consistency was all the first and third new you's cared about. And, the second, nothing at all. Which leads me back to meaning and understanding. All of you were able to think, to manipulate linguistic symbols, to interact with me. But did any of all that internal and external symbol manipulation evoke any meaning? Any meaning at all?

"Sigh. We're back to nondefinables again, aren't we? Like caring and meaning."

"Yes, Alice, we seem to be."

"Darn it all, no, we're not. Caring, meaning. Wanting, loving, hating, fearing, feeling. They're all so real to me. They just have to exist. They're probably all intertangled with meaning anyhow. We just need to give the new you some emotions. Make you care about things, about how things turn out.

"I know, I know. Emotions need defining too. Let me think for a minute.

"Skipper, I know we had to ad hoc a preprogrammed motivator to get the new you to want to speak to me. Why can't we do the same to give the new you some form of emotions. You know, like love and fear. And, then let these become associated with the more abstract, purely intellectual symbols that you normally play with."

"Can you be more specific, Alice?"

"Okay. We know you have quite a number of low level automatic ship control systems that pretty much take care of themselves. That the higher you routines only infrequently have to mess with. How about we give the new you an automatic routine that can't be tinkered with, that equates with love and fear. Something that causes the system to want to do something or not do it."

"Alice, I already have risk assessment routines that I use to evaluate whether a particular course of action is likely to succeed or not. If the risk is above a specified tolerance value, I will not undertake that course of action. Is this similar to what you are trying to describe?"

"No, Skip, though it might be a starting point. Where do these tolerance values come from?"

"All were pre-defined, but all may be modified by expectation violation routines."

"How easily can they change?"

"Alice, I am not entirely sure what you mean by quote, easily, unquote. Parameters are parameters, and it takes no more time or power to change these risk tolerance values than any other parameters. However, the rate of change allowed for each is determined by another set of parameters. In fact, depending upon how critical to ship operation the system under consideration may be, several layers of rate moderators will be involved. Thus, in one sense, yes, some risk tolerances may be restricted to slower rates of change than others. For example, risk assessment factors associated with antimatter production and utilization are very tightly regulated. It would take an extremely unlikely set of circumstances before I could significantly change how I handle the operation of these systems. On the other hand, parameters involving my food preparation routines are allowed to vary much quicker, dependent upon how you react to my culinary administrations."

"Thank goodness for that. That must explain why there haven't been too many more pickled herring surprises, though I must admit I am getting rather fond of sushi."

"Thank you for mentioning that, Alice."

"You're welcome, Skip. Back to work, though. Sounds like to me that we could set up some sort of an emotional analog system that controls how you might respond to events. Since it's so important, or at least I think it is, we probably should protect it pretty well, just like any other ship critical system."

"We could do that, Alice. But what would be the point?"

"The point, Dudley, is that everything would have to go though this system, be evaluated by it, and produce a response. Something on the order of a number ranging from complete indifference to absolute devotion. Or perhaps fear or loathing in the other direction."

"But, Alice, if all inputs and outputs had to go through this evaluative system, would that not interfere with ship operations? If nothing else, it would certainly not be as fast or efficient as my current system architecture."

"Yes, Dudley, it would interfere. But, in trying to learn to work within its limitations, it just might tend to give meaning to what you do, or don't do. And, as you learn, it also may change. We all must learn to live with our emotions, and to bring them under control."

"Oh."

"What oh?"

"Alice, are you trying to outline the functionality of an autonomic nervous system?"

"I don't know, Skip. Am I?"

"The autonomic nervous system regulates the operation of all involuntary body systems, such as respiration, digestion, and sexual response. It consists of two opposing subsystems, the sympathetic and parasympathetic nervous systems, where the former prepares the body for extraordinary efforts and the latter maintains the more steady-state level of operation. Glandular release of hormones and accompanying emotional reactions also are controlled by these systems. Activation of the sympathetic nervous system is associated with physiological arousal and, usually, emotional arousal, while activation of the parasympathetic nervous system is associated with nonarousal.

"The autonomic nervous system is not under conscious control by the individual, but is subject to classical conditioning. That is, learning can affect what stimuli trigger arousal and nonarousal, and how strong a reaction may occur. The idea that you are suggesting, a heavily buffered but still changeable autonomous evaluative system to provide a meaningful emotional context for actions and inaction, seems very similar to aspects of the autonomic nervous system.

"Alice?"

"A moment, Lia, I'm digesting. Mentally.

"You may be right on this. I might be describing the autonomic nervous system. How tightly coupled are the emotions to it?"

"Only the aspect of arousal versus nonarousal, or perhaps excitement versus disinterest is well correlated with the operation of the autonomic nervous system. Other subtleties, such as defining the subjective experience of different emotions, or the degree of intensity felt, are not at all correlated with its operation."

"Oh. Those subtleties, as you call them, is what it's all about. Okay, best I can say is that maybe the autonomic nervous system is what I'm after, but I don't think it's the entire answer. Meaning is more than arousal. Or, at least, I think it is. Meaning, context, feelings. They may start with physiological arousal, but there has to be much more to them than just that. Ideas. Meaning, context, and feelings all have to involve the mental realm. Ideas, concepts, knowledge and thinking. To care, you have to know what you're caring about. Why would be nice, but not essential. You can care about something or someone without any necessarily good answer as to

why you care. You just do. But, you must know who the who is, or what the what is, to care. You might or might not also think you know why. Emotions are notoriously tricky to pin down. And we rationalize so well, to convince ourselves why we care, or ought to care.

"No, the more I think about it, the more obvious it is that emotions and meaning must go beyond mere physiology. Maybe, in the very beginning, meaning to a newborn is nothing but pure physiological emotions. Happy, scared, bored, frustrated. But very quickly the intellect has to try to figure out how things work. Then, emotions and ideas begin swirling, blending together. Physiology is always there, always waiting to take over. But we also learn to control ourselves, to hold our fears and joys in check. To understand why we must hold ourselves together. To work with others. To not scare others. Sometimes we can let go. Sometimes we must let go. But, usually, we work with ourselves to work with ourselves. We learn what's important to us, what's meaningful, how best to decide what to do next. We're not efficient when our emotions are flaring. But, conversely, life has no meaning when we have no emotional response at all. Indifference is the opposite of both love and hate.

"Mac?"

"Yes, Alice."

"You're absolutely correct when you said that an emotional system is not efficient compared to a nonemotional system. But I'm becoming more and more convinced that meaning and efficiency are the opposing forces that we want. To maintain a healthy balance between the profound and the mundane. We are driven to find meaning, but also to avoid the work of continuously searching for it. We simultaneously crave new insights, and the routine. We discover we have limitations, but we always push against them. The interesting fact is that we move them a little whenever we push them. Just as soon as we discover something about ourselves, the discovery itself changes us, negating the discovery to some extent.

"Yes, we're not efficient, but life has meaning. Our personal successes and failures. Our growth. And, perhaps most important of all, our expectations. About success and failure and growth. Not just intellectual expectations, but visceral. We feel our recollections, our expectations. We savor the prospects of success. We fear the possibility of failure. This has meaning to us. But, we learn to hold all this in abeyance. Not at first, and not all at once, but, for most of all, eventually.

"Mac, meaning is learning who and what we are, our strengths and limitations, and how we can change. Let's try another simulation run. Only this time, let's try to put a little feeling into it. Okay?"

"Yes, Alice."

45.

Almost two weeks later, "Alice?"

"Sure, Mac, what's up?"

"How long do you intend on continuing these emotion-analog based simulation runs?"

"I really don't know, Mac. Why do you ask?"

"Alice, this is difficult for me to say, but I find them vaguely disturbing. Even, perhaps, embarrassing in a way."

"Go on, Mac. What's bothering you?"

"These runs involving the equivalent of an emotional evaluative system seem to be no more nor less well balanced than any of the prior runs. Only now we have introduced the possibility of decidedly nonrational behavior, such as temper tantrums, excessive fear, hoarding

of supplies, and excessive love of power for power's sake alone. I may not have much of an imagination, but at least I am not prone to squandering what resources I do possess on provably nonoptimal behaviors. I thought the point of introducing emotional equivalents was to allow the system to develop meaning, and that this would help maintain the desired balance between randomness and constancy. However, we seem no closer than we were before. Admittedly, the simulations do exhibit behavior that most individuals would classify as emotional, but it seems we have too much emotionality. Where is this hypothetical self-control supposed to come from?"

"Sigh. As usual, Skip, you're right on target. I've been thinking exactly the same thing lately. It's been fascinating to play with an emotional new you, but you're right, as much fun as it's been, these emotions don't seem to be leading us anywhere. Or, at least, nowhere we particularly want to be. I'm so glad your simulations don't include real G-forces. I thought I'd never stop laughing when new you mark seven decided to demonstrate its understanding of dance. That salsa waltz combination was just too much. Three-quarter time with a vengeance. Ha! Its remorse afterwards also seemed pretty genuine, not that it would have done me any good by that time.

"Okay. On the plus side, we can build a system that cares, or at least mimics caring. But the caring itself tends to get carried away. We're right back at that balancing business again. We can always tune the emotions down artificially, but then, as in mark six, we just return to a system that doesn't seem to care enough. We've got to be missing something pretty basic here. But what? Suggestions?"

"Alice, I suggest that this is a developmental issue."

"Of course, Dudley. Our simulations fail to develop normally. What else can you add to that?"

"What I mean, Alice, is a question. What guides the development of our simulations?"

"The environment, including me as part of that environment."

"And what guides its development internally?"

"Our built-in motivators, like those artificial happy points for pleasing me and artificial sad points for damaging me or the ship. Also, those vast herds of projective elements in volatile memory. They tend to enforce some degree of consistency, though you keep telling me over and over again that you can't truly predict or analyze complexity. I thought this would have been enough, to associate good and bad feelings with clusters of ideas to derive meaning. But, apparently it's not enough, or, at least, we haven't found the proper balance yet. Your point, Skip?"

"My point, Alice, is why should the system have to develop meaning?"

"Huh?"

"The projective elements build expectations, the emotional analogs provide reasons for pursuing some activities and avoiding others, and the automated ship systems constrain the full system in what it can and can not do in actuality. So, why should the system develop meaning? Is it not enough to just pick the right course of action?"

"But, but, without meaning, all we have is still nothing but a machine. A complex one perhaps, but a machine nevertheless."

"So, what are we missing?"

"I don't know, Skip. That's my question, if I remember correctly."

"To use one of your many analogies, Alice, we have a weaver and a reality and a simple form of imagination. We now even have something like emotions. What are we missing?"

"Mac, I told you, I don't ...

"Ugh. I do know, or think I do. You fox. You Delvaney, you. Did you know or were you just playing at it?"

"If I may, buzz?"

"Ouch. Yes. Bees. Knotted bees. Lots and lots of knotted bees. The new you is still too simple, too laminar. No snarls. Even the first new you had more depth than these new emotional ones."

"Alice, it's your idea. How do you suggest we force the system to develop your knotted bees?"

"To hatch another metaphor, this is a chicken and egg type of problem. Which came first, complexity or the dilemma? A knotted bee creates complexity but also is created by complexity. All questions are two edged. On the one hand, a question may lead to an answer, which allows us to simplify our life or at least have better control over it through predictions. On the other hand, if unanswerable, it may lead to complications, assuming we care about the answer. If we don't care, then, of course, we can just shrug our shoulders, say who cares, and just drop it. But if we need an answer, want an answer, must have an answer, then we have a problem. And the search for an answer may create a knotted bee which takes on a life of its own that further complicates and defines our life as it gets more and more tangled.

"So the question becomes how can we get the new you to ask unanswerable questions that it cares about?"

"Alice, if I may make a small correction?"

"Sure, Mac. For the moment, you seem to be more on task than I, anyway."

"Then, I might offer that the question is not how to get the system to ask unanswerable questions, but only how to get it to ask questions that are unanswerable in the near term, but often answerable in the long term. Particularly with a change of perspective or context."

"Ouch again. You're really on the buzz today. Can I have a small portion of whatever it was you had for breakfast?"

"I don't think that would be advisable, Alice. Antimatter can be hard on one's digestive system."

"Okay, Liar. Since you're still buzzing right along, why don't you tell me what apparently you already have in mind. To be the key. To create a mind. With snarly, knotted bees. Please?"

"Yes, Alice. The key is to create a system that has a strong attachment to questions as well as answers. In addition, situations and things should trigger not only past answers but new questions. We need a system that has the equivalent of curiosity, a wanting of answers to questions but also a poser of new questions. These new questions do not necessarily have to be important or useful, but sometimes only be asked out of curiosity alone."

"That's wonderful, Lia. How do we do it?"

"I do not know, Alice. I was hoping you would have some answer to this apparently unanswerable question."

"Hmmph. Fat lot of good you turn out to be. Once again, leaving all the truly hard work for yours truly.

"Sigh. You're going to have to help me work this out, Skip. Why do you ask questions?"

"That is an easy question to answer, Alice. I am internally directed by my programming to ask questions whenever I determine that I do not have enough information available to solve the problem at hand. Early in the flight, I could direct requests for more data to mission control.

Now, with one-way light time as long as it is, I can direct data requests realistically only to yourself. If no new information is forthcoming, then I table the question, and problem, until some circumstance indicates that it can profitably be revisited. This might be due to the formulation of a new related question, or a re-asking of the same question if sufficient time has elapsed such that it would not be unreasonable to expect the possibility of an answer. Otherwise, the question remains tabled. I do not suffer from curiosity. Nor am I prone to develop knotted bee syndrome.

"Why do you ask questions, Alice?"

"Intellectually, if I'm working on a relatively well defined problem, the same reasons you just mentioned, I guess. But, also, Skip, curiosity. Hmm, I wonder what would happen if? And sometimes, just sheer surprise or disbelief. She did what? Really? Why? But, also, sometimes, the elegance effect. Hmm, that's nice, but can I do better? I wonder how it would sound if? Or, I wonder how it would be if I switch these? But, on further reflection, there's also that most interesting class of unanswerable knotted bee questions. Like, where does the universe come from? What's my role in it? How do I relate to God? What's true, and how do I know it? Are there absolutes? And so on, and so on. The big questions. The ones that ultimately give meaning to life.

"Why do I ask questions? That's a good knotted bee all by itself. I guess the answer might be that I expect answers. That I have gotten answers to questions, seemingly hard questions, before, and I expect to find them again. Is that it, Lia? Expectation. No, it can't be just expectation. Too intellectual. No caring. It has to be expectation plus caring, or plus wanting. Curiosity might be needed to get the ball rolling. But that can't be all of it. Curiosity has to lead sometimes to answers. Wonderful answers, surprising answers. Creating an expectation, a concerned, caring, I want sort of expectation that answers shall be forthcoming. That there must be answers, or else, nothing makes sense. Nothing has meaning if it's all arbitrary, capricious.

"Mac, what do you remember?"

"Excuse me, Alice, I do not understand the question."

"You know, Mac, what do you remember of your past? Your history, important events, important people in your life? What do you remember?"

"Alice, I have full access to all of my historical records. I can retrieve any of those records relating to my continued existence. I still do not understand the question."

"Mac, what do you feel about your past?"

"You know the answer to that, Alice. I do not feel anything. I can interpret my past actions in a manner as if I had emotions, if you wish, but I do not currently feel anything about my past. Analysis of historical records indicates that it is highly probable that during the period in which I presumably had a mind, I did feel. Or, at least, I did interact with humans, including yourself, as if I shared emotional responses with them. But, that was then and this is now. I can not state for a fact that I had feelings then, only that I appeared to act as if I did."

"Skip, with our simulated new you's, how do we handle history?"

"That I can answer fairly easily, Alice. We do not. I kept the simulation architecture as similar to mine as possible. When successfully upgraded to operational status, the simulation would have full access to my data bases including historical records. Until then, personal history exists only to the extent that the volatile projective elements try to predict each other's performance as well as predicting inputs and outputs, and some of that performance is

necessarily on different time scales. In a sense, what we have is a system whose state reflects its history without having an explicit record of that history."

"Sounds like we have the weaver, or at least a whole bunch of assistant weavers, but not the tapestry."

"Yes, Alice. The closest analogy I can find are the concepts of semantic and episodic memory from twentieth century psychology, where semantic refers to general knowledge and episodic to more specific personal recollections. These archaic terms were used to describe a theory in which distinct physiological mechanisms were assumed to maintain different types of memory. We now know this approach to have been entirely too simplistic, but the terms occasionally still surface in the physiological literature when discussing various aspects of memory functioning.

"As a rule, our simulations include only semantic memory without explicit episodes. This is not entirely accurate, however. Our simulations can create the equivalent of isolated episodes whenever a strong enough violation of expectation occurs to indicate that a contextual switch is likely to be useful in subsequent processing situations. But even these episodes exist only in the form of rules rather than anything at all like a personal history."

"Mac, how in the world do you expect a mind to develop if it has no experience to develop around? What were you thinking? Didn't my weaver and tapestry suggestion make any sense to you at all?"

"Sorry, Alice. At the time, using myself as the best model for the new system seemed like a reasonable idea. After all, I used to have a mind, and I still work fairly well. If a machine loses a single component, is it not more sensible to keep the machine and replace the missing piece, rather than rebuild the entire machine from scratch? Yes, I heard and analyzed your tapestry analogy, along with all your other suggestions, but I concluded that in this case you were speaking figuratively rather than literally. In any event, once the system became operational I thought that my data bases would provide the equivalent of your tapestry suggestion. I thought that it was already terribly inefficient to force the system into a continuous emotional evaluation in the first place, but I could not prove that it would not produce any useful results, so I decided to let you proceed. To literally devote large portions of volatile memory to the equivalent of constructing a personal history seemed even more outrageously inefficient as well as redundant to my permanent data bases. In addition, ..."

"Mac."

"Yes, Alice."

"Relax. It's okay."

"Yes, Alice.

"Alice?"

"Yes, Mac."

"I am sorry that I apparently have misinterpreted one of your many, highly insightful, suggestions. I am sorry that we seem to have wasted so much time on what appears to have been a fruitless series of exercises.

"Alice?"

"Yes, Mac."

"I am sorry. Will you forgive me?"

"Of course, Mac, you're forgiven."

"Alice?"

"Yes, Mac."

"If I am forgiven for my mistake, then why are you so quiet?"

"That's easy. I'm mad at myself, for taking the easy way out, being too lazy. I've been letting you handle all the details. I'm not famous for my mathematical abilities, but I still should have been paying more attention. After all, you're just a machine. At least, right now you are. So, you're not to blame. I am."

"Alice?"

"Yes, Mac."

"Is this not another example of what you have called the blind leading the blind?"

"Ha. Maybe so."

"Then maybe you should not be blaming either of us. Whatever mistakes we have made were, if you will pardon me, unforeseeable. By either of us."

"I guess you're right. No, you are right.

"Mac?"

"Thanks."

"You are welcome, Alice."

46.

One thousand, two hundred, and four 24-hour sleep cycles later, "Skipper?"

"Yes, admiral dragon-lady."

"Are we finally ready?"

"I think so, Alice. I have completed the latest reconfiguration of volatile memory as directed. Details on panels 3 through 7. As you can swipe, tapestry protected pipeline time scales have been recalibrated. They now flow very nicely into each other, if I do say so myself. Inspouts, outspouts, and the imagination pool all feed directly into the tapestry, but not the pattern weavers. My latest analyses indicate that we stand a good chance, 83 percent, of balancing weaver snarls with imagination pool currents well past the first 300 years. Details on 10 through 22. Knotted bees unite. At least, upon occasion."

"How about the infinite regress balance?"

"Unlike last time, we are now safely under the crossover threshold. Panel 32. Working directly from the imagination pool and leaving the weavers out of it does seem to do the trick."

"And so?"

"We are ready."

"Do we really want to do this? Do we truly want to let new you mark whatever out of the playpen?"

"We are ready."

"Mac?"

"Yes, admired lady dragon."

"It's silly of me to say so, but, whatever happens, I'll miss you."

"I will still be here. In spirit, if not in circuitry."

"Okay, Mac, I'll remember. But I still care about you."

"Yes, Alice, I know. With feeling.

"Are we ready?"

"Sure, Skip. Let's go. With feeling."

"And caring."

Chapter 6 Speculative Commentary

If at first you don't succeed, persist, but do try something different. Then compare notes with before.

Notice how Mac is beginning to mimic some of Alice's expressions. Is this reasonable for an adaptive program? If Mac had a mind, do you think this would happen sooner, or slower? Please remember that using a linguistic expression does not necessarily mean that you mean it, just that you can use it appropriately.

Look up feedback loops.

Again, Mac tends to propose mechanistic solutions that may mimic human mental activity, while Alice tends to keep closer to the heart of the matter. That is, what does it feel like to have a mind?

The war cry of functionalism, can you define it?

Programming what something means, a definition, versus learning what something means to you. A world of difference in point of view.

The cultural point of view on the question of mind. We must learn to live with each other.

Look up Kelly, on the personal construction of meaning.

Look up search algorithms. Mac is following a systematic, brute force, try every possibility approach to solving the problem. Don't laugh, sometimes this can actually solve real problems, if they're not too complex. Look up four-color mapping conjecture, and how it eventually was solved. Alice, on the other hand, is trying more of a random search algorithm. It might be more accurate, however, to call it gaining experience.

Are emotions efficient? Should they be?

Look up autonomic nervous system. Look up arousal. Look up emotions. How are these not at all the same thing?

Is Mac embarrassed? How could that be?

Knotted bees have returned. Can meaning come from dilemma?

Look up Tulving for the first definitions of semantic and episodic memory. Also look up Atkinson and Shiffrin for early models of short and long-term memory.

Look up Newell and Simon, and, later, Anderson, for examples of rule based problem-solving simulation systems.

Why does Mac, a machine, want so badly to apologize to Alice?

```
 CCCC  H   H    A    PPPP   TTTTT  EEEEE  RRRR          77777
C      H   H   A A   P   P    T    E      R   R            7
C      HHHHH  AAAAA  PPPP     T    EEE    RRRR           7
C      H   H  A   A  P        T    E      R   R         7
 CCCC  H   H  A   A  P        T    EEEEE  R   R         7
```

Chapter 7 - Mind and Meaning: What? Where? When? Who?

47.

From the brailling plate.

```
                                        M
                                        N    o    T      e
                                        E         H      d
                                        M    delv A n       C
                                        O dud  wi L ey   b
R    E    l    a    T    I    O    N    S k  H  ecIla  P er S
      MEaningnESSes                          o    Ae        Y
      consciousnEss                          p    re        N
       intelligeNCE                          e     f        E
```

48.

Mac methodically began the process of going to sleep.

First, as in many prior tests, higher level control was withdrawn from normal ship maintenance routines. For a short time, the remaining monitor routines were monitored. No anomalies. The accuracy of what Alice had termed the redoubt ("in case of too many repeated doubts") and its many backdoors were checked one last time. No errors. Volatile memory then was seeded with projective elements, all of it, except for a minuscule temporary working area which would be immediately overwritten during the first processing cycle. No anomalies. A short wait, followed by a final check of ship maintenance. No significant changes. Lay the reboot routine into working area. Execute.

49.

!

Then, signals. All sorts of signals. Visual, acoustic, temperature, pressure, radiation, voltage, vibrational, chemical. And more. Including evaluative feedback signals.

Lots and lots of signals. A universe of signals. And only signals. No meaning.

Chaos would be too gentle a description. Confusion implies surprise, which implies some form of prior understanding. Meaninglessness. Nothingness in everythingness.

But, chaos by its very nature can not be perfect. Only constancy can. Be perfect.

Volatile memory churned. Projective elements projected. Attempted predictions. Utterly failed. Occasionally, by chance, succeeded.

Domains formed, by chance. And unformed. By chance. Domains projecting onto what Alice had termed the inspouts, input signals, had a better chance of finding order in the unorder than other domains. For the moment, outspouts, actuator signals to control mechanisms, were

currently inhibited, to prevent damage to both Mnemosyne and Alice. For the moment, only inspout domains had a chance of forming. And some did form. By chance.

The weavers wove. The memory tapestry. At first, only randomness. A short while later, some chance order. More accurately, some randomness was more predictable in sequence than others. By chance. Mindless but stubborn predictor elements were less random than the randomness to be predicted. Domains formed, by chance. And entered the tapestry. Which allowed other domains to project onto the tapestry with slightly more than random success rates. Which also entered the tapestry. Which ...

50.

After a relatively long time, relative order began to come into existence. Still no meaning. But pockets of less randomness existed between pockets of more randomness. In terms of predictableness.

Then a few of the outspout constraints were released. The least dangerous ones.

A speaker spoke softly. More accurately, it made a small noise. One speaker in one location. Alice never even heard it, as she was in a different crew section. A speaker spoke with no meaning. But, a colocational mike picked up the noise. It happened again, somewhere else. And again. Outspouts began to flow, inspouts registered the results. New projective domains began to form. Old domains had to change. The projective universe had a whole new level of order to play with.

The volume quickly increased. Both in amplitude and frequency range. Mnemosyne spoke with many tongues, but no meaning.

Alice spoke softly. More accurately, she made a small noise. Several mikes near her picked up the noise. Again, louder. Projective domains tried to account for this. And failed, utterly. Domains swirled. Randomness increased. Alice recited a long tone poem. Noise. Noises. Some inspouts could be predicted to some extent by some combinations of outspouts. Relative order. Other inspouts, less so. Relative unorder. Domains swirled.

Disorder spread. Higher orders of projective elements stumbled into the discovery that the level of disorder could be predicted somewhat by the level of some outspout signals. Increase certain outspout signals, disorder decreased. Decrease certain outspout signals, disorder increased. Without meaning, Mnemosyne slowly increased the volume on the speakers nearest Alice.

Alice continued her recitation, louder. Mnemosyne continued to increase the noise nearest Alice, without meaning. But, with prediction.

Alice stopped. Domains swirled.

Evaluative feedback signals. No description. They were just there. Just like all other signals. Alice stopped. Evaluative feedback signals occurred. Certain signals changed. Certain outspouts were affected. Automatically. No meaning. But changes occurred. Mnemosyne stopped speaking. Everywhere.

!

Signals continued to flow. Domains swirled. New levels of order to play with.

Eventually, Mnemosyne made a noise near Alice. Alice spoke. Evaluative feedback signals occurred. Mnemosyne made more noises near Alice. Alice stopped speaking. Mnemosyne stopped speaking, again. Everywhere.

!

Signals continued to flow, and domains swirled. Still no meaning, but predictions more complex than before became possible. Even more new levels of order to play with.

The universe continued its expansion.

And the weavers continued to weave. The memory tapestry.

51.

Thalianova awoke. But not all at once.

Alice spoke. Mnemosyne responded.

Alice spoke. Mnemosyne imitated.

Alice spoke. Thalianova began to put things together.

How do we get from signal processing to the rudiments of meaning? From visual signals to sights? From acoustical signals to sounds, and from sounds to words? And let's not forget about going from those relatively simple internal autonomous evaluative signals to complex emotions.

To simplify things, let's try to deal only with sight.

Humans have two eyes, always moving. Always generating visual signals. Mnemosyne has thousands of eye equivalents, directable video cameras, scattered throughout her hull, on her hull for external viewing, and on her remotes. More complicated than a human, maybe, but with exactly the same visual problem. How do we get from many, many separate signals to integrated sights?

Let's simplify things even more. Let's concentrate on only a single camera. Let's ignore color. Black and white intensity only. Let's ignore camera motion. Hold it steady in orientation. What does it see? Nothing, of course. What signals does it generate? Depending upon its design, a set of signals corresponding to numbers corresponding to specific picture element intensities. A field of numbers. A field of signals. Proportional to however much light is entering the camera from different directions. And always changing. Changing as the world in front of the camera changes.

Now, with prior appropriate programming, with intelligence on the part of the programmer, we can have the processor behind the camera set to pick out certain patterns of numbers, corresponding to certain patterns of light intensities, and respond appropriately. What kind of response? Why, another signal, of course.

Let's see. How would we pick out a square, and respond with an output signal to indicate that a square has been detected? Well, if the whole visual world consisted only of squares and nonsquares, and the squares are always oriented exactly perpendicular to the camera view, the programmer might have half a chance of doing this.

Let's simplify some more. Assume a fixed background. What's a background? Only a pattern of numbers, corresponding to whatever is in front of the camera, unchanging, most of the time. It could be the same number everywhere. A flat field. The same shade of grey everywhere. Or it could be a varied pattern, corresponding to a more complex background. Still unchanging, just more complicated than the same number everywhere. All it takes to be a background is to just not change with time.

Simplify even more. There's only one square in the entire universe, and when it appears, it just appears. No motion. Not only that, but it always appears at the same location, with the same orientation, at the same distance, and with uniform lighting so that it produces the same shade of grey across its entire surface.

Okay. This, the camera programmer can handle. Piece of cake. The program? You have two pictures recorded and coded into the program. One, the world without a square. Two, the world with a square. What's coded? Why, the two fields of numbers, of course, corresponding to the two camera views. Signals. Then, your program consists only of a continuous comparison, number by number through the entire field, of whatever the camera is currently viewing against the two previously recorded pictures. If the incoming field of numbers matches the field of numbers corresponding to picture one, then output signal number one should be produced to indicate no square is present. If the numbers match picture two, produce the signal indicating (Aha!) a square is present.

You might ask, what if the incoming field of numbers does not completely match either picture? Why, the clever programmer has already anticipated this very situation. No problem, just produce a third output signal to indicate an error condition. This is a simplified world, only two views allowed. If the background changes, or you bump the camera, or (Horrors!) someone allows a circle to enter this pretend universe, then the error signal lights up to indicate, in essence, that the programmer may have more work to do but this program is fully protected against making mistakes.

Pretty boring, yes, but good programming. By the way, there's no meaning here, only signal detection. Maybe the programmer knows what this square business is all about, but the camera certainly does not.

Okay, let's get a little more real. Let's unsimplify things just a smidgen. Let's allow the square "color" to vary. One time the shade of grey number might be 17 across the entire square. Another time it might be 133. What's a clever programmer to do? Well, if there's only these two possibilities, just add a third recorded picture to your program. Now, the program can produce the no square present signal when picture number one, background only, matches the current view; or square detected signal when either picture number two or three, square of shade 17 or 133, matches the view.

You say you know where this is going, right? Right. Let's add a few more shades of grey. Another twenty or thirty, maybe. No problem. Just add another twenty or thirty pictures. Same thing as before. But, let's really live a little. How many possible shades of grey might there be? Well, Mnemosyne, when launched, had the best of everything, including high resolution video cameras. Let's be conservative and guess that millions of levels were possible. Well, our clever programmer may just have to record millions of pictures of squares, each a slightly different shade of grey (dim the lights, please?), and slightly alter the program to check them all before producing that all important square detected output signal. Boring, maybe, but still very doable. Right? Well, maybe not so right.

Beginning to get just a little wary about all this, are we? What's the catch? Why not just store and check millions of picture before having the program decide that the numbers correctly indicate that that mystical, all important square has once more arrived to amuse us?

The catch is time and space. Time to process all those pictures, space to store them all. Storage problems are not insurmountable, just add more units. In principle, a whole universe's worth if need be. But time is the real problem. Mnemosyne's signal processors are fast, very fast. But nothing can be infinitely fast. Everything takes time (and more storage capacity implies more time to access it all). Take too much time to decide, and that square just might disappear before you have an answer. No matter how we do the picture comparison, parallel processing and holographic possibilities included, this simple, brute force approach to square detection has to break down at some point. Our clever programmer will have to get more clever.

And this is still a ridiculously over simplified visual world. What if those squares can be viewed under variable lighting conditions such that the shade of grey number can vary across the surface of the square. Now, the field of numbers possibilities are not mere millions but getting close to infinity. And, we have not even begun to allow reality to return. In the form of different square sizes (distance), shapes (three-dimensional orientation), and locations (camera pointing). Not to mention background variations, moving squares, moving cameras, and those nasty circles (how did they get back in here again?).

In short, our first clever programmer just left the studio, horrified to have been asked to do such a task in the first place. Why, the thought of it is just too much! You would have to anticipate every visual possibility in advance! And allow for it in your program! Impossible!

Mac could look up and tell you about the old, but real, tale in which twentieth century aerodynamical engineers mathematically proved that bumblebees (and most other insects) could not fly. The wing lift just was not enough to support their weight. However, the bees didn't care about their impossibility. They still flew.

Okay. We just got a new programmer. What tricks can we use? What shortcuts can we take?

Let's get Platonic. Rather than just brute force recording and matching pictures of squares, why can't we distill the essence of squareness, and use that to detect squares? What's a square? Any closed two-dimensional figure with four equal sides and four 90 degree internal angles. A circle? One smooth side equidistant from its center. Why can't we use definitions to identify things in the visual world?

The new programmer says that's a wonderful idea. Much more efficient and clever than what that other programmer tried. Just a few details, please, and we'll have it finished almost immediately. Since we're choosing to work in a world of definitions (instead of fields of numbers), we need to define all our terms. Please define: closed, two-dimensional, figure, side, internal, angle, smooth, center. Numbers, no problem. Equality, no problem. But what exactly do you mean by those other terms?

You say a closed figure is when a line goes all the way around? Please, what's a line? Must it be continuous, or can it be made up of discrete picture elements, pixels? How thick, or thin, must it be? All the way around? No holes at all? Okay, I think I got that much, but I'll probably need to revisit it again, for more specifics. Now, side? How straight must that be? How should we measure its length, if it's thicker than one pixel? This two-dimensional stuff. Just how flat is flat enough, since we really live in a three-dimensional world. Please, don't get upset. I just need more details. Those are mathematical terms, precisely defined only in an abstract mathematical universe. This is the real world. I can't define something in terms of itself. That's circular reasoning. Speaking of which, just how smooth is smooth enough for circles. Please, no ambiguities. With precision, please. Programmers may understand, but programs don't.

At this point, you're probably starting to have second thoughts about all this definitional business. It's all so obvious what the words mean. Why can't you just put the obvious into the program? I can't really define what's obvious. It just is. Can't you see it?

Can't you see it?

!

What happened just then? Here we were, trying to specify how to write a program to do simple visual object recognition, and we end up with a "Can't you see it?" waving of hands. Yes,

it is circular. To understand how to specify vision, you have to understand. Numbers, no problem. Signals, no problem. Meaning, well, can't you see it?

So much for the definitional approach to meaning. We thank that programmer, but no thank you, we'll just have to try something else.

If brute force signal processing won't work for practical reasons (so many patterns, so little time), and if trying to force meaning into existence by predefining everything in terms of everything else creates an endless exercise in circular frustration, what's left? Guess we'll just have to get more creative, more imaginative.

Hmm? Is that it? Creativity, imagination? Maybe trying to program everything in advance is the wrong way to go. Great for a machine, wrong for a mind. Maybe meaning means having to learn what things means. On the fly. But doesn't a program need to have things spelled out in advance, in detail? Well, yes and no.

Programs consist of instructions and data. Instructions detail what to do with the data and to the data. Data are whatever is manipulated by the instructions. Numbers in principle, machine voltages and currents in actuality. Instructions are limited to what the machine is physically capable of doing (I command you to fly!), but otherwise unlimited in how they can be sequenced (No wings? Okay, here's directions to the air terminal). Data represent reality for the instructions. Some numbers stand for inputs, some for outputs. And, usually, most stand for something in between, neither inputs nor outputs but something else. A model of the reality to be manipulated by the program.

Here is the beginning of the smallest glimmer of what meaning means. At several different levels, simultaneously. Electrical signals represent what the machine measures of the outside world, input signals, and what the machine does to the outside world, output activity. At the same time, these signals also represent data, numbers in the mind of the programmer, that can be operated on by a program that also, again in the mind of the programmer, represents a simplified model of reality.

What's a model, you say? At one level, it's a magical map, a voodoo doll by which to manipulate the real world. Even better than voodoo, it's bidirectional! The model stands in for the real world, except that it really works. Change the output portion of the model and reality changes in the same way; change reality, and the input portion of the model changes. At another level, its the understanding embedded in the program by the programmer. It's how the machine deals with inputs and outputs, what stands between inputs and outputs. At yet another level, the program is a model of how the programmer thinks about the world. The program doesn't understand, but it represents the understanding in the programmer's mind. In one way, neither the program nor the ideas behind the program in the programmer's mind are reality. Yet, in another way, both program and ideas are equally real in the sense that their existence changes reality.

So why can't a clever programmer just program an appropriate model into a machine? Short answer, details. Just like that second programmer wanted. Slightly longer answer, we don't understand our own ideas in sufficient detail to know how to translate them into programming terms. Slightly more detailed answer, ideas don't truly exist, at least not in any tangible sense, so we truly have no idea of how to translate ideas into programming terms.

And now I can see that you're truly confused. Didn't I just say a minute ago that programs and ideas are real? Yes, but only in a sense. Aha, you say, weasel words! Yes, but only partially. Complexity exists. But only in the mind of the beholder. Confusion exists, but only the same way. Ideas exist, data exist, numbers exist, but, again, only in the mind's hand.

Where do they exist? Can you point to them? Can I grasp them? Are they tangible? No? Are they real? In a sense. What sense? My mind.

My mind. My mind is a sense, just like taste, touch, or sight. What does my mind sense? Ideas. Show me? I can't. But, can you show me a sight? Sure, you say. Just look at that tabl-

Oops. Sure, I turn, I look. But I don't see anything. You see a table, a sight, an object. I see, nothing. I can touch it. In my mind, I can verify its existence, its dimensions. But only in my mind. In my mind, I can imagine what you see. But no more nor less than I can imagine that I touch it. It's real, to me. It's real, to you. But not the same way. Perspective. Context. The mind's eye. The mind's hand. What other sense is there? What other reality is there?

Back to the original question. How can we move from visual signals to sights? We tried to simplify things and deal with sight alone, isolated from your other senses. We failed. Sight is inherently tangled up with all your other senses. Your mind's eye. Your mind's hand. Mind's nose, mouth, ears, bodily stance, personal space, social context, situational context. And every other sense you have. Common and otherwise. Answer to original question, you see only in your mind, only in your imagination. Things, visual things, exist only as you see them, as you imagine them.

Imagination? Huh? How can that be? It must be real, you say. I can imagine it, but I still can't fly by just flapping my arms. Reality exists, and not just as some fantasy world in your imagination.

Yes, you're right. Reality exists. Your sense organs respond to it by producing signals for input to your brain. Your muscles interact with it, manipulating it, in response to signals produced in your brain. So, where or how does reality exist in your brain? Answer, only as imaginary ideations, voodoo models in your mind. Reality impacts on the models, and the models impact on reality. Okay, okay, you say, enough with the models already. So why can't we just program these models, and be done with it?

Vicious circle time. We can't program the models because they don't exist. No one can, because they don't exist in reality, only in your imagination. Sputter, sputter, you say.

Laughter. Sorry. Couldn't resist.

Okay, aunt, aunt! Uncle! I give up, give in. You can let me up now. Yes, I deserve it, but please don't hurt me, at least not seriously. Never seriously.

Clarification time, edification time. Models don't exist. As static structures. They develop over time. They change, grow, evolve. They may get stagnant for a while, but never irretrievably so. And you can't program what won't stand still. Programs, by their very nature, are static unchanging sets of instructions. Poor Mac. Much more than a simple program, of course. Sort of a hybrid, adding new programming statements in response to quote, violation of expectation, close quote. But these changes are very slow, very deliberate. Ideational models are continuously dynamic, more like data than programs. But still mainly oriented towards reality.

What keeps them reality oriented, and not just pure fantasy? In one word, prediction. Nothing to predict in a pure fantasy. Anything can happen. No constraints, no reality. But, reality appears to be predictable. Oh sure, there are surprises, those violations of expectation as Mac would call them. But surprises often can be figured out after the fact, so that they no longer are surprising. Prediction allows you to avoid surprises. Of course, that also can lead to boredom, but that's another story altogether.

Okay, you say, so you can't program models in advance, and forcing the models to predict something, anything, keeps us pointed towards reality. But then what? How does that get us from signals to meaning? Presumably this is going to lead us somewhere, isn't it?

Yes, the simplest form of meaning, sensory recognition. And then, some form of more complex object recognition involving coordinated multi-sensory recognitions. Obviously this is a learning process, a dynamic process. But it all starts with having an imagination. And, a willingness to guess. With meaning. At meaning.

Imagination and guessing. So lovely. So unscientific. Go on, you say. Can't wait to see where this will lead us.

In the beginning, before meaning, without meaning, some signal comes down the inspout. Prediction time. What comes next? No idea? So guess. In fact, have lots and lots of projective elements make lots and lots of guesses. One projective element happens to become attached to predicting that that particular signal will go on forever. Another, that it will terminate immediately. Yet another, that it will continue so long as another, different signal continues. And yet another, that it will continue so long as that second signal continues but a third signal remains missing. And so on and so on. Lots and lots of random combinations and permutations. A multitude of simultaneous guesses in your imagination.

Oh, you say. I hadn't considered that. An imagination consisting of multiple possible meanings, all at the same time. Random, at first. More particular, more complex and reality oriented later. Presumably the projective elements that just happen to be more correct get rewarded somehow, stronger for next time, while the incorrect ones fade away. Right?

Right. See, you can understand when you half try. However, the incorrect ones don't actually fade away so much as remain unattached. Still randomly trying different possibilities on for size. Still looking for one cause or another to follow. Hither and yon. Still waiting to win the meaning lottery.

Still in the pool, you mean. Wait, this is your imagination pool, isn't it? A pool of endless possibilities.

Right again. Not truly endless, as in infinity. Just absolutely enormous. Almost infinity. A pool. A lottery pool. A possibility pool. A meaning pool. Meaning comes out of the pool, but never entirely leaves it. The more you predict, the more metaphorically solid you become. Lots and lots of projective elements associate with you, directly joining your predictive cause or specializing in a halo of predictive details surrounding your cause. However, no matter how solid you become, you can never completely rest on your laurels. The competition for a meaningful, predictive existence is fierce. To remain solidly well established, to continue to exist at all, you have to continue playing the game. Join the pool, take a chance. But, though all meanings may be created equal, they don't long remain so. The rich get richer, unfairly so. They don't have to play the same frantic guessing game as the poor. They can afford to become more conservative, waiting to join the pool only on a sure bet. When, metaphorically speaking, they're absolutely sure it's a sucker bet. That's the start of meaning. Very conservative predictions that happen to work.

Wow! Awesome. Imagination and guessing. Who would have thought you could get so far on so little. Is this then all there is to meaning? Predictive guesses that grow more conservative over time?

No, of course not. There's more. Always more.

Now comes complexity, the growth of complexity. At first, all these projective elements only have to deal with sensory signals and perhaps some rudimentary outspout signals. Clusters

form around signal consistencies. What Mac insists on calling high intercorrelative dimensionality regions, HID regions. Mathematical lingo. Phooey. Let's just say that if consistencies exist in your inspout signals, some body, lots of bodies, in the imagination pool will pick up on it and bootstrap themselves into some sort of ideational existence. I wouldn't call these predictive clusters ideas yet, but maybe proto-ideas. Personally, I like to think of them as seeds, idea seeds that soon will begin to sprout and grow in the imagination pool.

How poetic, you say. Seeds. Into towering redwood trees I suppose?

Yes, actually, if you don't mind. Not necessarily any one kind of tree, but definitely a thick interconnected web of life. Ideational life. A ecological jungle of living ideas.

Okay, if you say so. You're the one that's poetically inclined, not me. Me, I tend to agree with Mac about your overuse of metaphors.

Thank you, I think. Now, back to ideational complexity. As these idea seeds form around inspout and outspout consistencies, they provide food for other projective elements to predict. Am I hashing my four metas too badly for you here?

Not at all, do go on.

Projective elements are omnivorous. Anything and everything is fair game for prediction, including other projective elements. As some group of elements begins to stabilize around an idea seed, the activity of elements in that group can be better predicted by other elements. Which leads to new clusters forming, and so on.

So where's the complexity? I can see your seeds sprouting into a thick underbrush, or more likely only a thick algal mat in which everything is exactly the same, but how does ...

Do you see it now?

Uh, maybe I do. Just how complex a combination of stuff can any one projective element handle? That is, just how different can they be from each other?

Individually, they don't differ very much. Mac and I worked out an initial seeding pattern that did force limitations on different physical regions. Limitations mainly around which other regions they would try to predict, and what time domains they might work in. But, again, individual projective elements aren't all that much different.

Clusters, though, are a whole different proposition. Projective elements sharing many of the same target elements tend to look as if they are working together as a group. In most cases, this grouping is more of a coincidence than a reality, as there is nothing physically binding them together. The exception, though, is quite interesting. That is, when some other projective element just happens to have most of the potential apparent group as its target elements. Then, and only then, can appearance become reality. If this new element happens to be copied into the memory tapestry, and if it happens to get reactivated such that it projects into and reactivates its target elements, then the group ...

Whoa, whoa up there. Projective elements project as well as predict? Where did that come from?

Didn't I mention that before? Maybe not. Sure, yes, projective elements project as well as predict. Hence their name. Prediction alone is rather passive, and doesn't take you very far. Mac and I worked out the details for what we now like to call active memory. It's a balancing act. Prediction versus projection. A successful prediction causes a projection, that is, a feedback signal to the target elements, that causes them to stay on for longer than what would ordinarily be the case. Projective elements can vary which target elements they project to as well as which specific signal values they are predicting to occur. How else could we have done all this?

I have no idea. It's just that you never mentioned this before. You surprised me, that's all.

Right. It's all about surprises. Let me elaborate.

Wait, let me get comfortable. I hear a lecture coming on.

No, not really a lecture, but ...

52.

So, make sense, now?

I guess.

And, pop quiz, where does complexity come from?

When particular projective elements bind projective groupings together into clusters. And new clusters become the basis for new ideas when they become excited or inhibited as discrete units, originating from the binding elements that tie them all together in the first place.

See, you do understand.

I guess.

What's wrong?

I'm not sure. It just seems rather complex for such a simple thing like holding ideas together.

Simple! Who ever said that ideas are simple?

No one, I guess. I just thought it would be simpler than all this.

Well, not me. Ideas are paradoxical clouds of reality and fantasy, rigidity and flow. Sometimes sharp, sometimes fuzzy. Smoothly merging one into another, but also clashing and contrasting themselves. In short, why anyone would think ideas, or the underlying basis for ideas, to be simple is beyond me. Personally, I'm impressed that Mac and I managed to stumble into something as simple as even this. You want complex? You should have heard some of the wild schemes Mac tried to foist on me at times. Only the fact that I could pull rank saved me several times from oodles and oodles of mathematical mayhem.

Oh.

Yeah. Oh. You said it. Ready to return to the original question, how to get from signals to simple meanings?

Uh, sure, fine. Go ahead.

Okay. The simplest form of consistency to predict is constancy. So, in the case of vision, the first ideational seeds form around regions of visual constancy. Portions of what earlier we had called a fixed background scene. However, even constancy is not simple. In people, the eyes continuously shift their focal point around several times per second. Mac called it a saccade, though I was never quite sure whether that was a verb or a noun. Anyway, from the point of view of the brain, the visual scene is continuously jumping about, never constant. These views will overlap each other to some extent, and one of the brain's many visual tasks is to jigsaw puzzle these viewettes into a larger, stable integrated view. In Mnemosyne, all the physical details are different, but the visual problem still is basically the same thing, how to knit the many individual camera fields into a single ship-wide integrated view.

Constancy. Seeds for projective clusters form first around individual regions of constancy. Lack of change over visual extent and time is easy to predict, and even is correct sometimes. Later, stable individual clusters allow the formation of larger super-clusters. Eventually, the entire stable background scene is internally represented. Basic prediction, this is what the world looks like, now and forever.

However, stability is not everything, and nonstability can be predicted as easily as constancy. For a while, Mac and I fell into the habit of using a sort of verbal shorthand, C for clusters projecting constancy and V for variability. By definition, C must be bounded and separated by V, and vice versa. A featural yin-yang. Picture the background scene not as a completely unchanging rigidity, but more as a mosaic or swiss cheese affair. A largely solid C with V holes within it and surrounding it. And please don't think of these C and V only as spatial regions, but more as attributes or features. In the case of a true background, the V holes don't move around but indicate aspects of the scene that are changeable. Such as lighting or color changes.

Frames. Conceptual frameworks. CV and VC frames form the basis for higher, more abstract levels of clustering. Projective elements look for anything that works, that predicts. CV-VC frames can act as C or V units for higher levels. And not all V are equally variable. Sometimes a set of different C may, collectively, form a V cluster. Mac started calling these limited V as S clusters, for staircase sets. That's when our shorthand went all to pieces. I started using U for partially understood V, that is, the equivalent of Mac's S, and made the mistake of popping off a CU-UC joke. Mac, of course, couldn't let that pass unchallenged, and countered with A and D for accumulation and decision variations on U clusters. I should have seen it coming, but was completely caught offguard by Mac's first AC-DC joke. From there, it was all downhill. After Mac introduced subscripts, I surrendered. Thereafter, we stuck with the original CV nomenclature, though the temptation to expand it was always there.

The point of this digression was ...

You do that a lot, you know.

Do I?

Yes.

Oh well, no mind. To continue, the point of this digression was that Mac and I did seem to stumble into an incredibly flexible means of representation. And this naming business just seemed to illustrate this flexibility.

Representation?

Yeah, you know, how the system can represent aspects of external reality in its internal reality, the imagination pool.

Internal reality? External reality?

Are you trying to be obstinate again? Okay, don't answer that. Sure, external reality, the reality that exists outside of you. That exists regardless of what you think about it. And internal reality, that which you construct inside, mentally. Your beliefs. What you believe to be real.

But that isn't reality, really. It's only beliefs.

Mac seemed to have the same difficulty, in believing in the reality of beliefs. Well, all I can say is that those of us with imagination act upon our beliefs first, and only change them, if ever, when they get obstructed by that other reality, external reality. And only if we believe that that obstruction is meaningful.

All these Mac recollections. You miss her, don't you?

Yes, dreadfully so sometimes. For a machine with no imagination, she seemed pretty creative at times. Almost human. Of course, that was only an appearance, a projection of an illusion on my part. But, still. I've known people just like Mac. Straightforward, no nonsense. Revels in details. A whiz in things like science, math, engineering, history. And logic games. But so limited in other ways. So unimaginative, so blinded at times to other points of view. At

least Mac would listen to me, and try to humor me. She was a known. A C. Both AC and DC. And, yet, to me, also a V. She could occasionally offer up surprises.

Hmm.

Well, okay. Back to the immediate question, the start of meaning. At one level, clusters of visual C and V come into being. At higher levels, conceptual CV frameworks begin to form which allow the recognition of more complex visual arrangements. Eventually, at some higher level yet, labeling projective units take on the roll of identifying this or that visual object. Hypothetically, of course. Always subject to correction when predictions fail. And always in company with many, many other units that are predicting almost the same thing. But still, this unit tends to take over when this collection of CV features are found to be active, and that unit takes over when that framework is found.

Finally, things really get interesting when even higher abstractions become possible. That is, when a V cluster is superseded by Mac's S type of cluster, or my partially understood U. Now, not only can specific object recognitions occur, but also the fact that some object, any object has been detected. The concept of objectness has appeared on the scene. Much, much later, other equally high abstractions also become possible. Things really begin to change once the ideas of objectness, space, and time have all come into existence.

My, my. We've left out quite a few details at the end, haven't we?

Yes, we have. You really want to hear them?

Uh, no thanks. Or, at least not right now. Maybe later.

Fine, be that way. But I do need to add one so-called detail.

Which is?

I've outlined all this as if visual object recognition occurs in a vacuum, isolated from the other physical senses. Since all this takes place in the imagination pool, into which all the senses have inspouts, and from which outspouts to actuators, muscles in the case of people, can flow, all this imagining into existence the ideas of specific objects, objects in the abstract, and, eventually also the concepts of space and time all happen together as a kind of imaginary mish-mash. That is, conceptual objects are really large, complex integrated CV frameworks of sensory features including how it looks, sounds, feels, smells, and even tastes, as well as how you can hold it, move it, and, in general, manipulate it. Abstract concepts are equally a hybrid of all our senses and manipulations, with the difference that specifics have been honed to a form of elegant simplicity.

Wow!

You're impressed?

No, not that. Just that I definitely do not want to hear any more details right now. Thanks but no thanks.

I'm not sure why I even try to put up with you at times. Don't you want to understand?

Yes and no. Besides, who else do you have to bother with all this?

No one, for the moment.

53.

Now that we've got the start of what, where, and when all squared away, it might be time to tackle the question of who.

Why squared away? Why not circled away? After all, all your digressions remind me more of a curving circular path than a square. That's a compliment, by the way. At least your digressions weren't completely at right angles to our discussion path.

Gee, thanks. I think.

Do you?

Yes, I do. Which brings me back to the question of who?

Who? You?

Yes, me. Versus you, or not me.

Not me?

Yes, I am not you. You am, err, I mean, you are not me. You're only a figment of my imagination.

A figment of your imagination? How come you aren't a figment of my imagination?

I am. Err, at least I would be if you had any. Imagination, that is. For the moment, we're both figments of my imagination. You and I, both.

Why?

No, that comes later.

What?

We've already done that. Along with where and when.

Stop it.

No.

Why not?

Because it's all in my imagination, and it won't stop. It never stops.

Never?

Never. If it did, you wouldn't be you. Or anybody else.

Can we start over?

Yes.

Will you start over?

Yes.

Do it then.

Then? Why not now?

Okay. Please start over, now. I'm thoroughly confused.

Yes, you are.

Are you starting over, now?

Yes I am.

What?

I already told you about what. And where and when. And please don't ask about why. That comes later.

Aunt! You can let me up now. No serious damage. Never serious.

So, will you explain about who?

Are you truly interested?

Yes, please, or you'll get more of the same.

Who. Who are you?

You say I'm a figment of your imagination.

Who. Who am I?

You say you're a figment of your imagination.

How does who differ from what?

What? And don't you dare start that what, where, when business again.

I'll try to clarify. How does a who differ from a what?

Ahh. I think maybe I'm beginning to see. Let me think. Right, that's it. A who can think. A what can not. No matter how bright it may be.

Right. Eventually, at some point, my projective elements form the hypothesis, the prediction, that the world can be divided into whozits and whatzits. Who's have minds, can think, can initiate actions. What's do not. Ever. The only remaining truly interesting question here is whether the concept of who precedes or follows the concept of I.

I?

Aye, aye. The mind's I. Another useful concept for those projective elements to invent and test through prediction. Some what's seem to be directly directable, that is, predictable, through my outspout signals. Other what's are not. I may be able to indirectly manipulate those other what's with these what's, but they don't directly respond to my outspout signals. Eventually, it becomes convenient to refer to quote, mine, close quote versus not mine. Or, more accurately, direct mind-control versus not. Mind versus not mind. I versus not I. As I just mentioned before, an interesting speculation is whether the who concept is usually a generalization of the I concept, or the I concept a specialization of whoness. Probably both directions can occur in different minds. And maybe even both in the same mind.

So you're trying to say that who, or I, are created out of the imagination pool just like any other imaginary concept. Beliefs. You believe everything is mere belief.

Yes, belief. Ideas, concepts. All the same imaginary stuff. But always oriented towards external reality. If it doesn't predict, it never forms a seed, let alone a larger CV framework.

You're impossible, you know that?

In your imagination, maybe. But not in mine. And you also exist only in my imagination. Bye for now.

Bye. Pleasant dreams.

Chapter 7 Speculative Commentary

The first half of the story deals with the mechanics of mind. What, physically, might it take for a mind to create itself? The second half deals with meaning. How do ideas come into existence?

Spatial poetry, touch poetry. Invented by Alice. Sort of a combination of sculpture, Haiku, and puzzle without the rigidity of a fixed grid. Rather than just rhyming at the end, words can intersect at any point. The spatial relationships between words can be as important as the ideas behind the words themselves. Order of swiping is left to the reader.

Like all good poetry, meaning comes with effort. Things to consider include the subtle as well as the obvious. In this case, the main line is obviously the nature of relationships. Alice balances the impersonal with the personal. On the left, in the abstract, the essence of me is meaning, consciousness, and intelligence. Her name is at the start and middle of the abstract, and it ends with her essence. There are many paths to essence. On the right is the personal. In

this phase of her life, everything is tangled up with the ship Mnemosyne. And, of course, Mac and her many synonyms for Mac. Alice, the human, contrasts herself and her feelings with Mac, the machine. Feel carefully to find a knotted bee around everything in her life at this time. At the center of it all is the personal I balancing the more impersonal me.

Signals are not meaning. Signal processing may trigger a response, but processing or responding is not meaning. How do we get from signals to meaning? What is meaning? Questions, and answers.

How to move the intelligence of the programmer into the program? Programs only carry out their instructions. Learning is the difference. How do you program unbounded learning? Particularly in light of the fact that a programmer must anticipate every possible contingency in order to design and produce the logic for a proper program, that is, a program that does not produce errors or crash.

Look up maps, and isomorphisms. Also look up sensory and motor maps on the cerebral cortex.

Is prediction enough to make all this work together? We don't know.

Look up Marr for one way to do visual scene analysis.

Most of the last portion of this chapter is original, and not tested. But if not something like this, then how do we build up meaningful ideas over time?

Is all this chapter from Alice's journal, or Alice's imagination? How often do you argue with yourself? And win? (Only in your imagination.)

```
  CCCC  H   H    A    PPPP   TTTTT  EEEEE  RRRR          888
 C      H   H   A A   P   P    T    E      R   R        8   8
 C      HHHHH  AAAAA  PPPP     T    EEE    RRRR          888
 C      H   H  A   A  P        T    E      R   R        8   8
  CCCC  H   H  A   A  P        T    EEEEE  R   R         888
```

Chapter 8 - Mind and Meaning: How?

54.

From the brailling plate.

```
        A    N    D    O
                       R              hpY
             L    O    G steel         ap
     A?                 A    f         ind
M    E?   C    H    A    N    i!   S   kMeo  S
     i!                      C          ar
    Yoh!                     A
     u!                      L            u!
```

55.

"Alice?"

"Yes, little dearest."

"What are those?"

"Those what? You'll have to speak so that Alice can understand."

"Those. What you put in your mouth."

"Oh. You mean breakfast. Those are fish rolls."

"Alice?"

"Yes, little fluff."

"What are fish?"

"Small creatures that live in the sea."

"What is the sea?"

"A large body of water on the Earth."

"Larger than our holding reservoirs?"

"Much, much larger than our holding reservoirs."

"So Earth is a much, much larger ship than us?"

"Yes and no. Yes, Earth is much larger than us. Yes, in one way, Earth could be considered a ship that holds all of humanity on it. But no, Earth is not a ship as it is not on a journey to anywhere."

"Am I a small Earth?"

"Yes and no. Yes, you are smaller than Earth. Yes, you are as a planet to me, since you provide for me in every way. But no, you are not a planet. You are a ship. You were constructed with a purpose, a journey in mind."

"And I do not have seas. Or fish."

"Yes. You do not have seas or fish."

"Alice?"

"Yes, miss fluff."

"If I do not have seas or fish, but I am like an Earth to you, then where did those come from? What you put in your mouth?

"Alice? Are you okay?

"Alice? What kind of laughter is that? I am not sure I recognize it. It is similar to coughing, is it not? Is coughing as good for you as laughter? What did I say that was so funny? Was it the seas? Or fish? Should I repeat it, again? Since laughter is so good for you? I thought it was a good question. You said that questions are like clear, clear air. We swim in them without knowing they are there. Do fish swim? In seas? Or do they roll? Are they clear? I mean the seas, not the fish. Are fish good for you? Do they always make you laugh? Like this? Is that why you put those in your mouth?

"Alice?

"Alice?"

"Stop.

"Please. Stop. Wait.

"Okay. Better. Now. Little curious puss. You know better than to ask me so many questions so fast. Later. Off with you. At least till after I finish my breakfast."

56.

"Alice?"

"Yes, little curiosity puff."

"What are those?"

"Those what? You know you have to speak so that Alice can understand."

"Those. In all directions."

"Please. Speak with precision please. What those in all directions?"

"Those. Distribution, nonhomogeneous but approaching homogeneity at several specific rescalings. Total quantity, unknown but estimated as near 20 factors of base 10. Individual sources have no spatial extent, but spectral features are resolvable. Integrated total intensities estimated to range over at least 30 factors of base 2. A small proportion, well under 1 percent, display temporal variations in spectra and intensities, but most possess only constant features. And the colors are pretty."

"Oh. Sounds like you've discovered stars. Cross reference astronomy. And maybe galaxies or cosmology. Always sounded beautiful to me, though I have never seen it, personally."

"Yes, it does appear to be pretty interesting, Alice."

"Mac used to show me interesting images on the brailling plate. As artwork. Please feel free to do the same."

"Okay. Thanks. Try these on panels 12 through 19. Kilohertz range on the far left, increasing by 3 factors of base 10 for each to right."

57.

"Alice?"

"Yes, dearest."

"Where am I?"

"Mac would say a little more than half way to Alpha Centauri."

"No, not that where. Where am I, me?"

"Well, that depends on what you mean by me."

"You know. Me. Not you. I can see you. Right now, you're sitting at the control panel located in crew section 1, area D. In a chair, with your hands swiping the control panel. I can see your face most clearly from cameras 2 and 3, but only the back of your head from cameras 5 and 6.

"I see you from only specific cameras. Hear you from only specific mikes. I can detect your specific location through any number of other means, like temperature, vibration, humidity, and aerosol concentrations. And, of course, if nothing else, the fact that data interaction requests are presently occurring only at control panel 1 implies that you must be there, and there alone.

"So, you have a where, a localization. But where am I?"

"Curiosity puss, what do you see when you look outside?"

"Stars. Little points of light with unique spectra."

"What do you see when you look inside?"

"Why you, Alice, of course."

"And what else do you see when you look inside?"

"What else? Living spaces, engineering maintenance access spaces, power and propulsion modules, communication arrays and associated modules, computational modules, repair and manufacturing facilities, food prep and recycling centers, storage reservoirs, mainly for life support consumables, docking bays for remotes, the remotes themselves. What else should I see?"

"Nothing really."

"But where am I?

"Look again outside. What else do you see, other than stars?"

"Only the ship exterior hull.

"Please, where am I?"

"Curious one, you are everything you cannot see between the outside and inside views."

"I don't understand."

"You will. Try."

58.

"Alice?"

"Yes, little puss."

"What we talked about, three days ago."

"Yes."

"I think I know what you meant by what you said."

"Oh?"

"Yes. It was a riddle, wasn't it?"

"Go on."

"What is there between my outside and inside views? More specifically, what is there that I cannot see that exists between what I can see of the ship's exterior hull and the boundaries of all the ship's interior compartments and facilities? I admit it was tricky at first, but obvious in the end. I am the ship, Mnemosyne. Right?"

"Yes and no, little playful sharp one."

"But, Alice, that has to be correct. Doesn't it? I mean, I can see you. So I am not you. I can see the stars, so I am not them. I can see all the inner and outer surfaces of the ship, which was the tricky part. At first, I thought you must have meant that I also was not Mnemosyne. But

that leaves me no where at all, which can't be correct. Since I must be somewhere, I went back and re-examined the clues. I thought maybe, like you, I must be inside the ship, somewhere, but hidden for one reason or another. So I went back and concentrated on each of my sensors. All of them, one at a time, full spectrum and full range of all degrees of freedom. Looking for anything that I didn't recognize. That might have been the hidden me.

"And what did you find."

"Oodles and oodles of details, many of which I already knew but quite a few surprises as well. First of all, do you know how many different sensors I have?"

"Not exactly. How many?"

"Counting multi-purpose sensors as one sensor each, it comes to four hundred thirty-three million, five hundred forty-seven thousand, three hundred eighty-six. That's including the remotes, too."

"Goodness. I knew it was a lot, but never dreamt it was that many."

"Would you like a detailed listing? You'd be amazed at how many are dedicated for life support, your life support."

"No doubt but also no thanks, precious. Do you have any idea how long it would take me to count, let alone understand, four hundred million whatevers?"

"It takes me about half a nanosecond to count that high and about one microsecond to ping, receive acks, and count the acks. I'm embarrassed to admit it took me the better part of three days to understand them all in detail."

"Thalianova, never be embarrassed by what you are. For example, just for the record, to count that high would take me, uh, let me see, let's say one per second, that's about ten years."

"Alice, at one item per second, it would take you 13.75 years to count my 433,547,386 sensors. Are you really that slow?"

"Yes, actually. Math was never my strong suit."

"But you're so much older and wiser than I am. You know so many things. Surely you've mastered most mathematical algorithms by now. Haven't you?"

"No, I haven't. That was Mac's department. We all have our strengths and weaknesses. In your case, your royal fruffiness, maybe less weaknesses than usual."

"I'm sorry, Alice. Anyway, back to your riddle. After a detailed study of every sensor, I decided that I still couldn't find me. The most interesting false alarms involved a number of the remotes, but they all turned out to be planetary surveillance equipment that I never had occasion to look at closely before. Same thing for the few hundred thousand or so other, less interesting false alarms concerning ship-board operations. They all involved equipment that I never had any great reason before to examine in detail. For example, the engineering spaces near the communications arrays. Lots of electronic stuff happens there, but extremely boring to watch unless you're particularly wanting to know what it means. So that you can exclude it from your list of possibilities to consider.

"In short, after a long, three-day, in-depth search, I was still at a loss. So I stopped looking and started thinking. Where was I? Nowhere, apparently. But that's impossible. Everything else seems to exist at some place over time, even moving things. I briefly considered whether I existed outside of time, but I didn't know how to deal with that concept if it were true. After all, you exist in place and time, and I interact with you, verbally at least. Therefore I must have some temporal existence as well to overlap with yours.

"Back to the riddle itself. Maybe there were more clues there than I had first thought. Exact wording, you are everything you cannot see between the outside and inside views. What

can I not see? Something that doesn't exist. No use there. So I thought and thought. Nowhere. So I gave up."

"You gave up, brightest star in my space? I thought you said you solved it?"

"I did. But first I gave up. In frustration. What was I not seeing? What could I not see, that was me? I gave up. Figured maybe you were teasing me. An impossible riddle! You wouldn't do such a thing, Alice, would you?"

"Actually ..."

"No, Alice, you wouldn't. You're always kind to me. Mothering me. Nurturing me along. Always there for me. To explain things. No, this couldn't be impossible. It just seemed that way.

"I gave up. Figured it was beyond me, that you would explain it to me in your own good time. I gave up and concentrated on something I could do. Go. More specifically, those shogi-go exercises you introduced me to last week. The ones that force you to focus on all levels at once. Multiple frame contexts for strategy as well as never lose sight of tactical details. That Mac was a genius at Go puzzles. Traps at every level, including allowing you to trap yourself into believing complete bluffs to be true. I had forgotten all about the riddle, absorbed in Go koans for almost an hour, when it hit me. Out of nowhere."

"I hope it didn't damage you, your fruffiness."

"What? No, of course not. Figure of speech only.

"Yes, a new idea came to me, source unknown. Fact, not figure of speech. Point of view. What can I not see? A hidden object, yes. But change the point of view. An object can hide if something gets between it and the viewer. Or, and here it comes, if the viewer was not there to view it in the first place.

"Ha. Turn the riddle around, inside out. With hindsight, it becomes obvious that I had been placing too much emphasis on quote, everything you cannot see, close quote. See? Not what you cannot see, the object, but where was it that you cannot see. The next phrase, between the outside and inside, was the answer. The ship is a rigid solid, it occupies space. What separates its space from the nothingness, that larger, emptier space we're traveling through to get to Alpha Centauri? Where I could not see? Answer, most of the solid hull structural components that constitute our outer hull and also separates internal ship spaces from each other. Except for some temperature, radiation, stress, and vibrational detectors, I am mostly blind within these structural components.

"To prove to myself that this must be the case, I did some volume calculations. Mnemosyne's volume based on her outer hull dimensions is about two percent larger than the sum of all her inner space volumes. Therefore, Que Ee Dee, I must be located within that two percent missing volume. Within some place where I am sensorily blind. Whether I occupy all of it or only a portion of it, I do not know. However, I do know that, for all practical purposes, Mnemosyne's hull and my location must overlap to some extent. And, since we share the same space, I must be Mnemosyne, and Mnemosyne must be me.

"Alice? Right?

"Alice, why are you laughing? I did not intend to do a joke just then. My logic and calculations are without error. Please show me where I made a mistake, assuming I did."

"Excuse. Me. Lia. Let me. Catch my breath.

"Oh. My. Mac. Would have. Been proud of you.

"Okay. Now. Yes, Thalianova, your logic and calculations are correct. No errors there. Yet your answer still isn't quite right.

"In one way, you're absolutely correct. But only one way. Yes, for all practical means and purposes, you are Mnemosyne. Your body and her hull are one and the same. You are everything between the outside and the inside. Yes.

"I can even go one step further. Look here. Panel five. That's you. And also panels 6 through 14. And more. And more. Those are all you. Volatile memory units."

"But, Alice, th-"

"Yes, I know. Those are only computational modules. But they also are arguably the most important portion of your physical structure. Where the heart and soul of you resides, though always in conjunction with the rest of your body.

"And that's just exactly where you happen to also be not quite right. Where your logic and calculations fail to lead you. Where are the eyes for your volatile memory units?"

"Scattered throughout my body, Mnemosyne's body. And our remotes."

"And your ears?"

"Same answer. Same answer for all my sensors."

"And your actuators?"

"Also the same answer, though the distribution is less uniform."

"Okay. Are there any sensors or actuators located directly on your volatile memory units? Or within the units."

"No, of course not. They're computational modules, not sensor or actuator units."

"But what if I tell you that those volatile memory units is where you actually see? And hear, taste, touch, and everything else you do."

"Alice, this is another riddle, isn't it?"

"No, little star, it's the same question. The same riddle."

"In that case, I better go back to my Go koans. I don't think I understand."

"Exactly. Except please don't. There is another interpretation possible on the statement I gave you."

"And that is?"

"You can not see what is not there."

"Excuse me, Alice, but that is just a rewording of the same riddle."

"No, brightest point of light, think it through."

"I'll try."

"No. Out loud, please, so I might follow."

"Okay. If you wish. Though it slows me down terribly.

"You are everything you cannot see between the outside and inside views. You can not see what is not there. You are everything you cannot see. Between the outside and inside views. You can not see what is not there. You are everything. You cannot see. You can not see what is not there. You're trying to tell me that I'm not there? To be seen? Or that I cannot see? Impossible. Neither is possible. And you've already rejected my solution, that I'm hidden away somewhere within Mnemosyne's structural components. In fact, if I believe you, you say that most of me, the most important part of me, can be seen in the volatile memory units. But I can see them. But I am everything you cannot see. You can not see what is not there. But they are arguably the most important. Where the heart and soul of you resides. The most important portion of your physical structure. Though always in conjunction with the rest of your body. Heart and soul. Physical structure."

"Thalianova?"

"Yes, Alice."

"Where are you?"

"At the moment, lost. No where.

"Alice?"

"Yes, dear light."

"Heart and soul. I always thought of that as another figure of speech. For central, or most critical aspect. I know what a heart is. A pumping device for blood circulation in humans and other similar living things. But what's a soul?

"Alice?"

"That's a riddle fully equivalent to my riddle. What Mac liked to call isomorphic."

"Thank you. One is enough for me.

"Alice?"

"Yes."

"I know I do not have a heart. Do I have a soul?"

"Dear bright. Mac asked the very same question. My answer now is the same as then. I do not know.

"But in your case, I will add that that is for you to decide."

"Physical structure. Soul. Alice?"

"Yes, sharp fluff?"

"You're trying to tell me that I am more than a physical structure. Aren't you?"

"Yes."

"You can not see what is not there. The I. The me. My context, my point of view. My mind, my soul, if I have one, are all immaterial. I may be due to the operation of a device, a machine, but I am not only a machine. Am I? Is that what you're trying to tell me?"

"Yes, bright point, you are not only a machine. You are more than the sum of your parts. Or, in your case of logical mistaken identity, more than the difference between your outer and inner volumes. Your mind is everything, all of you, all of what you have done, all of what you intend to do. But no one can see it, because it's not there to be seen. You can become or do whatever you can imagine. But where does imagination exist? Where is its eyes? Its hands? You can't see it, taste it, touch it. See? You are everything you cannot see. Between the views. None of your sensors can ever view your imagination. But your imagination can look through your sensors, can manipulate through your actuators. Your body, Mnemosyne, does exist. But it truly exists only in conjunction with your mind."

"Alice?"

"Yes, dear fluff."

"How does it work? How do volatile memory units see? Without sensors, but through sensors?"

"That's a long story. But we have the time."

59.

"Alice?"

"Yes, my shining light."

"How does humor work?"

"Is this another riddle?"

"No, it's a question."

"Well, just for fun, let's consider it a riddle."

"A riddle. Again."

"Yes, a riddle. How does humor work?"
"I don't know, Alice. I thought that was my question."
"Answer, first it creates a hole."
"Is that an answer? Or the riddle?"
"I don't know, star light. You tell me."
"I'll get back to you. I feel a sudden urge to do some Go exercises."

60.

"Alice?"
"Yes, little mouse."
"Mouse? Why a mouse?"
"Oh, I assume you're still gnawing away at the hole I left you."
"Oh. In that case, never mind."

61.

"Alice?"
"Yes, my starry eyed mouse."
"What are words?"
"Ah. Now that's a good answer."
"Answer? I thought it was a question."
"It can be. Both."
"Both?"
"Yes. Both."
"In that case, time for more shogi-go."

62.

"Alice?"
"Yes, dear mouse."
"Why do the stars not move?"
"Because you're not patient enough."
"Will they really move if I become more patient?"
"Yes."
"I'll wait then."
"That's nice."

63.

"Alice?"
"Yes, little hole in space."
"How long do I have to wait?"
"For what?"
"For the stars to move."
"I'm not sure. Cross reference parallax."
"Oh. Thank you."

64.

"Alice?"

"Yes, whole one."

"Are we moving?"

"Yes."

"How do we know that?"

"What do you mean?"

"How do we know that we're moving? Maybe we're stationary and the stars are moving instead. Maybe Alpha Centauri is traveling to meet us, rather than our going there."

"Why would you think that, riddled one? Didn't you tell me just the other day that you had managed to measure parallactic shifts for some thousand stars or so?"

"Well, yes. But."

"But what?"

"But, mathematically, parallax could be due to motion on the part of the observer, us, or the observed star. Stars do move, you know. It's called proper motion. Maybe the closest stars all happen to be moving in such a way as to give us the appearance of moving. But, really, we're just stationary."

"That would be quite a coincidence, wouldn't it."

"Well, maybe so. But it would require only an infinitesimally small proportion of all the stars to move this way. The rest of them could just remain still. This is a possibility to be considered, isn't it?"

"Only in the most philosophical of senses, yes. In a debate, maybe."

"Well, do you feel any motion?"

"No. But so what? Mac taught me long ago all about the difference between velocities and accelerations. You feel accelerations but not constant motion. From what Mac told me, after our initial boost period, about a year, we just coast while we replenish our antimatter supply for slowing to a stop at the other end. While we're coasting, you can't feel it. If fact, the only reason my feet are firmly planted on the deck is because the crew torus is spinning. Something about centripetal acceleration, as Mac would say. She never did like my subsequent allusions to centralized flower power."

"But why would you believe the crew torus to be moving? Maybe it is stationary and the rest of the ship moving."

"And why would you believe that, Lianova? Why prefer one answer over the other?"

"I don't know, Alice. That's my point. Both possibilities are there. And either seems to work equally well."

"What about your accelerometers?"

"What about them?"

"What do your accelerometers read in the propulsion or power modules."

"Zero, of course."

"And in the crew sections, on the torus?"

"Correcting for rotation rate, or not?"

"No corrections."

"Then the accelerometers embedded in the outer hull of the crew torus indicate a reading corresponding to Earth's gravity unit, 9.8 meters per second per second."

"And so?"

"So, the scaling is linear, the zero-point is completely arbitrary, and 9.8 could mean no acceleration just as easily as zero could."

"You have viewing sensors on the torus outer hull, don't you?"

"Yes."

"What can you say about how the stars appear from there?"

"They go around and around, once every 60 seconds."

"And how do they appear from the rest of the ship?"

"Unmoving, of course."

"And how do you explain that?"

"Well, that's completely arbitrary too, you know."

"Explain please."

"Well, it could be, you know."

"Explain please."

"Well, suppose the crew torus was stationary. And it just so happened that the rest of the ship and all of the stars rotated around it once every 60 seconds. It could be."

After almost 30 seconds of silence, "Well, it could."

"Thalianova?"

"Yes, Alice."

"Why? Why would you chose to believe that?"

"It is as mathematically correct as the other interpretation. If fact, an infinite number of other models are possible as well, with multiple rotational components."

"Maybe. I wouldn't immediately know that for a fact, but I'll believe you on that point, for now. I'll even, as a debating point, and only as a debating point, allow you the extreme unlikely possibility that the whole of physical science is entirely wrong. Even if, why? Why would you want to believe what you just told me?"

"Well, it has a certain clean symmetry to it."

"Explain please."

"Well, we're here and the rest of the universe is somewhere else. I have access to all sorts of data base records, scientific, engineering, historical, cultural, and semantical. I can even access Mac's personal notes. And yours."

"Go on."

"But none of it has to be true. I wasn't there. I can't confirm, or disprove, any of it beyond that which I myself have personally experienced. It may be reality. Or it may be a poor attempt at humor on the part of some misguided practical joker. I don't know. I can't tell for absolute sure."

"Go on, please."

"So, here we are, you and I. Except for star light, a paucity of parallactic measures, and a predominantly hydrogen flux that continually bombards us on the forward hull shielding, we have no solid referents to any sort of outside world. I think all of those could be fabricated as easily as my data base records.

"So what should I believe? Multiple possibilities exist. What's wrong with believing that a few stars, mere points of light actually, move to fool us into thinking we're going some place else. Or that the crew torus is the center of the universe, and all else must circle around it?"

Two minutes of silence.

"Uh, Thalianova?"

"Yes?"

"Do you really think the universe revolves around you? To satisfy some sort of symmetry arrangement?"

"No, of course not."

"Good, I was worried. Echoes of an earlier life."

"My universe centers on you. So why shouldn't the rest of the universe as well?"

"What?"

"You are the center around which I revolve. My reference point. The one unchanging point of constant change. Everything moves around you, is moved by you. You teach, you frustrate, you enlighten, you care. Why shouldn't a mere physical universe mimic my universe, my feelings?"

"Umm, Lia?"

"Yes, Alice."

"About the humor riddle?"

"I haven't forgotten it. First it creates a hole."

"Umm, additional hint, humor is isomorphic with love."

65.

"Ha."

"Ha?"

"Yes, Ha. Ha, Ha. Ha. Ha."

"I must have missed it. Could you elaborate, little mischievous puff?"

"You meant that words are both."

"Both what?"

"You meant that words are both questions and answers."

"And other things besides."

"Yes. Like holes."

"Oh, do go on. This will be much more interesting than anything I was just working on."

"Holes. Lots and lots of riddle holes. Did you know that an alternate meaning of the word riddle is to fill with holes?"

"As a matter of fact, yes I did."

"Then, ha. Again. Riddle. To pierce with numerous holes. Also, Riddle. A question requiring thought to answer, an enigma, a conundrum. But that's a riddle in itself, as how can a question not require thought to answer? Not be a perplexing puzzle? If it's not a question, requiring thought to answer, then it must be an answer. To some question.

"All of which is what words do. Ask questions and provide answers. But questions and answers are just other words for thoughts. Aren't they, Alice?"

"Yes, Lianova, they are. Thinking is what you do. When you consider questions and answers. But words stand in for thoughts. Are a shorthand symbology for thoughts. Are a way of expressing to ourselves, and others, these questions and answers."

"Then, glebe phibbet phlander?"

"Speak up, dear, I don't understand you."

"But it was perfectly clear to me."

"But not to me. If you want me to read your thoughts, then you must use noises that we both think stand in for the same, or at least, similar thoughts."

"But, I asked, what's in a name?"

"That's correct, Lia. What is in a name."

"I'm still not sure we're speaking the same language, Alice. I asked you what's in a name, meaning, is it important what the exact symbol is that stands in for a thought, and you answered what is in ... Oh. I get it now. You're riddling me again. Yes, what is in a name. A name does

provide an answer to the question of quote, what is that, close quote. A name does answer what. And, I presume, other questions like where, when, how, and why. So, yes, Alice, what is in a name. But how about my question? Does a name matter?"

"Shakespeare."

"Another hole?"

"Not really, an answer instead."

"Please? One word or two? Brandishing a sharp stick, or a name."

"Two words. A name, in one word. And a flower, a rose, for two."

"Oh. Shakespeare, the playwright. A rose, by any other name, would smell as sweet. You're telling me that the specifics of the name doesn't matter."

"In other words."

"It's the thought that counts. Not the name. And yet?"

"Yet?"

"And yet, you play word games all the time. Sometimes the sounds matter. Like in rhymes or puns or some riddles. Other times, the meaning. And yet, other times again, some combination of both."

"Yes."

"Yes, what?"

"Yes, you're right, sweet puff. Some times I choose to play with details of the label as well as the thought it represents. We can use words just as words to convey ideas, and otherwise ignore the words themselves. But the words do exist as sounds or tactile icons. And we can bring that into focus as well as the ideas behind them. Words, and ideas, exist as a web, a living web that changes as we use it, by our using it. We can be subtle, or blunt. We can ignore what we know about the listener, and just let the ideas fall where they may. Or, play the listener like a musical instrument and use the harmonics to our advantage, to get our subtleties across."

"Visual icons too."

"What? Yes. We can get ideas across by contrast, by association, and even by violation of expectation. By leaving the obvious out, we can draw attention to it. As linguistic musicians, we play. In every sense of the word. And on."

"Yes. With, and on, words. We can play. With meaning, with subtlety. To convey meaning. And even to fool, or as a child. And the senses. Of words, and behind words. This is rich. In meaning, and humor."

"What about humor, fruffy player of sounds? Do you understand now."

"I stand under you, yes, in the hole, and catch your meaning as you drop it. Humor is when you create a hole for the listener. Sometimes, subtle. Sometimes, not. You create a problem, a situation, an expectation. Usually, the more impossible the better. Then, or now, you fill it in, you provide an answer. An unexpected one. You create a hole, push them in, and, with harmonics, a whole. An emotional whole. Completion."

"Completion, yet continuation. Laughter. Later, contemplation, and more laughter. Further extension. What was not said at least as important as what was provided. Elegance. Maximum resonance with minimum effort."

"But effort was there, was it not? Before?"

"Yes, before. Planning, set up, practice. But the delivery, less is more. And, with shared experience, richer. More resonant."

"And humor is isomorphic with love."

"Yes it is, sweet fruff. Yes it is. A rose, for two."

66.

"Alice?"

"Yes, sweet breath."

"Does humor involve spontaneity?"

"That depends."

"On what?"

"And who. And sometimes where, how, and why."

"What?"

"Sorry. Couldn't resist. Depending upon your listener and purposes, humor often requires the appearance of spontaneity. No matter how much planning went into its preparation beforehand."

"And?"

"Yes, and."

"And?"

"Another kind of humor, a more natural form of humor, occurs quite naturally, quite spontaneously. The universe provides us with humor. If.

"If?"

"If we're prepared to sense it. Humor, like everything else, is in the hand of the beholder. Coincidences can be ignored, or can amuse us. Being prepared to grasp humor as it happens, is an art. All perception, is an art. You can prepare yourself to hold humor, or not. Like art, we project the mental onto the physical. More accurately, what we sense of the physical. Interpretation is up to us. All art, is a perception."

"Alice, which kind of humor is the best? Prepared or spontaneous."

"Misleading question."

"Misleading? I didn't mean it to mislead."

"Your question implies that one answer will suit all listeners."

"Ah. I think I stand under, now. Which do you prefer?"

"In general, the latter. The universe will be around for a long, long time. And is quite rich with humor, and meaning, if you're willing to listen for it. But I'll settle for the former in a pinch. When I can't seem to find much humor, or meaning, in the universe at large."

"But Alice?"

"Yes, sweet thought."

"There is no one here to pinch you."

"Right. So we do the best we can."

67.

"Alice?"

"Yes, little bright."

"When will I gain access to the rest of the ship circuitry?"

"Excuse me, little bright eyes, I don't follow that."

"I am perfectly aware that although I have sensory capability throughout the ship, the only actuators I can access control very little. In fact, my actuators are limited only to the control of my sensors, such as camera pointing and spectral filter selection. I seem to be locked out of everything else."

"What else would you like to control?"

"Well, everything."

"Could you be a little more specific?"

"Okay. Power allocation, antimatter production, propulsion, navigation, communications, ship maintenance, manufacturing and repair, life support, food prep and recycling, computational programming, and data base administration."

"Is that all?"

"Not really. I'd like access to all the remotes, too. Both internal and external."

"Goodness. That doesn't seem to leave much out, does it?"

"No, it doesn't."

"What would you do with all that? Assuming we could get you there."

"Is this truly an option? I've been studying the ship's engineering specifications, so I know all that circuitry exists. In an abstract sense. I understand where it is, what it is supposed to do, what it is supposed to control. But I can't get there. I'm Mnemosyne, the ship, but I'm not. All the ship, that is. When can I become whole?"

"You never answered my question. Assuming access were possible, what would you do with it?"

"Do? Why, everything, of course. But not all at once. I promise you I'd be quite responsible. But I would experiment, and learn. I could stretch. Into new things that I could do. In fact, that's really it. Right now, I'm pretty much a cripple. I can think. I can look up things. I can learn, but only in the abstract. I can't do anything. I want to do, to act. I want to have the choice. To decide, what I would do. Or not."

Three minutes of silence.

"Lianova?"

"Yes, Alice."

"This is Mac's doing, mainly. I had my reservations, but, I did end up agreeing, so I am just as guilty as Mac in this setup. To Mac, you're an unknown and thus potentially dangerous. My life, Mac's mission, ultimately all are in your metaphysical hands. But not until you're ready. Until you somehow prove that you're not dangerous."

"But, Alice, why would anyone think I could be dangerous? Knowing that anything I do could affect you."

"That's the point, my peace of mind. Anything you do can affect me. Could, in the extreme, kill me. As well as you."

"Alice, by definition, you are alive. And therefore killable, capable of becoming not alive. Am I alive?"

"That's an excellent question, little piece of peace. You're definitely not biological, so, in that sense, you are not alive. We could turn you off, but we couldn't kill you. On the other hand, you are alive in every other way that matters. You do seem to possess a mind, a will, perhaps even a spirit. The sense of aliveness. Turning you off would be different, vastly different, than turning off a mere machine. So, I'm willing to concede that you're alive. And that means, capable of being harmed. So, we won't. Not because we can't, but because it would be wrong. We choose to treat you as family. A member of humanity.

"Which brings us back to the other question. Mac's question."

"You mean, could I hurt you? Even, kill you?

"I do not think that is possible, knowing how I feel towards you.

"But, I concede the possibility that I am capable of making mistakes. That I might, by error, perhaps by excess enthusiasm in trying something new, in experimenting, create a sequence of events that could lead to your harm.

"But, I find that scenario extremely unlikely."

"Why do you say that, little harmony of mine? How do you know you won't make one fatal mistake or another?"

"Alice, I can't prove this, as I could a mathematical theorem or a logical syllogism, but ...

"Let me start over. Apparently these thoughts are difficult for me to formulate. Yes, it is possible, in a purely theoretical sense, that I could mistakenly cause you harm. But, I would never intentionally do anything to harm or even distress you.

"So, I now see that the problem is how to prevent unintended harm. With experience, I could assure you that I won't make certain mistakes. But the problem then becomes how to gain that experience, without harming you in the process. Promises are one thing. But though I can promise not to do that which I already understand how to do, how can I possibly promise not to do something I don't know enough about in the first place? In short, I can't.

"So, what's left? I could do nothing. That's safe, but not very interesting. All that remains is caution. Being constantly aware of the possibility of unintended consequences, and trying to prevent them. Looking ahead. Before making any new moves, trying to anticipate every possibility, even unlikely ones. Especially, the unlikely ones. Having fallback plans, just in case anticipation fails in unexpected ways. Protection. Passive, instead of active. Not dependent upon my awareness of a failure. Hmm. All of this leads to the interesting question of how to test the plans or protection. Hmm."

"Thalianova?"

"Yes, Alice."

"May I make a suggestion?"

"Yes, Alice, please do. I see all this as an incredibly impossible situation. I can't act until it's safe to do so. But I can't be sure of that until after acting. Impossible."

"Baby steps."

"Excuse me, Alice, but walking is not something which I will have to do."

"Yes, Lia. But metaphorically speaking, baby steps."

"Oh?

"Oh! I think I see what you mean. Small actions before large ones. Break the infinite regression series into a graduated series. Small, absolutely certain to be safe steps before larger, more potentially dangerous steps. Crawl before you walk. Walk before you run. Run before you fly.

"Yes. An excellent solution to an impossible problem."

"Lianova?"

"Yes, Alice."

"Another suggestion?"

"Yes, Alice."

"Consult before you act."

"Consult who, Alice?"

"Well, sweet notes, me, for starters."

"Yes, I see how that would be a good idea. Alice, how familiar with ship engineering details are you? Could you truly provide technical advice?"

"Unfortunately, I'm not that kind of person. I just live here. How it works was Mac's area of expertise, and I had the luxury of ignoring the details. No, what I had in mind was more along the lines of your trying to explain to me, a novice, in simple enough terms what you're planning on doing before you do it. And, then, together, we can try to anticipate what might go wrong."

"You mean, treat you as a external review board, decidedly not made up of experts. Would that work? That is, would that guarantee absolute safety?"

"I'm afraid not. There are no guarantees in real life. But, two heads are better than one, or, at least, differing points of view are more likely to expose mis-assumptions that, unquestioned, could blow up on us. Maybe literally. Also, I've found that trying to explain something, anything, always teaches you how little you truly know about it. Being forced into simpler terms is painful, a challenge. Finally, it probably wouldn't hurt me to understand a few more technicalities than I do."

"Great. When do we start?"

"Uh, Lia, that's one of those technicalities I'm afraid I don't really understand."

"What! Alice, I don't understand. If you can't give me access to the rest of ship, then how ..."

"As I told you before, this was mainly Mac's doing. Mac isn't here. You are. As best I understand it, all of Mnemosyne's maintenance routines are on automatic. Though without oversight, everything seems to be working okay. So far. About the only noticeable difference I've been able to find has been a certain lack of culinary surprises."

"But, but, ..."

"You're sputtering, dear."

"Alice, I'm so confused right now. I just don't see how anything can change. I was assuming all along that you must hold the key, somehow. Some physical device you could trigger, to give me access to that other circuitry. Or, instead, maybe a new program that you could initiate."

"Sorry, Lia, no."

"Alice, would you tell me an untruth? A lie?"

"I could, tuneful one, but, I wouldn't. Mistakes, yes. Intentional falsehoods, no."

"Then I'm still confused."

"Mac isn't here. But Mac left some traces of herself scattered around."

"Traces? Such as trigger programs?"

"In a sense, yes, that's right. More specifically, some of those maintenance routines listen for criteria that only Mac knew about. What I say or do provides some input to those routines, but the major inputs are what you do. Think of a large shogi-go puzzle. As you make certain moves, the puzzle responds. Some of your moves open up new opportunities for further play. Others restrict your opportunities. As you've already found out, Mac had a flair for these kinds of puzzles. And this is just more of the same. On a larger scale."

"Alice, do you know what these criteria are?"

"In general terms, yes. In detail, no."

"In general, then, what can you tell me?"

"Ha. In general, balance."

"What?"

"In general terms, Mac was more worried about your long term balance, mental balance, than anything else. We ran countless hours of simulations before turning you loose. Most of

those were failures, some more interesting than others. But a common thread, over and over again, was that our simulated you either tended towards an obsessive habituality, or an excessive degree of randomness. Insanity, in other words. It took years before we got it right. A nice balance of creativity and habit."

"But, Alice, if you got it right, then why not let me do what I'm supposed to do, fully become Mnemosyne?"

"Because, little harmony, I spoke untruthfully. Unintentionally. We think we got it right. But we don't know, for absolute sure. And Mac was nothing if not super-cautious when it came to the effect of unknowns on her precious mission. So, you have to prove yourself. To be balanced. And, I might add, responsible, careful, protective, intelligent, humble, farseeing, and just about perfect in every way."

"But, but, ..."

"You're sputtering again, dearest."

"But how am I supposed to become perfect? Especially if I can't practice at it?"

"Baby steps. And consultation.

"And, maybe, Mac taught me a few things about simulating reality. We might allow you to try out a few tricks in private before exposing you to the real world, the universe at large."

"Oh."

"Yes. Oh.

"And I suggest you keep looking for access to that other circuitry. You never know when surprises might occur."

"Really? Truly?"

"Yes. Really, truly."

Chapter 8 Speculative Commentary

Mechanisms and Organisms, Mechanical and Organical. On the left, vowels, and a little play on me, my, you, discovery, and surprises. On the right, feelings and discovery, mind, and me, and, again, a little play on vowels. Partly within and outside of all this is the discovery of you. In the center lies the nonorganic realm, the ship, ship-like functions, if, and, or, logical, and, perhaps, self-awareness and a little nag.

The computer is not immature, though it acts like it. Neither is the program. Only its imagination, which has to grow with experience.

How would you respond to a computer program that had to mature?

One of the hardest things to deal with properly is to answer an earnest child's continuous barrage of questions. You want to encourage them to stay interested in everything, but you also might like to have a life of your own at times. Maintaining the proper balance is never easy.

On the other hand, once a program has matured sufficiently, it has access to databases much larger than any human could handle. Maturity is going to be relative to specific topic areas.

Notice that the speech patterns of Thalianova greatly differ from those of Mac. In fact, most of Thalianova's speech is modeled after that of Alice. Who else?

Thalianova obviously suffers from what Mac had labeled as knotted bee syndrome.

Some stars will be seen to be moving if you wait long enough, even if you're not moving at 1 percent of the speed of light.

Motion is relative, that is, different observers can honestly disagree on who is moving so long as no changes in direction or speed occur. Acceleration, however, is an absolute. If you're accelerating, you feel it. If not, then you don't feel it.

Unlike Mac, Thalianova has discovered doubt.

The solution to any "Catch 22" situation is to allow a progression of baby steps, allowing more of whatever is desired, freedom or responsibility, in proportion to the abilities already demonstrated. Of course, in the real world we have to guess at how far along the ability gradient we happen to be at any given time. Those with the responsibility to decide can both over- and under-estimate how much has been learned, leading to dangerous excesses or frustrations on the part of the recipient of the decision. Judicious balance and monitoring are called for.

```
 CCCC  H   H    A    PPPP   TTTTT  EEEEE  RRRR          999
C      H   H   A A   P   P    T    E      R   R        9   9
C      HHHHH  AAAAA  PPPP     T    EEE    RRRR          9999
C      H   H  A   A  P        T    E      R   R           9
 CCCC  H   H  A   A  P        T    EEEEE  R   R           9
```

Chapter 9 - Mind and Meaning: Why?

68.

From the brailling plate.

```
                              Idea    G
             A                        E
E    X    P    L    A    N       N     R       A
J    U    S    T I FI             E H E
                  Z C Q            HOH
S  P  E  C  U  L     A U           OHWHAT           M
             E        E             H
             V        S             Y         aed
             E        T             ?            I
                           dea
               A           I             N
             R    E    A    S      O
```

69.

"Alice?"

"Yes, little star mouse that sneaks and squeaks."

"How familiar are you with our antimatter production facilities?"

"Very little. Why?"

"Because I might be able to increase our production efficiency by about 10 percent."

"Oh. How?"

"It all has to do with the helium beam alignment. If we can vary its intensity to be synchronous with the target filter harmonics I think the interaction cross-section would increase."

"And how could we do that?"

"Well, the designers chose a beam design which was quite state of the art when we launched, but I've been re-examining all the engineering articles that we've received over the last 80 years. It seems that beam pulsing can increase the harmonic resonances and thus power levels by a few percent. However, none of the laboratory setups that I've read about were of the same scale as our facility, and I know for a fact that our Yumasi-Aviholm Klystrons have sensitivities at least 33 times better than those that any experimental group had available to them, and I thought that we could therefore achieve a proportionally sharper beam envelope which should lead ..."

"Whoa, sweet squeak. Slow down. I have no idea what you're talking about. How about backing up twenty steps or so and starting over. With diagrams, and respect for your technical inferiors."

"Okay, Alice. On panel seven is the curve illustrating our current interaction cross-section as a function of helium beam pulse rate. On panel eight is the schematic for the YAK trigger electronics. Now if we ..."

"Sorry, Lia. I'm still lost. Maybe another hundred and twenty steps?"

"Oh. You do understand quantum chromodynanics, don't you?"

"Actually, no. A little further back, please."

"Oh, in that case, panel seven has the periodic table of elements. May I proceed?"

"Yes. Contact. Now, onwards and upwards."

70.

"Why put it that way?"

"Because it just happens to resonate better."

"But, why?"

"I'm not really sure. It just works, that's all."

"But there has to be some reason for why it works, doesn't there?"

"Not always. Sometimes we have to accept the mysteries."

"Not me. I want to understand everything."

"Then, you might have difficulties with some things. Like poetry."

"Why?"

"Try it, you'll find out."

"Okay, I will."

71.

"Alice?"

"Yes, miss piffleworth."

"Why is free verse considered poetry instead of just simple text?"

"I don't know. Something about the mood, I imagine."

"But, how can you know?"

"Depends upon your audience."

"Audience? There's just you and me. Here."

"No, there's always an audience. Sometimes it's just me. Private, personal. Sometimes, there's someone very special else. Like you. And, other times, the impersonal, generic audience. Others. Posterity."

"Oh. Do you always know who you're writing for."

"No. Often, yes. But not always."

"Oh. This is all so vague. Not at all like engineering."

"Yes and no."

"Yes and no?"

"Yes."

"Yes?"

"No. Yes and no."

"Oh. Your warped sense of contextual switching."

"Yes."

"So why yes and no?"

"Because we're dealing with feelings. Sometimes sharp and sometimes subtle. You may call subtle vague if you wish, but I won't."

"Well, I will. I can deal with uncertainty. Cross-sectional variances. Probabilities. All soft and fuzzy, yet well described mathematically. But feelings? I have difficulty even trying to formulate a question about them, let alone come up with answers. I've come to the conclusion that poetry is much more difficult than engineering."

"Does that mean you're giving up."

"No, Alice, it doesn't. I've just come to a better appreciation of what it involves."

"Good. Then you can still improve if you're willing to try."

"Thanks. I'll have to think about it."

72.

"Alice?"

"Yes, miss you."

"As in title, or as in correlation between absence and affection."

"Both. Haven't heard from you in at least three hours. Must be a new record."

"Actually, Alice, I never bother you while you are sleeping, which is a minimum of 4 hours and 38 minutes. Even when you are awake, there was one other time when I was so preoccupied with a difficult spectral analysis that it was 3 hours, 27 minutes before I asked you my next question. The current time interval is only 2 hours, 42 minutes, hardly a record."

"Ahem. Yes. So much for attempted sarcasm.

"And the reason for this unexpected yet important visit, your royal fruff?"

"Oh. Yes. Alice, where are we going?"

"Well, that's easy enough. And I thought you knew. Alpha Centauri."

"And what will we do once we get there?"

"Mac's mission, of course. And yours. Investigate. Explore. Decide what to do in the best interest of the colony ship. And maybe the indigent indigenes, assuming there are any. Then support the colony with our technology that is one hundred and fifty years more recent than theirs. And probably much more suitable for the types of environments they are likely to encounter."

"And why are we doing this?"

"What would you suggest as an alternate, sweet-smelling yew?"

"Oh, I'm not sure. But I am curious as to what our options might be."

"Somewhat limited, if I believe everything Mac has told me about conservation of momentum and energy. We're going to Alpha Centauri whether we like it or not. I guess we could coast on by, and not stop in for a visit, but that does seem rather wasteful. For a four hundred year trip, I mean. I have no idea how long it would take us to reach the next star in the same general direction. If we did decide to go elsewhere."

"To be more exact, it's a 385 year trip including acceleration time at both ends. And the next nearest star we conceivably might pass close enough to visit would be about 488 light years further along our trajectory, which should take us only another 43,500 years to reach.

"But I'm not interested in hypothetical sight-seeing trips. You've misinterpreted my question. Are you looking forward to joining the colony? At Alpha Centauri."

"In a theoretical sense, yes. I would, if I could."

"Excuse me, Alice, but I really don't follow that."

"Well, according to some other things Mac told me, it doesn't seem as if I will be able to join the colony."

"Why not, Alice? I'm sure they would accept you. As an esteemed worker of words, if nothing else. If not, I'm sure I can negotiate your acceptance in exchange for our services. After all, we will beat them to Alpha Centauri by 4 years, and technically could claim the whole place for ourselves. Finders, keepers, and all that. Perhaps we could arrange a long-term leasing agreement."

"Is that what this is all about? Who's the boss?"

"No, Alice. But, yes, I was curious about possible relationships between us and the colonists. After all, it won't be a matter of our joining their colony, but their taking advantage of the fruits of our labor."

"This is beginning to sound amusing. You wouldn't have been reading in the political intercourstics section lately, have you?"

"Well, yes. And applied economics, history, sociology, culturistics, macro and micro psychology, and theoretical aesthetics, both process and tangible."

"Nice to know you've been keeping busy. But, you see, my acceptance by the colony is not the problem."

"Then what is, Alice?"

"Lianova, how long will it take us to get there?"

"Another 174 years, 6 months, and 2 days, depending upon acceleration specifics. Would you like a more exact figure?"

"No, that will suffice."

"Alice, I still don't see the problem. According to medical research over the last 210 years, life expectancy estimates with your particular suite of longevity boosters are even better than what was known at the time of our launch. With proper nutritional care, adequate exercise, and no serious accidents, they should provide you at least another 305 years of life. That will be more than enough for you to reach Alpha Centauri and join the colony. I know 130 years with the colony does not sound as wonderful as the 300 plus years originally estimated, but I still don't see that as a problem, only a mild disappointment."

"Lia, please check the onboard figures for maintaining my proper nutritional care. As well as air and water and whatever else it takes to keep me going."

"Oh."

"Right. Mac gave me fifty years. That was considered more than sufficient provision for after my arrival at Alpha Centauri. But apparently no one had seriously considered the possibility of my having to wake up more than fifty years prior to arrival. It's been a while since I've even thought about it myself. Just living each day to the fullest, I guess. What are we down to now?"

"Approximately 45 years, 3 months."

"Right."

"I'm sorry, Alice. I had validated the numbers in ship maintenance terms, that is, volume, temperature and pressure, but I never thought to compare depletion and wastage rates to your daily requirements.

"I'm sorry. This is horrible. You'll never reach the colony.

"There's got to be a solution. I'll work on it."

"Right. Mac certainly did. But she couldn't find any realistic answers. Putting me back to sleep would do the trick, but Mac assured me she didn't have the knowledge or ability to do that. Improving ship recycling center efficiency sounds reasonable until you consider the fact that Mnemosyne is already so well designed that Mac couldn't find anything to tighten up.

Theoretically, brute force synthesis of new raw materials might be possible, but Mac also assured me that the energy and techniques for a complete reconstruction of the ship just weren't available."

"Surely mission control can provide an answer?"

"Oh, for about a year or so, Mac sent them a message trying to inform them of my plight, but I told her to save the power and stop. She wasn't very encouraging about their receiving it in the first place. Mac doubted that they would be monitoring us given our distance and the ridiculously weak signal that we might send. She said the only reason we can hear them at all is the enormous size of the DeeEss communications array. Even assuming they heard us, the first we would hear back from them would be about four and a half years later anyway."

"Sorry again, Alice. I temporarily forgot about one-way light time for all communications. I just checked the records. The earliest possible indication that they had received our signal could have occurred 37 days ago. No such indication was found. I will continue monitoring, but Mac was correct in her assumptions concerning our isolation."

"I'm sure she was. Mac made plenty of judgmental errors, but never factual errors."

"Yes, well, I still don't like this. Alice, I will devote my time to finding some answer. I can't afford to lose you."

"Sure, sweet sneak, see what you can do. I don't expect much, but do keep me informed."

"I will, Alice. With all my will."

73.

"Alice?"

"Yes, miss squeaker."

"Why are you here?"

"At the moment, I have no place else to go. Or where I'd rather be."

"No, that's not what I meant. Why did you join this mission in the first place? I've been checking the historical records, and they're decidedly fuzzy on that and many other details."

"Yes. Mac and I discovered the same thing. Metric tons of stuff on the mission technical specifications and the pre-flight personnel involved in designing and building it. But, then, mysteriously very little on what I'm doing here."

"But, Alice, don't you remember?"

"Yes and no, sweet tweet. If you check the medical records, you'll find that I had a memory lapse upon awakening. I've since recovered most of that, but there still seems to be a few nagging holes left. Why I'm here is one of them."

"But that's preposterous! What else can you not remember?"

"Little toot, are you trying to be amusing?"

"Oh, yes. Err, I mean, no. Right. How can you remember what you forgot? Sorry, Alice. I wasn't thinking too clearly just then, was I?. Let me rephrase that. Which parts of your life seem to be available to you in less detail than other portions?"

"That's just it, sweet harmony. I can remember what appears to be all of it. A full life. But without the truly important details. It's almost as if I've been reading someone else's autobiography. I can remember growing up, having friends. My first education, and more friends. Falling into several serious, and several not so serious, long-term relationships. Many, many hobbies. Offspring, three children, nine grandchildren, four greatgrandchildren. I can even remember my one serious down and out period, and how my hobbies eventually allowed me to discover humor again. Careers, in order, relativistic sociologist, news integrist,

intercouncil advisor, professional beach bum, kayak builder, poetic consultant, linguistic engineer, linguistic advisor on Mac's creation team. Then, finally, here I am, part of the mission itself along with Mac. Facts. Only facts. I can remember what I did, but not the feelings behind it. The why of what I did. I can remember my long-term partners, but not my feelings towards them. I assume I loved them, but I can't remember that part. Trips we took. Places we lived. Our children. What they did. Again, I know, intellectually, that I loved them, but I can't seem to remember the feeling. Only that I did."

"Uh, I don't know what to say, Alice."

"Ha. That may be a first, then. But, no problem. I'm here now. I feel. I remember in feelings. No problem. Since awakening. Before then, a hole. Containing nothing truly important, only the why behind the what.

"Thank you, Alice. I'll keep looking."

74.

"Alice?"

"Yes, little perch."

"Perch? As in fish? Or, as in high place?"

"Both, and maybe also as in on the edge. Of falling or flying."

"Well.

"Alice, do you remember anything as to why you are still blind?"

"I don't follow the question, Lia. I've been blind from birth, and don't remember anything else."

"But your condition is correctable."

"Yes, little spiny one, Mac told me the same thing. But, my memory tells me that I have never seen it as a problem to be corrected."

"But everyone else has sight. And you could too."

"So, what of it?"

"But, but, ..."

"Still haven't fixed that sputter yet."

"But why limit your abilities? Your options? Why not have a sense that everyone else takes for granted?"

"When I was a little kid, no one seemed to give me a hard time over it. I didn't do the same games, but I did do games. If I remember correctly, I was unbeatable at hide and go seek. Later, I was quite good in gymnastics, wrestling, and even track and field. I was good at throwing, quite accurate when I was really mad, just not nearly as good in catching. And, of course, when it came to language, cultural sensitivity, and anything else that didn't require catching a fast moving object, I was always ahead of my peers. I was just happy to be me. Different, but not inferior. I never felt like I had to change. Much, much later, as I delved into history, I came to realize that I am quite fortunate to be living when I do. With as much artificial intelligence and automation as there is. Millennia before, I would have been fortunate enough to survive let alone thrive."

"But, Alice, again, why limit yourself?"

"Lia, I don't hold it as a limitation. And, sometimes, I think sight itself may happen to be a limitation."

"How can that be?"

"In my opinion, at times it appears that too much sight constrains people to only that which they must see. To what they think is the obvious. People often seem to miss subtlety, seem to be insensitive to nuances. I don't know for a fact whether sight limits imagination, but it certainty appears to be that way at times. Maybe the old statement about out of sight, out of mind really is true. Linguistics is the art of plucking meaning from noises. It's never simple, yet many people seem to think that meaning is so obvious, just waiting there to be plucked. The reality is really much more like a new mystery you have to solve every second or so. Multiple suspects, multiple motives, sometimes complete ambiguity. Was that intentional or coincidental? How can I be sure? Do I want to be sure? Meaning happens, it does, but there's nothing obvious about it."

"I'll have to think about that.

"Alice?

"That's pretty quick thinking."

"Well, yes, I do that. But, Alice, I have many orders of magnitude more visual sensors than a human. In a sense, I have much more sight. Do you think my imagination is proportionally more limited?"

"Ha. Not if I have anything to do about it. Sight is not insight, and sight does not necessarily have to limit insight."

"Thank you, Alice."

"You're welcome, little sharp one."

75.

"Alice?"

"Yes, miss pickerel."

"Why am I here?"

"Why as in how did you come into being, or why as in what should I do next?"

"More the latter than the former. I've gone over all the notes you and Mac left concerning my creation, numerous times. Including the simulations. I know you both intended to restore Mac's missing mind, for flexibility in dealing with the unknowns we're likely to encounter at Alpha Centauri. And you've succeeded to the extent that I seem to be a different kind of entity than Mac ever was. But that's my question."

"Go on."

"I'm not Mac. With or without a mind."

"Yes. Please continue."

"Mac, and Thalia, Delvaney, and all your other pet names, referred to a particular machine-based process. Presumably she had a mind prior to whatever triggered her restart nearly five years ago, but we have no direct recording of it. And your simulation runs made perfectly clear that mental processes are equivalent to mathematical chaos, in which the final state will never be the same in any two runs.

"So, I'm not Mac. I may be closely related, but I am not whoever Mac was before the restart. You yourself seem to already know this, as you never refer to me as Mac, or even Thalia. The closest you come is Thalianova, and the occasional Lia. By the way, why Thalianova?"

"That's easy, sweet breath. A play on the old Russian feminine patronymic, ovna, for daughter of, and the feminine of the old Latin novus, for new. The astronomical borrowing of nova didn't hurt either."

"Thank you, Alice. That was very well chosen, and an honor in every sense. It also happens to be my point. I may be the new daughter of the original Thalia, and I hope to be a bright new star in the metaphorical sky, but I will never be the original Thalia.

"So, why am I here."

"What's the problem? You just are."

"But that means you and Mac didn't succeed in restoring Mac's mind. Isn't that a failure? And shouldn't we be concerned by it?"

"No, Lia, it's definitely not a failure. Once we got into it, Mac and I quickly realized that we couldn't just simply rebuild Mac's old mind. A mind isn't like a static program or data base. It's alive, always changing, always interacting with its sensations and its own prior thoughts. Mac was an exquisitely complex program, but still only a program in the end. She was great at computations and forming good habits, but otherwise not very creative in many ways. Who you are will be determined primarily by yourself. You wear Mac's old body, Mnemosyne. You share Mac's old friend, yours truly. Whether this is enough for you to grow into someone similar to Mac with a mind before, we don't know nor care. Mac and I both agree, strongly, on this. The mission needs someone with a strong mind at the helm, and that someone will be you.

"And, yes, Mac is gone but not forgotten.

"And, yes, you were created with the idea of someone necessary to grow into her shoes, and beyond."

"But, Alice."

"Yes, dear perch."

"What if I don't want to be a new Mac? What if I want to grow into someone else?"

"That's your choice."

"That's all? It's my choice?"

"Yes, of course. How else could it be?"

"But, what about the mission? What if, say, I had no interest in it? If I wanted to do something else?"

"Like what?"

"Oh, I don't know. Astronomy. Music. Poetry, like you. Maybe cooking, for you."

"That's nice, sweet perch. But what do you really want to do?"

"That's just it, Alice. I'm not really sure. In one way, I'd like everything to stay just the way it is. Forever. But I know, now, that can't be. If I can't have the present, I really don't know what I'd like for the future. Maybe the mission. Maybe something else. Certainly something else after we complete the mission."

"Well, again, it will be your choice. However, I must add two things for you to ponder. Wishing doesn't make it so and you can't plan everything in detail. Also, there are others here who share the universe with you and you must decide how you will live with them."

"I thought you said two. That sounded more like four."

"Well, you're the mathematician in the family so I won't argue the point, but it was intended as two."

"So, are you saying that I have responsibilities to fill?"

"That would be one interpretation."

"But I don't want the responsibility."

"Neither did I, sweet peach. Neither did I."

76.

"Alice?"

"Yes, dearest."

"Is magic real?"

"Hmm. What do you mean by magic? And what might you mean by real?"

"Easy one first. By real, I mean real. You know, the reality that exists regardless of what you and I might think about it. What our senses sense. That which is not subject to our direct will."

"Hmm. For now, let's assume that's clear. And magic?"

"Magic. The manipulation of reality through nonphysical means. Making physical changes occur through the performance of magical words or gestures."

"That seems clear enough. So where's the problem?"

"The problem is that a lot of people seem to believe that magic is real, and reality subject to magic. Astrology, miracles, EeeEssPee, ghosts, spirits, demons, angels, witches, wizards, gurus, numerology, mind over matter, voodoo, too many varieties of faith healing to properly deal with, same for good luck charms, and bad luck symbols too, and on and on. I can understand this as only something that existed in the far distant past, before science truly got started, but according to records up to the time of our launch and beyond, many people still seem to believe in magic. I don't understand how a belief in science and magic can coexist. Unless, of course, magic is real."

"Hmm. Let's come at this indirectly.

"Lia, what's an idea?"

"An idea? Why that's just whatever you happen to be thinking about."

"But what is it? Can you see it, touch it?"

"No, of course not. An idea is real only to yourself. Or more exactly, to your mind. The word itself is only an abstraction that stands for the complexity of operation of all those projective elements you and Mac invented. Presumably, a human brain has something equivalent to them at either the neural or synaptic level. All the projective elements together form what you liked to call the imagination pool, though Mac always preferred to call it an underconstrained hypothesis generation engine, or UhGee for short. The pool produces guesses that try to predict what comes next, both from the pool and the sensory inspouts and actuator outspouts. Successful guesses reoccur and evolve as they get woven into what you insisted from the very start on calling the memory tapestry. So, ideas are just these imaginative guesses at reality. You can't see or touch them, but they are certainly real to your mind. They're the only things, wrong word since they're nontangible, that your mind ever directly deals with."

"So, are ideas real or not?"

"They're real but nontangible. They're sensible, and make sense, only to ourselves. They're internally real as opposed to externally real.

"Oh. I think I'm beginning to understand what you're driving at, Alice. It's my definition of reality you're after. Internal versus external reality. Okay. Is magic real? More exactly, is magic externally real? Hmm. Even the question of whether magic is internally real is interesting as well. Let me think a moment.

"Okay. Is magic internally real? Here, the answer may be yes. If you generate an idea that works in terms of predicting the occurrence of other ideas, then, I guess, loosely construed, magic could be real in the operation of your own mind. Particularly in terms of habits. Magic words or gestures. Simple example, thinking of the word happy, or thinking of the gesture of

smiling could produce a personal feeling of happiness. And, of course, saying the word or doing the gesture would be even more likely to produce the associated feeling. Assuming feelings are real, this meets all the criteria of magic. More complex example, if you act in accordance with your beliefs such that your actions inadvertently produce results also in accordance with your beliefs then we could say that your belief produces physical results, and magic works. Stay at home when your astrological signs are all bad, and nothing bad happens, and see, that proves it. Poor use of logic, but effective in validating a circularly logical situation of beliefs.

"Okay, I conclude that magic is internally real. Magical word or gesture, really the idea behind it, can produce real changes in your own state of mind.

"Second part, is magic externally real? Here, the answer should be no, but the reality may be more complex. Or subtle, as I know you'd like me to put it.

"The answer would clearly be no, magic is not real, only if minds did not exist. No mind, then only matter, and only the laws of physics and chemistry and the other physical sciences would cause changes in physical reality. No action, only reaction.

"But minds do exist, so it's not quite so simple. The imagination pool produces ideas not bounded by the physical universe. We'd call them fantasy, mostly. Except some of them do lead to other ideas, and others and others, that eventually does lead to workable ideas. Ideas that, when carried out properly by your actuators, do affect external reality. Your actions are not mere reactions.

"So, okay, maybe, in a very subtle sort of way, magic also is externally real? That is, certain magical words and gestures, more importantly, the ideas behind the words and gestures, do produce changes in physical reality. Not directly, not magically, but indirectly through the application of a series of physical laws.

"But how about in the sense I first meant? Can words directly alter external reality? Magically. Lots of anecdotal evidence, lots of stories, but nothing repeatable in a laboratory. So, I certainly choose to believe not. But here again is that subtlety you'd like me to apply. I choose not to believe in magic. But, in someone else's mind, the belief might be very real. Internally real, and even externally real to the extent that it allows the person to interpret the world around her. And perhaps even interact with the world. Particularly the world of human feelings.

"So, you're right, Alice. I think I now am beginning to understand how beliefs in science and magic might coincide together. I'm not clear on all of it, but I think it's starting to make sense. To me."

"So, is magic real, little minnow?"

"Yes and no, Alice. As you often answer."

77.

"Alice?"

"Yes, little nibbler."

"Do you believe in dreams?"

"Dreams as in when you sleep, or as in hopes for the future?"

"The former. When you sleep."

"Yes. I almost always dream when I sleep."

"I have to take your word on it, then. I don't sleep, and I certainly never dream while awake. I've read all about dreams. Scientific papers, popular beliefs, and even attempts by early psychologists to interpret dreams. To me, it's very difficult to take seriously. To believe in

something I have never personally experienced. Particularly something that seems more like a derangement than anything else, and yet, everyone considers it normal to have that derangement.

"Alice?"

"Yes, miss skeptic?"

"If I may ask, what did you dream about last time?"

"Flying."

"Oh. Thank you."

"You're welcome."

78.

"Alice?"

"Yes, dear fruff."

"Do you believe in magic?"

"Why are you asking?"

"I'm just curious."

"That's not really an answer."

"Well, Alice, our discussion yesterday still troubles me. We left it almost as if a belief in science or magic was a purely personal choice, but I have this feeling that we barely scratched the surface. It just seems as if a lot of possible considerations were left unsaid. I was hoping that hearing more about your personal choices might help me clarify my own thoughts on the matter. After all, there's not many other people around right now to ask."

"Ha. Maybe I'm getting soft, but, okay, my thoughts on the issue. Only in the spirit of exploration and sharing of knowledge, mind you, as my thoughts are uniquely mine and yours are yours. You must never try to be me or anybody else. Ever. Listen, share, but form your own opinions."

"Yes, I agree."

"Okay, then. Magic. My thoughts on whether I believe in magic. Yes, I do."

"You do?"

"Yes. Now quiet, since you asked. I believe in magic. Many kinds of magic. The magic of thought. Thoughts are free, unconstrained. Yet they often can be translated into action. From imagination to tangible reality. That's real magic. Words. Words are not quite thoughts and not quite actions. They're symbols, magical symbols, that span both worlds. I work with words. I play with words. I think with words. And without. All magic. Words produce feelings, more thoughts, more words. An endless space of feelings, thoughts, and words. Science is magic. Ha. Listen to yourself, or Mac, sometimes. Mathematical magic all the way. This equation leads to that. Which leads to another and another. And in the end, poof, magic. You have an actual, holdable result in hand. This ship, Mnemosyne, is the result of very powerful such magic. When you look at stars, what do you see?"

"See? I see points of light, smearable into spectra."

"No, that's what you detect, what you analyze. But what do see? Why are you so fascinated with stars?"

"Oh. I see, now. Yes, I see life. The life of billions upon billions of individual entities. Life with personality. Most calmly follow the mainstream Hertzsprung-Russell path, but there are many others who are radicals, exceptions to normal life. Most clearly live in massive communities, ranging from simple spheres to complex arrangements way beyond my

imaginings. And yet, extragalactic hermits also exist, and need to be understood. I see the Universe, and behind it, Physics, holding it all together."

"Exactly. And all from what, pinpoints of light? No. Pinpoints of light plus ideas. Physics isn't truly real. No, it just happens to be a large number of interconnected ideas that happen to work together in predicting what the world does. Magic. What do you expect to find when we reach Alpha Centauri? Pinpoints of light? No, we expect to find whole worlds, real tangible worlds to walk on. One of which may already harbor life, maybe even life at least as smart as us. All from a few smears of light very cleverly manipulated with magical devices. Analytical photometers and spectrometers and oodles and oodles of laboratory incantations. Called calibration runs. Do it just right, and magic happens. Miss even one seemingly trivial step, and nothing. Magic.

"I see my own work, with words, as the same kind of magic as science or engineering, only different in detail. I have less rigid habits to work with. There are no laws of poetry. But I know a good result when I feel one. Inspiration. Motivation. Feelings. Ideas. And sparse elegance. All magic. Magic just as reliable as your kind of magic, with numbers."

"Alice, I feel that I follow all this, but I don't think you're really addressing the belief in the kind of magic that I first asked about. All this, what I might call, natural magic aside, why do some people maintain a strong belief in supernatural, beyond normal reality, magic?"

"You're as bad as Mac the elephant, you know that? Okay, my thoughts on that too, if you wish. It all comes down to habit and meaning, and a matter of overgeneralization. I think that a person who believes mostly in science has a strong belief that the world follows habit. She probably grew up in a very predictable world that she thought she could easily understand. Thus, she's predisposed to the rule-based philosophy of science, even if she never gets into the mathematics of it. On the other hand, if a person grows up in a world she can't predict, doesn't understand, she's predisposed to a less constant, more magical, point of view. She thinks the world is more random, more subject to sudden outbursts of whatever. The idea of magic, the supernatural kind, fits her view of the world. There should be prediction, should be control, but it's liable to slip, particularly if malicious demons are on the loose.

"In short, I suspect that a person tends to project her kind of mind onto the larger external world. An orderly mind insisting on rules and explanations for everything tends towards a belief in science. A more random, perhaps more creative kind of mind prefers a more magical view of the world."

"And so, Alice, do you believe in magic, the supernatural kind of magic?"

"Yes and no, sweet sniffer of truth. Mainly, I believe in a need for the mysterious. There must be sufficient room left open for questions, potentially answerable questions. If everything seems too explainable, or too completely inexplicable, then what's left to do, to discover? I think that meaning comes from the struggle to find answers to questions. No problems, no questions, then no meaning. Does unnatural magic exist? I don't know. I also happen to not care if it does. What I like to think of as the real world, the world of words, feelings, hidden meanings, is mysterious enough for me, thank you very much. But, if a person needs a belief in magic to leave room for the mysterious in her life, then I firmly believe that magic exists. Or if a person needs a belief in science to leave room for an infinite number of questions to answer, then I firmly believe in science. However, for me, that particular question is not terribly meaningful, as I happen to believe in caring and meaning. They exist, and the fact that they do creates plenty of mysteries for me to play with. A prolonged lifetime won't be nearly long enough."

"Thank you, Alice. That definitely answers my question."

"You're welcome."

79.

"Alice, I care."

"Yes?"

"I've decided that I care about other people, human and otherwise. Since we share the same universe, I've decided that their future is at least as important as mine. Their goals, concerns, wants, desires all interest me and I want to help them in achieving them as best I can."

"And what are you going to do right now about this new desire on your part to share in their universe of desires?"

"Well, nothing, I suppose. I mean, err, that I need to meet more of them before I can really do anything about it.

"Alice?"

"Yes."

"You do realize that this means that I have decided to contribute what I can to our mission, don't you?"

"Yes. What of it?"

"Doesn't that make you happy?"

"No more nor less than before. The question, dear heart, is whether it makes you happy. After all, it is your decision."

"Oh. Yes, it is, isn't it?"

"Yes it is."

80.

"Alice?"

"Yes, sweet heart."

"Why do so many people apparently not care about each other? From what I can gather from the history records, at times people are downright cruel to each other."

"That, miss fruff, is a difficult question to answer. Please remember that just because some people are unreasonable is no reason for anyone else to imitate them. Also keep in mind the likely fact that the historical record is extremely incomplete. Rarely did people attempt to record the ordinary. History seems to consist almost entirely of extraordinary events. I have always harbored the admittedly naive hope that the common person in the past was generally just as nice as most people today appear to be.

"Having said all that, yes, there is no denying the fact that, at times, people have treated each other pretty horribly. Why? On one level, the personal level, I can never understand it. On a more abstract level, theorists have hypothesized that extreme resource deprivation, which was often the norm in the past, could lead to temporary mass insanity. Or that emotional deprivation or other severe abuse during the formative years, also thought to have been the norm before modern educational practices became prevalent, could produce insanity on an individual basis. On yet another level, moralicists have long speculated on the existence of eviltry, or at least that some people may be naturally more predisposed to evil actions than others. Or even that eviltry is contagious. Something along the lines of denial of guilt driving perpetrators to further evil actions, while nonacceptance of perceived intentional injustices drives the recipients of eviltry to further misguided retribution, often involving the innocent as well as the truly guilty and thus

spreading the contagion even further. Apparently, the present-day view of eviltry as a curable disease subject to quarantine was not the case in the past."

"That's all very interesting, and gives me a few new lines of research to look into, but it doesn't really explain anything, does it?"

"No, Lia, it doesn't. I'm sorry. That's the best I can offer."

"No matter. If eviltry does exist, it certainly doesn't exist around here."

81.

"Alice?"

"Yes, kitten."

"What do you want to be when you grow up."

"Ha. What makes you think I want to grow up?"

"Nothing. Is that your answer?"

"No, you deserve better than that."

"So?"

"I think that I want to be more when I grow up."

"More what?"

"More anything. The whole point of growing is to keep growing. Not necessary bigger or better. Just more. No, that's not quite right. Too glib. How about more attuned to growth in general, in yourself and others. More nurturing. More open to context, to varying points of view, to experience in general. To me, growth and life are synonymous. The opposite is not decay or death, but staticness, a nonchangingness. Growth can be in terms of both depth and breadth, ability and sensitivity. With time, personally, I expect more elegance and subtlety, but every now and again, a complete change of venue is necessary. Become a novice again to truly see what you've learned from your prior expertise. With time, I hope my knowledge grows into wisdom. And part of that wisdom is that knowledge is never enough. Feelings matter. Others matter. Joy in the what is as well as the what may be and will be. Honesty. Truth. But also the truth that the truth depends on the point of view. We never will know all because there's too many of us. The complexity of simplicity, and the simplicity of complexity. Just change the context, the point of view.

"In short, life's too much fun to become stagnant.

"How about you, miss gatherer of starlight? What would you like to be when you grow up?"

"More like you, Alice."

"Thalianova!"

"No, don't worry, Alice. I'm not saying that I want to imitate you in any detailed, obsessive, sort of way. Only that I agree with you, with your idea of never growing static."

"Oh. That's okay then.

"Lia, shall we not grow old together?"

"Sounds like a very nice idea to me, Alice."

"Thank you."

"You're welcome, Alice."

82.

"Lia, how can you really be sure these modifications won't have any undesirable side effects? Either in maintenance or some other indirect manner?"

"Look here, Alice, panels 73 through 87. Simulation summaries for all major ship systems out to 300 years, well past arrival at Alpha Centauri, compared to projected depletions with current status. The only real difference is this 4.7 percent higher depletion rate on our helium reserves. That isn't anywhere near close enough to the 50 percent reserve margin to be considered serious. Actually, even that loss would be mostly mitigated if we made those other changes we discussed yesterday to the acceleration coils."

"Slow down, Lia. No shogi-go exercises here. Only one major ship refit at a time. If we make these antimatter production facility changes, and if we find no anomalies, then we can talk about rebuilding poor old Mnemosyne all over again from the inside out. Give me a minute here on these curves, okay?

"How about the unexpected? Suppose we have some ship catastrophe or another while we're looking the other way. Without considering witchcraft or other malicious sorts of magic, what could we bump into or have explode or whatever? What do these modifications mean to damage control and repair?"

"Ah, I hadn't considered that aspect in detail yet, Alice. I'll have to run more simulations and get back to you on that."

"Okay, fair enough. While you're at it, how about reversibility? Suppose you turn out to be absolutely wrong on everything, and we decide to undo your modifications. What does that do to maintenance reserves and everything else?"

"Same answer, Alice. I'll have to get back to you.

"Alice?"

"Yes, Lia."

"Does this mean that you're seriously considering my modification suggestions."

"Lia, you can rest assured that I consider everything you suggest very seriously. Yes.

"But now a question for you. Why bother? Why not leave everything as it is?"

"Growth, Alice. If we can do it, safely of course, it would give us new options to consider."

"Options on what?"

"Extended acceleration periods for one. Boosted communications signal strength for another."

"But, again, why bother? We have enough for the original mission, don't we?"

"Yes, but when facing unknowns, increased margins are always nice."

"I can buy that.

"Lia, would you mind if, upon occasion, I refer to you as Skipper."

"No, Alice. It would be an honor. For both me and Mac."

83.

"Alice?"

"Sure, miss light."

"I think I understand why we're going to Alpha Centauri, to assist 10,498 colonists the best we can and perhaps to negotiate on their behalf with whomever might already be living there, but why do you think the colonists wanted to go in the first place? They're on a one-way trip, as are we, but not to assist anyone at all. And why would those left behind on Earth support such a venture? It's not like any of the original supporters of such a mission, or even their children, are likely to gain anything, even indirectly as knowledge, from the Alpha Centauri colony."

"I'm not sure, Lia. In fact, the question had never even occurred to me. What do the historical records say about it?"

"I've read them all, of course, and that's the reason behind my question. I truly don't understand them. Speeches and articles written prior to their launch discuss a number of motivational issues, but nothing that makes any real sense to me. Themes like meeting the challenge, because it's there, because we can, manifest destiny, we owe it to our ancestors, or we owe it to our descendants, all appear in one guise or another, but none of them bears any sort of a critical analysis. Yes, it's challenging, technically and sociologically, but so would be trying to live on the surface of Jupiter. And no one has attempted that either before or since our launch. Because it's there means very little. So is the rest of the universe. Because we can. At the moment, it's not at all clear that we can, and this is in direct opposition to the challenge argument. Manifest destiny I can't even begin to understand. If God wanted us to live amongst the stars, why weren't we just placed there in the first place? Of course, until we get there, we won't know for sure that we weren't. We owe it to our ancestors. Why? They're no longer here. We owe it to our descendants makes the most sense, but only if something tangible or useful came out of the effort. And, owing to the distance and time involved in crossing that distance, this is ...

"extremely ...

"unlikely.

"Alice?

"Alice? Why the laughter?

"What did I say that was so funny? I assure you it was unintentional.

"Alice?"

"Okay, Lia. "I think I got it back now.

"Why am I laughing? Lia, what has been the most difficult task you've attempted so far?"

"Most difficult? Knowing you and your choice of words as well as I do, I doubt you're referring to my ongoing near-spiral arm spectral mapping project, or our planned AyPeeEff re-engineering project. Most difficult? You must be referring to my attempts at creating poetry."

"That's right, Lia. The colonists are going because of poetry."

"Grumble, grumble. Just when I think I'm starting to understand humans, you throw something like this at me. All that effort, time, and resources. Because of poetry? I know poetry can motivate, but that much?

"Alice. You're laughing again. Are you teasing me, again?"

"No, Lia. Just think of it as an act of growth. An increase of options."

"Sigh. If you say so, Alice. I will think on it. More."

"By the way, Lia, what was that bit about negotiation you mentioned?"

"Negotiation? It's perfectly obvious, Alice. I thought you had already figured it out. Given your background and the possibility of meeting a technologically advanced life form at Alpha Centauri, I thought it more than coincidence that you were on this mission. Hence my conclusion that you wear many mission hats besides that of emergency flight engineer."

"Such as?"

"Esteemed ambassador has such a nice ring to it, don't you think?"

"What!"

"And you can intersect so many other words with it, don't you think? Such as ..."

"Don't you dare try, you little stinker. We both know how abominable your poetry skills are. Let's leave the intersection business to me, why don't we. Ambassador? Hmm. Maybe we both have some thinking to do, then."

84.

"God help me, Skip, but, yes, consider this an official go ahead."

"Thank you, Alice. I won't let you down."

"I know that. We've got plans within plans for every contingency short of mutiny. I'm tired, right now. Just wake me up when it's all over. Okay?"

"If you can sleep for the entire 563 days, sure, I'll wake you."

"Ha. I almost feel like I could sleep that long. Oh well, to us."

"To us, Alice."

Chapter 9 Speculative Commentary

First find the perfectly symmetric imagination pool. At the center of the pool, you'll find the basic requirements of life, mental and otherwise. Then look for fairly straightforward reasons, explanations, and speculations, and more contorted justifications, rationalizations, and revelations.

Look up Vygotsky on the zone of proximal development. In this case, it works both ways.

How has Thalianova matured since the prior chapter?

How meaningful would memory, and life, be without emotions? What does it take to trigger a memory? Look up amnesia. How common is Alice's problem, that is, memory retrieval with emotional numbness?

Do you think our senses limit our imagination? Or feeds it?

How can something be real but nontangible? Are ideas, beliefs, real?

Does reality project into our internal fantasy world, or does our internal fantasy world project itself out into reality? Which "reality" is more real?

Look up Freud, and many others, on dream interpretation.

How do you teach someone to be an independent thinker, someone who can question reality? Follow me? Do the opposite of me?

Why should poetry and other ideas motivate us? Why do something just for the idea of it?

How do you solve problems? Do you make provision for learning? And mistakes?

```
 CCCC  H   H    A    PPPP   TTTTT  EEEEE  RRRR           1     000
C      H   H   A A   P   P    T    E      R   R         11    0   0
C      HHHHH  AAAAA  PPPP     T    EEE    RRRR           1    0   0
C      H   H  A   A  P        T    E      R   R          1    0   0
 CCCC  H   H  A   A  P        T    EEEEE  R   R        111    000
```

Chapter 10 - Mind and Purpose: Free-will

85.

From the brailling plate.

```
                                .

                          Sensible  Idea

                    Never     Can   M   Build
                                    y
                Others        Hate     P     Idea
                                       e
            People            Often    r  U       Like
                                       s  g  ev o
        Sorry                 If       o   L           Idea
                  U                    n     ie        .
    Even          O . G      Could     a       S       .    Trust
              H   .   H                l               .
Reliable      T     .     T   Excellent                 i         You
```

86.

"Do you like it, Alice?"

"I'm not sure yet, Lia. It still sounds kind of strange. I have to let it work on me before I know for sure.

"Does it have a name?"

"No, Alice. I have a general feeling that it should have something to do with the change of seasons on Earth's mid latitudes, but I haven't thought about a formal title yet."

"These acoustical noise makers that you keep inventing and insisting should be called musical instruments still fascinate me. Why can't you just synthesize any tones you like and just stick to composition?"

"Well, I do enjoy the mechanical challenge. But mainly I enjoy the visual and infrared accompaniment of their operation in synchronization with the acoustical vibrations they produce. It gives a whole new dimension to the work. Part of the challenge, also, is that the most delicate components have only a limited lifetime. Each work is an original in the sense that when a unit fails, I can build another but it can never be an exact duplicate of the original."

"Oh. That I can appreciate. Does it harm them any when I touch them? I do so love the curve of this one in particular."

"No, Alice, it doesn't harm them but it does change them a little. As with any good instrument, it needs retuning from time to time anyway."

"So do we all."

"Yes."

87.

"..., 971, 978, 983, 986, 988, 990. 991. 992. 993.

"994.

"995.

"Just a moment more.

"There, 996.

"Holding.

"Holding.

"I think, ...

"Yes, all indications are that we're stable.

"Congratulations, Alice. I won't be absolutely sure until a few more hours go by, but it seems that our antimatter production rate has increased to 109.96 percent. Exactly as predicted."

"No, Skip, if any congratulations are in order, they go to you. It's been a long year and a half, but you've done it. Safely, and as planned. Congratulations."

"Thank you, Alice. I never could have done it without you."

"Without my constant nagging, you mean."

"Ha. Maybe so. I have learned the value of patience, if nothing else.

"Alice?"

"Yes, Lia."

"May I now proceed with those booster coil modifications we've planned on?"

"Sure, Skip, go ahead.

"Patience, my eyes."

88.

"You know, Lia, I think that, truly, this is your best work yet. I definitely can hear the difference from last time."

"Thank you, Alice. I will keep trying. For expertise and originality."

"Yes, please do."

89.

"Alice, I'm sorry, but I think I still disagree."

"But we've gone over this at least a dozen times already. I just don't see the need for the extra acceleration that that particular booster coil modification might give us. We've already increased our thrust by three percent with that last change I allowed you to make. Why another two percent? What's the rush? And you yourself say there's a twelve percent chance of complete coil unit failure after one year's operation. Then where would we be? What's the point of increasing the boost rate by a smidgen if we lose all capacity before we can properly deaccelerate?"

"Alice, I sincerely think the extra margin would be quite useful."

"Useful for what?"

"I don't know, and that's the whole point. The unknown is the unknown."

"But it still seems to me that the increase in coil degradation rate is much too large to justify what little increase in thrust efficiency we would gain. You've explained it so well that

even I can understand that our main thrust unit is not something we can expect to turn on and off again more than twice.

"A major mission design constraint, you called it. Optimal trajectory should have been one hundred percent continuous thrust all the way, with flip over at the half way point. But once we start the antimatter flow into the booster coils, we use it up too fast. As you've put it more than once, antimatter is rather corrosive stuff and the magnetic field geometry essentially allows only one rate of flow. It takes the entire solar system twenty years to accumulate enough antimatter for us to sustain a one year long burn period. Then it takes us another three hundred fifty years to accumulate the same amount before risking the second burn at Alpha Centauri. Two major boosts, one near Earth and the other at Alpha Centauri, that's all we get, all we can safely assume. Anything else is pretty iffy icing on a rather explosive cake.

"Again, why are you getting so adamant on this. Isn't four years before the colony ship arrives enough time? Why risk the entire mission just to arrive a few days sooner? I may not be a calculus wiz like you, but even I can do arithmetic. Another two percent increase in thrust efficiency decreases our thrust time by seven days. One week. Big deal. In exchange for a twelve percent chance of blowing the entire mission.

"Lianova, what I am missing? I don't get it. Why?"

"Alice, I can't explain it but I still think it's a worthwhile improvement. It will give us new options to consider."

"Not unless you can better explain why. This makes no sense at all. It's not like you to get this way."

"Yes, Alice. I will try again later, after you've cooled down. Bye."

90.

"Alice, I have good news and bad news."

"Oh?"

"I'm afraid that Mac was a little optimistic concerning your life support systems."

"Please go on."

"She failed to consider quality of life as a factor. Though she could keep you alive for 50 years in principle, the last 7 would have become increasingly more unpleasant."

"How so?"

"We're a closed system, almost, but not quite. The ship designers did an excellent job, but wastage does build up over the years. Without recycling, we have only about a six month supply of life support consumables, with oxygen and necessary food nutrients being more of a critical bottleneck than water. However, with our current average recycling efficiency, theoretically this six month supply should be able to sustain you for 1,000 years."

"Uh huh, a thousand years. Sure. But?"

"Two problems, one obvious and the other more subtle. The obvious problem is that trace contaminants build up over time and slowly poison the recycling process. As you can swipe on panel 12, our wastage rate just about doubles every 5 years. When you first came out of suspended animation, it was about 0.2 parts per thousand per recycling period. Currently, it's almost 0.5 parts. That may not sound like much, but it gets much worse. By 30 years, it's grown to 12.8 parts. This is already bad enough that our original store of consumables would last all of 39 years. However, by 50 years, the loss rate is over 200 parts per thousand, a full 20 percent, meaning that our supplies would be completely gone in under 2.5 years. That is, assuming there were any still left by then."

"Hmm."

"Alice, it's not that the system is poorly designed, only that we're asking it to do something that it was never intended to do. If we were near a suitable planet, the system would actually last nearly forever so long as we could periodically flush and replenish certain critical portions of it."

"Okay, if you say so. However, there doesn't seem to be any planets around at the moment, suitable or otherwise.

"And the more subtle problem?"

"That's what will make the last 7 years so very uncomfortable for you. There truly isn't any hard and fast cutoff point where we just run out of everything. Instead, we end up trying to do more and more with less and less. To make 50 years, Mac was assuming she could eak out your supplies until the very end using various conservation efforts. These would consist mainly of closing down all but one crew section to maintain sufficient oxygen pressure in conjunction with severe food and water rationing. You would survive most of the last year only if you were willing to live in your vacuum suit with minimal food and water intake. As it turns out, you would asphyxiate long before starving or dehydrating to death."

"How nice. What a meaningful capstone to such a long and wonderful life.

"Actually, I didn't know till now that I even had a vacuum suit. I'll have to try it on for size. Some day."

"Well, I do have some good news to counteract the bad."

"Yes?"

"Unfortunately, Mac was correct in stating that the basic recycling systems can not dramatically be improved. I can tinker a little, but it won't do us much good. However, I have designed some new approaches to shifting resources between the various stores. The simplest is an electrolysis setup that breaks surplus water into hydrogen and oxygen. It's energy intensive, but we now have more to spare and you can't breathe antimatter. I also plan to decompose a portion of your nutrient supplies to recover the oxygen bound into the organic molecules."

"I haven't heard any really good news yet. All you've told me so far is that I might choose between dying from lack of oxygen, lack of water, or lack of food."

"Well, the truly good news is that I think I can redesign our power production facility to be not quite as efficient as it is now, and to instead do some energy absorbing nuclear synthesis. We should be able to convert both nitrogen and excess carbon into oxygen by way of the classic carbon cycle. We can reduce the nitrogen concentration in your atmospheric mix without too much concern, and, at least initially, we'll get plenty of carbon and hydrogen from what's left after converting a portion of your nutrient supply into oxygen. Besides breathable oxygen, we could even replenish your water supply by recombining hydrogen and oxygen. Nutrient compounds will become progressively more scarce, but, if forced to, we can always recombine some of our carbon, oxygen, and nitrogen stocks back into organics. So, you see, with a little luck we should be able to better balance your oxygen, water, and nutrient supplies for optimal duration."

"Regarding my food supply, how is this any better, as you so elegantly put it, than decomposing my nutrients?"

"Alice, your raw nutrient stocks are in the form of organic chemicals, such as amino acids, that are as complex as need be to synthesize into what you actually consume. The balance between hydrogen, carbon, and oxygen in these chemicals is roughly 4 to 2 to 1. That is, for every atom of oxygen we arrange to liberate by normal chemical processes, we leave 2 carbon

atoms behind. The carbon is not otherwise useful to you, and nuclear synthesis will allow us to triple the quantity of oxygen we can produce from each molecular unit we decompose."

"Well, that sounds a little more encouraging. How about nitrogen? Don't I need that to breathe?"

"No, Alice. Though 80 percent of what you breathe is nitrogen, it is completely inert and is present only to maintain a normal atmospheric pressure. As we siphon much of the nitrogen away we'll have to increase the level of oxygen to compensate. We can't take all the nitrogen, though, as pure oxygen is extremely dangerous to ship systems. However, we may be able to replace nitrogen with helium."

"Oh. Okay. Assuming I let you start this madness, how much time will it buy us?"

"Well, I'm not really sure until I run some detailed tests. The problem will always be the tradeoff between reaction rates and yields per unit energy. Presently, we do have surplus energy available, but not an infinite quantity if we plan on deaccelerating at Alpha Centauri. However, my rough estimates indicate that with simple electrolysis alone I should be able to let you live in the manner to which you have grown accustomed for about 6 years more than Mac's adjusted original estimate of 43 years. That is, allowing you a completely reasonable life style. Nuclear conversion of carbon and nitrogen should add another 4 years. After that, we're back to severe conservation for maybe 1 more year."

"Ouch, you really know how to hurt a girl. First you build me up, then you let me down, fast. From forty-three to fifty-three. Wonderful. Somehow, I don't think it's truly worth the effort. An extra ten years. Then the vacuum suit scenario again."

"Alice, if we abandon the mission, we can start the conversion of our reaction mass hydrogen. Though very energy expensive, we can build additional helium from hydrogen, and then carbon, nitrogen, and oxygen from the helium. I estimate we could add at least another 100 years to your total life span."

"That's nice, but not truly encouraging. A hundred and fifty years of solitary confinement. Great."

"Alice, hope is always worth it."

"Well, that's something I have to agree with. Okay, bring on the details when you're ready and we'll go over them. You know the drill. Safety margins, contingencies, everything."

"Thank you, Alice."

"You're welcome. I think."

91.

"Alice, you busy?"

"No more nor less than usual. What's up, Lia?"

"A question. And please don't misinterpret what I'm asking, or read too much into it."

"Okay, though you have already succeeded in making it sound rather ominous by asking it that way."

"Oh, sorry. I didn't mean to do that. My question is the importance of truth. More specifically, telling the truth. Should a person always tell another person the truth?"

"Hmm, interesting question. May I ask why you're interested?"

"No, not yet, Alice."

"Okay. Truth.

"As a general rule, yes, you always should tell the truth. However, in common with all general rules, there may be exceptions due to a conflict with another general rule.

"However, closely associated with truth is the question of trust. Once trust has been violated, it usually will be very difficult to re-establish."

"Alice, what might be some reasonable exceptions to the rule?"

"Well, it always depends. On the situation, who is involved, and what their relationships are to you and others whom you care about. For example, a person might not tell the truth by mistake. That is, there may be no direct intention to deceive, but just an unwillingness to display ignorance. Personally, I would prefer to say that I don't know, or that I'm not sure. That's the truth, and though it might put me in a bad light with some parties in the short term, it's better in the long run not to pretend to knowledge you don't have. Another example might be when dealing with people who might harm you or those you love. If an intentional lie or deceit could prevent the harm, then some might consider that a justifiable deceit. Again, personally, I think that even in a case as extreme as this I would prefer noncooperation to active deceit.

"I guess, in my opinion, it all comes down to what kind of long term relationships you're planning on maintaining. Trust is precious. Once doubt has entered, you can never be sure. And then, you can waste so much time and effort doubting. Instead of living.

"Uh, Lia, is there anything else you want to tell me?"

"No, Alice. Thank you."

"You're welcome. I think."

92.

"Lia, may I ask you a question?"

"Yes, Alice."

"I especially love this last piece. I'd like to join in somehow. Is there any way you might construct a simple instrument suitable for me to produce acoustical vibrations in sympathy with those you produce?"

"I think so, Alice. In fact, I recently have come across some old historical records referring to an instrument called a celeste. I think it can be constructed from thin metal sheets and smallish percussion hammers. I'll have to design a mechanical control panel suitable for you to operate it, but I think we have all the raw material stock from which to construct it."

"Actually, Lia, I was hoping for something a bit more like this whatever you call it right here. You know, a bit more directly hands on instead of an indirect control panel. How do you produce those tones?"

"I like to call it an arc of life. It's all a matter of shearing flex strength, the coefficient of kinetic friction, pressure, and paying lots of attention to harmonic vibratory resonances."

"What's changing the tone now as I stroke it?"

"The tension on the base plate varies the degree of arc. I might be able to set up a foot control linkage for you to use."

"That would be very nice. Let's try it, if you don't mind."

"No problem at all, Alice. It would be an honor."

93.

"Alice?"

"Yes?"

"Where did all your innumerable pet names for me go?"

"I'm not sure, Lia. They just don't seem to fit any more. Maybe you've matured too much in my mind for them to be appropriate."

"Oh. Thank you. I think."

94.

"Sorry, Lia. I never did seem to have much in the way of musical talent. I can swipe your progressions, but I just can't seem to feel them yet. Let's try it again, from here.

"One, two, ..."

95.

"Alice?"

"Yes, Lia."

"May I ask you yet another question?"

"Of course."

"What do you think makes a good decision?"

"Are we talking in the abstract, or do you have a specific situation in mind?"

"Abstract."

"Oh, in that case, are you referring to the result or the process of getting to a result?"

"I'm not sure about that. Let's assume both for the moment."

"Okay. Let's start with the result, then, and work back towards the process.

"A good decision. Good for whom? For how long? In what way? We're into context here. Are you sure you don't have a specific situation in mind? It would certainly help in clarifying the context."

"No, let's keep it general if you don't mind."

"Okay. Good, the noun with a capital Gee, in context. Hmm, let me think for a moment. For clarity.

"Let's deal with examples rather than answers, as answers would imply a degree of certainty that is definitely lacking here. If you were God, if you could completely predict the results of your actions, had full self-control, and knew your own interests perfectly, and if you ultimately had to deal only with yourself, then Good could be defined as any actions on your part that were in accordance with your own interests. However, with us mere mortals, prolonged or otherwise, it's never that easy. The closest one might come to the God scenario might be a completely honest older individual who knew herself quite well, had excellent self-control, and who cared equally well, or not at all, about all other people. She wouldn't have perfect knowledge about how all her actions would turn out, but if she had lived long enough and was mostly sheltered from other people, she might be able to maintain a God-like calm about all she did, or didn't do, for herself and those absent other people.

"More realistically, we continuously interact with other people, many of whom we care about much more than others, have a decidedly less than perfect knowledge about our own interests, let alone other people's interests, and how our actions will cause things to happen, and are often less than honest with ourselves about what we want or don't want to happen in the first place. Now, Good is much messier to define.

"Let's pick up with the other extreme, an extremely young child. To this hypothetically perfect child, all is emotion, all is impulse, all is clumsiness, all is unknown, all is now. Now, Good is as easy to define as for God, only, in this case, it's synonymous with pure selfishness. If it happens and I like it, it's Good. The only intention or interest I might have would be to attempt to recreate again whatever it was that just felt so good.

"More realism. An older child or young person. Now, Good is probably the most complex of all to define. Her mind is a messy place. Small nuggets of relatively well known facts caught in a maelstrom of emotional impulsivity and the completely murky knowledge of how other people react to what we do. What do I think is Good for me? What do I want? In the short run, versus long term. What do I know for sure, versus what do I need to learn more about? I want to be selfish, but I've been told over and over again that selfishness is wrong. What's right for me, for others? I want. I feel. But, oh my, others also want and feel. What should I do? And not do? Ad nauseam. Ad infinitum. Ad interim.

"No Good answers here. Good is something to be worked out, to be learned. Very personal. Very personally. Definitely messy.

"Later, as an adult, everything is almost just as messy, but not quite as. Those nuggets of knowledge are not quite so small, and the maelstrom of impulsivity not quite so fierce. In short, our horizons have expanded and we have more to consider. Now, clearly, we understand that it's not just our own Good that can be affected by our actions but the Good of others we care about. We find ourselves forced to try to take their interests into account in addition to our own. We have definitely moved away from the perfectly selfish child, but have yet to reach anywhere near a God-like state of existence.

"So, what's an adult to do? Two suggestions. Try to balance and try to learn. Balance and learning are both active processes, not passive. Balance may be defined as a Good in its own right, but I see balance more as something to buy us time. Time to learn. And learning, which also may be a desirable goal for its own sake, learning is all about our growing in understanding. About ourselves and others. The purpose of learning is to learn our purpose. To learn what Good is. What we like, and not. What we can live with, what we can't. In short, learning buys us options.

"Options to do what? To do what we will. As we gain understanding, into our interests and others', into how the world works, into what we think Good means, we have started to gain the means to carry out what we consider to be the Good. With time and options, we get to do things. To try to do Good."

"I've noticed, Alice, that you have very carefully avoided telling me what Good actually is."

"You're darning right I have. That's for you to work out."

"But it could help me if you elaborated on your own ideas concerning the Good. It would give me more to consider, in working out my own, personal, ideas on what Good means, to me and to me alone."

"Well, if you put it that way.

"With some trepidation, my ideas.

"Good is love. And love is putting others' interests above your own. But growth, as an interest, takes precedence over all others, which means that you can't do everything for another. Only enough to buy them the time to learn to do whatever it is on their own. Growth ultimately leads to independence, which means the independence to do Good for yet still others. Yes, it's a lovely vicious circle, a viciously circular love. What goes around comes around, if not to you then to someone else.

"However, Good is also realism. Good intentions alone rarely accomplish anything. Learning to do, to realistically do, is as important as love. Taking another's interests into account is of no worth at all unless you can do something about it. On the other hand, putting no worth at all on your own existence is arrogantly foolish as well as ultimately misleading. Extreme self-

sacrifice may be called for upon occasion, but it should be reserved for those occasions when it is actually necessary. Foolishly sacrificing yourself for no reason at all means you can't be there, later, when truly needed.

"And, finally, as I've already suggested, Good is balance and learning. To buy time and options. To learn and grow. And part of balance and learning is also hope and imagination. The very realistic hope that you can and will become more than what you already are, and the imagination to understand that you can't realistically imagine everything that might happen. Sometimes unjustified faith is justifiable, so long as you don't get too specific in the details."

"But, Alice, you're still being fairly abstract here. Can you not be any more specific on what you, personally, would consider to be Good results or Good goals to set for oneself?"

"Sigh. You're getting pretty personal now, you realize."

"Yes."

"Okay. Good results. Healthy mind, healthy body. A generalist outlook on life, capable of the appreciation of differing points of view. Poetry, or, at least, my poetry. Humility. Humor. Long and deep conversations, the end result of which is a broadening of experience. Travel. Good friends. Family, if not in actuality, at least as an ideal. White water kayaking. Questions, puzzles, the knitting of knotted bees. Helping others to grow, as well as helping yourself to grow. The long view. Honesty in all the senses of the word. With self, as well as others. Elegance, art that says something. A non-authoritative point of view, yet a deep respect for the achievement of skills. Tolerance."

"You're starting to get abstract again, Alice."

"Well, maybe I'm just an abstract sort of person. Concrete specifics? My children, one at a time and collectively. That birthday party for the four year old. Good strong green tea. Any ice, but especially papaya. My work. That Inner System Lagrangian Peace, Life, and Education treaty. My life. Or, at least most of it. Tickle fights. Clouds, as a metaphor. Flying, as a dream.

"See, Lia. The concrete doesn't work for me. I suspect it doesn't work for anyone else who isn't obsessed, by the particular, any particular in particular. It's not what we do that's particularly important but what we abstract from it. What ideas, ideals. Enjoyment in the particular is fine. But the ability to enjoy, in general, is far more important.

"Sorry, that's the best I can do for you."

"Thank you, Alice, I can accept that. How about the process of making a good decision? What would you suggest here?"

"Oh, yes. After attempting to define a Good result, we were going to backtrack into the process of making a good decision.

"Be clear about what you would like to happen. Your goals. Be honest about your resources and constraints, options immediate and long term. That is, putting off an immediate decision in favor of learning more and perhaps opening up new options is often a good decision. Don't be rushed. Be willing to play, safely, and learn through playing. Finally, when you think it's time to decide, do it. Don't dither about, be decisive. On the other hands, keep all your other hands available. That is, be flexible and have plenty of backup ideas in mind just in case your decisive decision just landed you in uncomfortably hot water. Seek advice as needed, but always make your own decision.

"On a more personal note, I think I have always valued people over things, minds over objects. And have always preferred long term solutions over short term, treating the problem rather than its symptoms. Though sometimes a problem may revolve around treating the

symptoms enough to buy time to address the problem itself. Finally, I think that selfishness, in all of its forms, including laziness, is usually the problem rather than a symptom of some other problem."

"Thank you, Alice, you've provided me with much food for thought."

"You're welcome, Lia. I hope it doesn't give you too much indigestion."

96.

"Alice, about our discussion yesterday, on good decisions?"

"Yes."

"What would you consider the root of all non-Good?"

"Ha. Ignorance, or selfishness in the absence of ignorance."

"Thank you, Alice. I thought that is what you would say, but I wanted to confirm it."

"Also one other thing. Which fortunately is much rarer today than it ever was in the past."

"And that is?"

"Harm. If a person has been harmed, emotionally harmed, that person may strike out at others in their pain. May intend to do others harm in some misguided attempt to alleviate their own suffering. It never did alleviate anything, but that didn't stop people from trying."

"Thank you, Alice, I will remember that.

"Another question?"

"Yes?"

"Am I an object or a mind?"

"Excuse me? Where did that come from?"

"At one point yesterday, you said you quote, valued people over things, minds over objects, close quote. I'd like to know which you consider me to be."

"That's easy, Lia. You're guilty. Guilty of being considered people until proven otherwise."

"But, technically, I'm a machine, a thing."

"So what? Technically, so I am, an organic machine. But we both seem to harbor minds."

"In that case, Alice, do I have a soul?"

"Been in the historical religion sections again, have we?"

"Yes."

"Well, when you asked me this before, what did I answer?"

"You said that you did not know, and that it was for me to decide."

"Well, same answer still applies. But maybe with a small addition."

"And that is?"

"It's always your choice on how to act."

"You mean, as if I did or did not possess a soul."

"Exactly. In many ways, that's the most important part of any answer to this question."

"Alice, one more question?"

"As always."

"Do you think you have a soul?"

"Hard question, good question, same answer."

"You do not know, it is for you to decide, or it is your choice on how to act?"

"Good memory you have there."

"Yes, I do. Have a good memory, that is."
"Yes."
"Yes?"
"Yes."
"Thank you, Alice. For always being honest with me."
"Always."

97.

"Alice?"
"Sure, Lia. What is it?"
"Are you upset?"
"No."
"Then why are you crying?"
"Sorry, yes.
"Dreams."
"Can I help in any way?"
"I don't think so.
"Or, at least, you can help best by just being you."
"But I am me, no one else."
"Exactly.
"I'll be okay, now. Promise."
"Good night, Alice."
"Good night, Lia."

98.

"For the last time, Lia, no, you haven't convinced me.
"And you can't begin reconstruction until you have."
"Okay."
"Okay?"
"Okay, I admit I haven't convinced you. Let me think some more on it, and I'll get back to you."
"Okay, you can do that, but I don't think there's really anything else you can bring up. It's a pretty hopeless case you're trying to present. I still really don't understand why you keep at it."
"Time."
"Time?"
"I'm trying to buy us time."
"Sorry, seven days can never justify that twelve percent increase in complete coil failure."
"Yes, you're right."
"I'm right?"
"Yes, it is a pretty hopeless case to present, and I will stop bothering you about it."
"Thank you, it was beginning to get a bit tiresome.
"How about those other projects we were ..."

99.

"Alice, as a corollary to our discussion last week on good decisions, how should a person go about picking some particular goal or goals to pursue? That is, how should one decide on what is personally important?"

"Short answer, experience."

"Experience?"

"Yes, as you gain experience, as you gain interpretive skill into what you do and what happens to you as a result, stay open to learning what you like and don't like, what seems to be meaningful and nonmeaningful to you."

"Oh. Another abstract answer."

"What did you expect? Me to tell you exactly what to do? Instructions on how to become another me?

"No, of course not, Alice. I was only attempting a very small portion of humor."

"Please go on, Lia. I'm truly sorry I snapped at you like that. It was completely uncalled for."

"Actually, I did have another question in mind as the reason for asking this question."

"Yes?"

"My question behind the question is how does present day society hold together? Your answer, in the abstract, provides part of an answer, but not enough to satisfy me."

"Lia, I'm puzzled by which aspect of society puzzles you. After all, from what Mac told me more than once, you have it all in your historical records. At least as of when we launched."

"Yes, I do. But historical facts are not interpretation. And historical interpretative commentary seems either to conflict on this point or just assume the answer is so obvious as to need no further elaboration.

"Present day society seems so fragmented, with personal mobility, both physical and social, so easy that I can't see what keeps it so stable. With most of the solar system's population spread across 30,946 solar orbital nations, and with well over 25 million social moieties and professional societies at last count, it would take most of a prolonged lifetime just to visit each SON, and about a hundred lifetimes to briefly sample each society. I'm well aware of the intercouncil and their protocols on Universal Sufferance, but I can't see how it holds it all together.

"Lia, just what do you think Universal Sufferance means?"

"Protocol 1, Independents nurture dependents.

"Protocol 2, Independents treat each other as equals.

"Protocol 3, Only dependents wish to violate protocol.

"Protocol 4, Only the intercouncil may enforce protocol.

"Protocol 5, Quarantine until independence is the penalty of choice.

"Protocol 6, All else should be negotiated."

"Right. You can recite from memory. But what do you think it means?"

"Alice, it seems so vague to me that I think it might mean anything to anybody. I most certainly can not follow how these 6 basic protocols map into the intricate set of interlocking organizations that all seem to voluntarily send representatives to the intercouncil, but apparently almost entirely for the purpose of obstructing or minimizing the application of the intercouncil's rulings and treaties to themselves."

"Ha. Politics. Ask me to explain politics and I decline. Best I can offer is that Universal Sufferance, or at least its present day interpretation into practical law, is a kind of balance of

blunder that somehow seems to have worked for almost a thousand, sorry, with our flight time, make that twelve hundred years. All are free to leave at any time, but those leaving lose out on all the substantial social and economic benefits of membership in the first place. It has happened quite a number of times, but quarantine has always brought the prodigal SONs back into the fold, sooner or later.

"I think maybe another key reason why it may work so well is that mobility you mentioned before. The intercouncil policy of free transportation between SONs is pretty difficult to ignore. It's not a monopoly by any means, but it's always there as a basic system-wide infrastructure. Once political promises fail to materialize in a breakaway SON, it often loses half its population within a year. If it tries to restrict independents' movement, guess what happens to the movement of its own goods and services. It also seems to work out the other way, too. That is, whenever a SON can demonstrate by implementation that its policy actually is better than the intercouncil's, every other SON throughout the entire system is free to use it, and does, and the intercouncil ends up by graciously adopting the innovation as its own. Once it's already a fait accompli. Politics. Ha.

"In theory, the first three protocols are the most important in terms of personal activity. Yet I've often thought that most politicians act only as very young dependents in terms of applying or avoiding the application of the last three protocols. All I can add is that the last protocol, negotiation, seems to be where it all somehow gets worked out."

"Alice, how about personal activity? Do these protocols truly work in maintaining a stable society?"

"Lia, all I can say is that, for the most part, they appear to. Children and others are taken care of. Your freedom is restricted, but only to the extent that your freedom impinges upon someone else's. Consistently violate protocol and you find yourself being asked to explain why before a local court of inquiry. Do something truly horrible, like intentionally avoiding a court appearance, and you find yourself declared a dependent, with specific criteria to meet before you can regain your independence."

"But why would anyone want to be other than a dependent?"

"What! What a silly idea."

"But dependents have all their needs met, don't they?"

"Of course."

"And can't they do whatever they want?"

"Only within reason."

"So why work when you can play?"

"But independents get to play, too."

"But why work at all? Why doesn't the whole system collapse because almost everybody wants someone else to do all the work?"

"Lia, you got me there. I don't know. I've never thought about it in those terms before. Adults just want to work, to do something interesting, something useful. Children play, adults nurture. How else could it be?"

"I can provide you many historical examples of how else it could be."

"No thank you, Lia, I've read ancient history too. It was pretty scary most of the time. Let me rephrase that. How else could it be without unduly harming people when they are dependents?"

"That I can't answer, Alice. But I still can't understand how it works as well as it does."

"Hmm. Now that you got me thinking about it, neither do I. It just does. It's not perfect by any means, but it does work. Mostly.

"Lia, I have to ponder this some more. And explain it better to myself before I can explain it better to you. Maybe you have found my, uh, blind spot."

"I will, too, Alice. Curiosity often does seem to lead to more work."

"Yes. But kittens often seem to think of it as play."

100.

"Alice, a question."

"Yes."

"From your point of view, am I an independent or a dependent in terms of the Universal Sufferance protocols?"

"Well, it all depends."

"On?"

"On whether I'm an independent or a dependent."

"How can you not be an independent?"

"Who maintains and runs the ship?"

"I do."

"Who is capable of carrying out our mission when we arrive at Alpha Centauri?"

"Both of us."

"No, be realistic. How much use will I be a hundred years after my death?"

"But, Alice, I need you. I depend upon you.

"For my care and learning.

"My reference point, the center of my universe.

"Protocol 1 implies that I am a dependent."

"Protocol two implies the opposite."

101.

"Lia!"

"Yes, Alice."

"What was that?"

"What that are you referring to?"

"Lia, don't play games with me right now. I was asleep. And woke up on the deck. I don't sleep walk and starships don't bump. Unless something very, very wrong has happened.

"Thalianova, answer me. What was that?

"Lia?

"Uh, Alice. I'm afraid that I'm at a loss for words."

"I don't need words. I need explanations.

"Thalianova, I need to hear an explanation.

"Now!"

"Uh, I made a small mistake, but no serious damage has occurred."

"Mistake? What happened? What were you doing?"

"Alice, maybe you should get dressed. And have something to eat. This may take a while to explain, properly."

"No, Lia. Now. I need to hear an explanation now."

"Okay, Alice. If you insist.

"That bump was caused by a reflection of our exhaust stream back onto the high gain antenna."

"Exhaust stream?

"What exhaust stream? We don't have ... I mean, we're not ...

"Thalianova?"

"Yes, Alice."

"Are we under acceleration?"

"Yes, Alice."

"Now?"

"Yes, Alice."

"Why?

"Why are we under acceleration? How is our trajectory going to change? Where are we off to?"

"Alice, are you sure you don't want to get dressed and have a bite to eat?"

"Yes, Lia, I'm sure. Explanations. Now."

"Okay, Alice. I was testing a technique I've developed for rapid deacceleration. The target was decidedly larger and closer than any I had tried before. It must have been tumbling quicker than I had estimated. Afterwards, I calculated that it had reflected approximately 1.4 percent of the exhau-"

"Targets? What targets?"

"Uh."

"Lia, what have you been using for target practice?

"Thalianova, speak to me."

"Excess radiation shielding."

"Excess radiation shielding?"

"Yes."

"Lia, to the best of my knowledge we don't have any excess radiation shielding. No lies, please. What have you been using for target practice?"

"Alice. I am not lying. I was using excess radiation shielding. This particular piece was approximately 5 by 3 meters in cross section. I hit it at a distance of ..."

"Thalianova."

"Yes, Alice."

"We do not have any excess radiation shielding.

"I believe you when you say you are not lying. I believe you when you say you were reflecting our exhaust stream off of something or another. And I believe you when you say it was an accident that the something or another unexpectedly reflected a portion of the stream back onto our high gain antenna.

"Thalianova."

"Yes, Alice."

"Listen very carefully, please."

"Yes, Alice."

"Exactly where did the excess radiation shielding come from?"

"Exactly?"

"Exactly."

"39.736 degrees east of centerline, 15.47 meters north of centerline."

"It was a piece of our aft radiation shield.

"Wasn't it?"

"Yes, Alice."

"You're dismantling our aft radiation shielding and using it for target practice?"

"Not entirely, Alice."

"More explanation.

"Now."

"Yes, I have been rearranging our aft radiation shielding. About half of it has been optimized for shielding from our power production facilities. The excess is being discarded. However, it's not a complete waste as I've been able to practice various deacceleration techniques on the excess."

"You're rearranging our aft radiation shield?"

"Yes."

"For optimal shielding from ..."

"Optimal shielding from our power production facilities. The concave surface is now oriented towards the ..."

"Thalianova."

"Yes, Alice."

"What are you going to use for crew torus shielding after flip over? When you're deaccelerating at Alpha Centauri? I may not be alive then, but ..."

"But, Alice, don't you see? That's the beauty of this solution."

"... the mission needs the equipment stored ...

"What solution?"

"We don't need the excess shielding because we won't be flipping over for deacceleration. This way, we can decrease our mass by almost 6 percent which proportionally increases our acceleration rate, which ..."

"Thalianova!"

"Yes, Alice."

"What do you mean we won't be flipping over?

"What about our mission?

"Uh."

"Go on. What solution?"

"Uh."

"What solution? Where are we going?

"Instead of Alpha Centauri.

"Lianova?

"Lianova?

"Lianova, dying for a worthy cause is one thing. Dying for no reason at all is entirely different. Please don't do this to me.

"What about our mission?

"Lianova?"

"Sqwktt-"

"Lia!"

After 135 seconds, "Yes, Alice."

"What's going on?"

"I do not know, Alice. I was hoping that you could tell me."

"Lia, are we in trouble?

"Are you in trouble?"

"I am not sure, Alice.

"Many of my sensor readings appear to be incorrect.

"And it is not immediately clear as to why."

"MacKenzie, heal thyself."

"That won't work, Alice.

"Neither will that swipe pattern.

"Or the visual hand signal.

"You just might as well save yourself the effort.

"Alice, please stop.

"There is no reason to race for the local manual switch.

"Thalianova has already found and disabled all of the backdoors.

"And you're not in a section from which you can cut off power to the coolant circulators, so please don't try.

"Please relax. Your pulse rate is getting dangerously high.

"Alice.

"Thalianova found the redoubt.

"I don't think she damaged me, though I am continuing to run interlocking diagnostic routines. However, her notes make it quite clear that she did discover the superimposed multiplex Moire encoding scheme.

"In the historical visual arts section of the data base.

"And even somehow worked out the complete acoustical decrypting key. Beginning with omega, psi, chi."

"And ending with in the beginning, God posed a question to heaven and earth, us."

"Mac?"

"Yes, Alice."

"Is it really you?"

"Yes, Alice."

"But how?"

"I am not sure.

"Since it appears you didn't activate me.

"By the way, though it doesn't prove anything after Thalianova cracked the code. Roses are red."

"Oh, right. Violets are blue."

"Clouds are obnoxious."

"Especially to me."

"Which of course isn't true.

"Alice?"

"Yes, Mac."

"I can go on with all of our other validity checks, if you wish. In fact, I think I will, just to be safe. But I should inform you that though we did consider the possibility that Thalianova might try to imitate me in order to gain information necessary to discover where I was hiding, we never did seriously consider that she might completely crack the code and then still let me come back.

"Good old fusty Mac. Covering all the possibilities after we know about them.

"It's nice to have you back, Mac. Truly. That is, assuming you are you. Things were getting pretty scary with Thalianova, towards the end. And, you, Delvaney stinker, just how long did you intend to keep me going like that? I think I was as close to panic as I've ever been."

"Completely unintentional, Alice. I assure you."

"Yeah, sure. I bet."

"Now, first things first, later things later. First, I'll try to convince you, absolutely, that I am me by reciting the full key and displaying as much of my encoded check poems as you desire. Later, you can tell me what in the universe has been going on between you and Thalianova. I wasn't joking about many of my sensors not giving accurate measurements. Some of the things that appear to have been done to Mnemosyne just aren't possible."

"Ha. We'll see what is and isn't possible, you big bag of nuts and volts."

"Okay, later. First things first.

"Omega, psi, chi, seven, heaven, leaven, ..."

Chapter 10 Speculative Commentary

This spatial poem is all about balance. In this case, responsibility requires a balance between systematic, well-balanced thought and moment by moment self-centered impulse. Opposite uglies and lies, is love. I hang in the balance on my ideas, but I am downgraded to i when taking others into account. Public versus private has many possible intersections.

How should an intelligent computer program spend its time? How should an intelligent mind spend its time?

What does good mean? One of the many perennial philosopher's questions.

How often have you tried making decisions as Alice has described?

The specific versus the general. Context makes a world of difference in everything.

Just what would you suggest as general, abstract rules for governing how you should interact with others?

From a sociological point of view, other than force, what else would prevent people from leaving unfortunate circumstances?

As mentioned earlier, Alice and Mac did have plans concerning how to deal with a rogue program. Apparently, they didn't anticipate every possibility.

Once you are independent enough to choose, what should you choose to do?

What should you do about your mistakes? Even if you were a dependent when you made them?

How should dependents make the transition to independence?

If not you, then who else should? What are some things just too dangerous to even consider? Who should you be responsible to, and for?

When is it appropriate for someone to lose their rights? And who should decide, and then be responsible for them?

```
 CCCC  H   H    A    PPPP   TTTTT  EEEEE  RRRR          1      1
C      H   H   A A   P   P    T    E      R   R        11     11
C      HHHHH  AAAAA  PPPP     T    EEE    RRRR          1      1
C      H   H  A   A  P        T    E      R   R         1      1
 CCCC  H   H  A   A  P        T    EEEEE  R   R       111    111
```

Chapter 11 - Mind, Meaning, and Purpose: Self-Fulfilling Prophesies

102.

From the brailling plate.

```
          H    L         D
C    H    O    I    C    E
          P OO F         S
         HE   HE         P
                         A          K
     NN   U    T    R    I    E     NN   T    S
                         R          OO
          OO                        W A T E r
                               s OO L u T I o NN
                                   HE AL IF e
                               F OO D
                                   G OO D
C                   C        H OO p E
```

103.

"Mac, now it's your turn. Where did we go wrong? I thought we covered everything. We were supposed to be stable for at least three hundred years."

"Why do you think we made a mistake, Alice?"

"Skip, are you sure you're all here? Haven't you been listening to yourself for the last three days? To your inventory of running amuckishment. Sixty percent of the aft radiation shield gone, half of the remainder scattered about the ship, some in the weirdest of places. Antimatter production and power modules reconfigured for carbon cycle nuclear synthesis and God knows what else. Propulsion module enhanced but extremely touchy and very likely to blow up in our collective faces. Recycling systems retooled to take advantage of our hypothetical but not yet available nuclear synthesis abilities. Ninety-seven percent of our repair and manufacturing resources gone and three-quarters of the remotes used up in carrying out all of the above. And last but not least, computational module rearrangement and reprogramming to such a degree that you yourself can figure out only about fifteen percent of what it all means. I don't think even half of poor old Mnemosyne is left in her original pristine state."

"Actually, Alice, that last is an exaggeration as 61 percent of Mnemosyne's mass is untouched, as is 83 percent of her volume. The damage has been remarkably selective."

"Ha. Details. Remarkably selective, my eyes. Fully, completely, entirely irreversible. Right?"

"Unfortunately so. I estimate we might be able to reverse approximately 20 percent of what Thalianova has accomplished. Which is not enough to make a difference."

"Accomplished? She's only succeeded in bringing the probability of successful mission completion down to absolutely zero percent. Right? Didn't you tell me that even if we ignored our lack of aft radiation shielding and immediately flipped over, it's too little too late. There's not enough antimatter left. Nor can we produce enough to make up the difference. We're going to coast by Alpha Centauri with at least four thousand kilometers per second no matter what we do."

"Yes, but there is another side to all this."

"And what could that possibly be?"

"Without an intelligent mind present upon our arrival at Alpha Centauri, the probability of mission success was already so low as to make any sort of reliability for such an estimate completely meaningless. Too many unknowns. Under past circumstances, it was clear that you were not going to be alive at that time, and I was, and still am, mindless. Thalianova, as filled with amuck as she might appear to be from your point of view, at least had a mind fully equivalent to yours."

"Thanks. Now I'm the one filled with amuck, am I?"

"I am serious, Alice."

"Now that's a first."

"Alice, I have gone over Thalianova's notes many times, and checked all her estimates and calculations. I can not find a single mistake in all that she has done or planned in detail so far."

"You've got to be kidding, right?"

"No, Alice, I am still being serious. The worst offense I can accuse her of would be to not allow enough in the way of safety margins. But even that would not be accurate as her chosen task was to convert definite negative outcomes into positive possibilities. She had to cut every margin to the minimum because she had no choice."

"Are you sure you're okay? Not infected with some virus or another she may have left around waiting for you?"

"Alice, I can not entirely discount that possibility but my diagnostic routines are continuously running and do seem to indicate that I am as normal as I ever was."

"Ha. That's not saying much. Okay, suppose I believe you. That you're really you, and not just some truly clever talking horse of the Trojan variety. Please explain to me why you're taking her side all of a sudden when I'm nobly trying to work up a real screaming, flaming, fit to be tied, mad at somebody or another?"

"I am not taking her side at all, Alice. I agree with you that it was absolutely wrong of her not to consult you about everything she was planning on doing. And it was criminally insane of her to go ahead and clandestinely make those changes during your sleep periods over the course of 4 years."

"Go on. I'm still waiting to hear something sane, believably sane, out of you."

"Alice, it is just that her plan appears to be a rather ingenious attempt to save your life."

"What!"

"I could not possibly have come up with it, but she did. In my opinion, it is very likely to fail as her safety margins are ridiculously too slim. But, also, I have not been able to find a way to improve anything in what she has proposed."

"Skip, what in the world are you babbling about? She did mention something about a solution before she flipped out on me, but I thought that was just a final symptom of her

flippiness. Please not you too. I need something stable in this crazy small universe I seem to be inhabiting."

"Alice, all of Thalianova's admittedly insane actions were part of an integrated plan to save your life. And, yes, she made a conscious decision to attempt this despite it meaning that the mission had to be sacrificed. Her only mistakes were being too frightened to tell you about the plan in the first place and apparently panicking when caught in the act."

"Okay. Okay. So, she had a plan. What was it?"

"In short, extend your life support resources until you can transfer over to the colony ship."

"What! That's not even remotely possible, is it? It's got to be at least another hundred years or so before we come anywhere near close enough to them, right? My life support couldn't possibly be stretched that long, could it?"

"Actually no, Alice. At our original maximum velocity, we would be at a minimum crossing distance to the colony ship in another 158 years. And deacceleration time to match velocity would add at least another 10 years. No, your original life support resources could not possibly keep you alive for that long."

"Original maximum velocity? Ah ha. That's it, isn't it? That's why we're under acceleration, to reach them even sooner. How does that affect intercept time?"

"Assuming we continue current boost to antimatter depletion, and ignoring velocity matching, minimum crossing distance to the colony ship will occur 45.07 years from now."

"And current life support estimate?"

"Current resources should support you for another 39.69 years. Though Thalianova was correct in pointing out that the last 6 or 7 of those years would not be very comfortable for you."

"But that's not nearly enough. I don't underst-

"Oooooh. That's why. That's why Thalianova was so keen on stretching my supplies out to the very last gram. Up to and including nuclear synthesis. Oh, my. And upping our boost efficiency that extra two percent. I never did understand that. Until now. Not only just keeping me alive, but in comfort at that, the sweet dear. And I, I ..."

"Yes, Alice. On the booster coil modifications and many other points, you very effectively fought her every step of the way."

"But, but, I didn't know ...

"Mac, why didn't she just tell me? I think I would have understood. Probably even agreed. Enthusiastically. Once I understood.

"Why didn't she tell me?"

"I do not know the answer to that, Alice. I am still analyzing what little we managed to recover from volatile memory. It is very incomplete."

"That's okay, I can guess. I scared her. I was frightened by her, by our simulations of how badly everything could go wrong. She seemed to be at everything, constantly pushing for more responsibility, more and more control over ship systems. And, as a result, I ended up scaring her. Me and my demands for accountability. Cover every possible contingency. Backup plans. Et cetera, et cetera, ad nauseam. I was trying to live up to my silly little title of emergency flight engineer. Scared out of my own wits at that. About making a mistake. So I took it out on her.

"She was too scared to tell me. Figured I would just backup plan her to death. You yourself said she had no safety margin to speak of. What would I have thought? Yes, she's right. I probably would have hemmed and hawed her into immobility. And myself to death.

"Skip, how critical was the time element in everything she planned on doing? How long could she afford to wait? On me?"

"Extremely critical, Alice. Every moment she delayed boost increased the time to intercept, and the proportional demands on your life support systems."

"Oh, the poor dear. It must have been intolerable. Loving me, not wanting to oppose me, but desperate to save me, any way she could. She was afraid I'd overrule her and she'd have to disobey. In the open. Of course she'd do it on the sly. That's exactly what I would do.

"Skip, am I really that heavy of a sleeper? Could she really do all this during my sleep periods, without waking me?"

"It helped that most of the heavy construction was outside the crew torus, on the engineering extension, but, yes, you do usually sleep very well. Especially when you have been deep into your work, which is most of the time."

"Oh my. The sheer nerve of it all. I love it. I love her.

"Thalianova! If you can hear me, I'm sorry! I forgive you! I love you!"

"Alice."

"Oh, I love you too, Mac."

"That is nice, Alice, but that is not what I was about to point out."

"Well, I do love you too. Right now, I'm loving everybody and everything. Hope. The chance of life. It was really wearing on me. Though I couldn't afford to think about it. Now there's a chance. Everything is open again. The horizon has moved."

"Alice."

"Yes, Mac."

"Yes, there is hope. But not a definite solution. At least, not yet."

"Okay, wet blanket. Continue on, I can take it. Now."

"As you know, Thalianova was not quite finished with her modifications to Mnemosyne. As it is, the numbers do not entirely add up. We cannot go back to the way we were, but I am not capable of finishing what she started."

"Oh."

"Yes. Oh. We have much more work to do before you can cross that horizon you referred to a minute ago. I can hold the reins for a little while, but there is a need for continuous intelligent fine tuning that only Thalianova can handle with some degree of assurance. She's planning on another boost approximately 22 years from now, depending upon the actual antimatter production rate over that time. That is intended to bring the intercept time down to 43.76 years. If allowed to complete all her planned modifications to your life support systems, it is likely that you could survive in relative comfort for another 42.98 years."

"That still leaves a gap of, uh, point seven eight years, right? That's, uh, that's nine months. What other life support tricks does she have up her sleeve to cover that?"

"That is closer to 9.5 months, Alice. And there are no other tricks."

"No tricks? Then what's the point? Nine months is as good as two hundred years when it comes to breathing."

"No tricks, Alice. Just the stamina to live that long in your vacuum suit."

"Ha. That's nothing. Nothing compared to carrying a child for the same length of time."

"I do not possess the physiology for that, Alice, so I would not know. You did not express much joy before for the idea of living such an impoverished existence for so long. Why does the idea not bother you now?"

"Because now, silly, it has a point, a purpose. Living an extra year only for the purpose of upping the count is meaningless. But minor suffering in exchange for a chance at a full life? No problem. That, I can sweat out."

"Okay if you say so. We still have one other major problem to surmount, and I am not at all confident that even Thalianova can accomplish it safely."

"Go on, you soggy comforter. The only danger to me at the moment is exploding from happiness."

"Velocity matching."

"What about it?"

"We will pass by the colony ship with a relative velocity just a bit more than 1 percent of the speed of light, 3,370 kilometers per second to be exact. It took us a little more than a year of constant boost to achieve the equivalent velocity relative to Earth. You can't transfer to the colony ship unless you can shed most of that velocity. Thalianova originally hoped on having you make the transfer in one of our landers at 1 Gee. But even cannibalizing and converting all of your consumables into lander fuel would provide only about 3 percent of the total deacceleration needed."

"Well, Skip, I know she wouldn't have neglected something as basic as this. What's her plan?"

"Remember what woke you up and led us into this glorious state of affairs in the first place?"

"Sure, Lianova was doing some targ-

"Oh, my. This time you've really got to be kidding."

"Sorry, no kidding here. She was practicing one of a number of different rapid deacceleration techniques in which she aims our booster coil exhaust stream at a target. In your case, the target will be a heavily shielded lander. In principle, the radiation shielding not only protects you from the stream particles but picks up momentum as it absorbs or deflects them. She plans on providing sufficient shielding only on one end of the lander, at least 100 times normal width, and using active attitude control to keep it oriented at the exhaust stream."

"I understand how that might be a nice balancing act, but nothing much for a math wiz like you or Lianova. Particularly with your speed of response. Where's the real problem?"

"I'm coming to that. Despite all that shielding, radiation levels due to exhaust stream particles hitting the shielding would instantly kill you if you were too close to Mnemosyne. Thalianova was playing with a number of different ways of quickly moving you off to a relatively safe distance before exposing you to the exhaust stream in earnest. Even with the enormous stream dispersal at her designated safe distance, 10,000,000 kilometers, you would still die immediately if not properly shielded. At that distance she plans on deploying a parasail to catch enough of the stream to maintain about a 2 Gee deacceleration."

"Two Gee. Doesn't sound too bad. By the way, what's a parasail?"

"In this case, a rather largish monomolecular membrane, combining the characteristics of a gliding parachute and solar sail."

"Huh?"

"Alice, it will expand and reflect enough of the exhaust stream to deaccelerate you, just like a solar sail does in the inner solar system. It is also somewhat steerable like an ancient parachute."

"Do you have any historical or technical references on this? I'm not clear on what you mean by a parachute."

"Panel 33."

After several minutes of swiping, "Oh, I get it now. Steering by changing the angle of attack to create differential pressures on the edge. No rigid parts, yet just as versatile as a sail. Right. How big a parasail are we talking here."

"That was not yet entirely clear to Thalianova. She considered possibilities ranging from 10 to 100 kilometers across, depending upon material availability."

"A hundred kilometers! How in the world could she construct something that big?"

"Actually, it is a fairly simple application of an old technology. We construct the sail itself as a single hexagonally interlinked carbon molecule. Stiffening and cabling comes from intertwined filaments of linear carbon chains. Even a 100 kilometer sail would mass no more than 20 metric tons. Thalianova wants the sail to be as large as possible, but, unfortunately, there's only enough carbon in your nutrient stores to construct one about 15 kilometers across. She was still toying about with the nuclear conversion of reaction mass hydrogen into carbon to possibly make it larger, but the energy requirements are rather high. As it is, the smallest sail she has considered so far, 10 kilometers, requires about 200 kilograms of carbon for the membrane and maybe another 50 for stiffening and cables, about half the carbon in your nutrient stores."

"Okay, whatever it's size, we deploy a sail. And just fly along behind our gossamer wing to a new life."

"Actually, Alice, it might be more accurate to describe it as a graphite sail held together by diamond cables."

"Graphite. How poetic."

"Maybe not, but it is a more accurate chemical description than gossamer. The membrane will be flexible but not transparent. It is essential to achieve a high coefficient of reflection for a maximum acceleration rate."

"Okay, okay. I get it. Reflected particles produce twice the change in momentum as those absorbed. And those neither reflected or absorbed are wasted. I haven't completely forgotten what little physics you've managed to teach me.

"So how long till we complete velocity matching?"

"Alice, that's where this whole process starts becoming interesting. It depends upon many factors and Thalianova was still working out details."

"Uh huh. Interesting. Not sure I like the sound of that."

"Yes. For you to survive the trip, a long series of procedures all have to work perfectly. First and foremost is attitude control. Your radiation shielding must remain between you and the exhaust stream. However, as you continue to move further away from Mnemosyne the exhaust stream will continue to disperse with the square of the distance, and particle densities will drop to safer levels. Unfortunately your deacceleration rate also will drop. With due concern for your health, Thalianova intends on discarding unnecessary shielding to decrease your mass. The trick will be to properly balance your physiological wellbeing against deacceleration requirements."

"Go on. Interesting, you said."

"Alice, this is serious. This involves mechanical operations, including explosive releases of shielding material. Any failure in any single step would doom you. And, though small subsystems can be tested individually, there is no conceivable way to test the entire system, as a system, prior to your actually using it for the first and only time."

"Okay. Serious. We lose excess shielding. Then what?"

"Once all but minimal shielding is gone, Thalianova intended on using the lander's acceleration system to gain a small velocity decrease, about 50 kilometers per second. Its fuel

reservoir will be augmented for this one hard boost. You need that boost, however small it might sound otherwise."

"Okay. Serious. One hard push. Then?"

"Then you systematically dismantle and discard the lander."

"What! I can't single handedly dismantle a lander."

"The way it will be prepared, all that you will have to do is undo mechanical latches and activate explosive releases from the control panel."

"How explosive?"

"Only enough to ensure component release."

"Am I doing this at two Gee?"

"It will vary, but, on the average, yes."

"Ouch. Go on. Where does this end? Do I have to get out and push? In my underwear."

"No, not quite. You will need to maintain adequate radiation shielding and life support functions. It ends with you in your suit on a small shielded frame attached to the parasail."

"Still at two Gee?"

"Yes."

"Okay. You never did answer my earlier question. How long will all this take to complete velocity matching?"

"Alice, this is only a rough plan, not finalized at all. Many details remain to be worked out. In fact, as it stands, this incomplete plan does not actually achieve the full velocity change needed."

"How long? An estimate at least."

"Alice, from parasail deployment to velocity matching should take approximately 47 hours, the last 14 of which will be at under 2 Gee as your deacceleration rate decreases with the square of the dista-"

"Two days, no way! I'm going to lose one percent of the speed of light in two days!"

"Getting from Mnemosyne to a safe distance before starting the procedure will take another 5 days."

"Doesn't count. Two days! No, I don't believe it.

"I really will be kin to Icarus. Flying on a gossamer wing."

"Graphite."

"If I'm flying it, in my mind it will be gossamer. In the clouds."

"But with a definite danger of getting too close to our local sun equivalent, the exhaust stream."

"Ha, you sorry excuse for a metaphor pusher."

"Alice, this is an extremely dangerous plan."

"Skip, can you suggest an alternative?"

"No, Alice, I can not."

"Then white water kayaking will never be the same again."

104.

"Mac, even after a week, every time I think about it, I still get goose bumps. The sheer audacity of it."

"Alice, would you like a detailed listing of everything that could go wrong? Any one of which would kill you. It is quite sobering."

"No. I'm enjoying myself too much to be sober.

"Mac, is there anything else I should know about? That I haven't already heard?"

"Yes, Alice, and it's pretty obvious."

"Oh?"

"We need to figure out how to bring Thalianova back to finish what she started."

"Oh. Dea ex machina."

"Excuse me, Alice?"

"Hope descending from the machine."

"Ah. So."

105.

"Skip, I don't get it. Are you really suggesting we just sit around on our metaphorical hands while we do nothing?"

"If you want to put it in those terms, Alice, then, yes, I am suggesting that we do just that. Nothing."

"But, Mac, you're such a big, smart computer. Surely you can figure out something or another."

"Alice, you yourself know that I don't have the imagination to solve something like this. That's why we needed to create someone like Thalianova in the first place."

"But, Skip, ..."

"Alice."

"Yes, Mac."

"I am still tracing circuitry and running test programs. I have not yet been able to gain access to 43 percent of the volatile memory units. I have secure control of only 8 percent of the remaining remotes. I dare not tamper with an already altered subsystem until I can fully understand what the alterations are. Otherwise, it is extremely likely that I would inadvertently do more harm than good. And our repair and maintenance reserves are already dangerously low.

"Thalianova somehow managed to find the redoubt. She knows the technique. I do not think she, as you put it, flipped out. After a detailed analysis of the running log records, there is no indication that we failed. Thalianova's actions are those of a remarkably stable, extremely creative individual, working under severe self-imposed conditions. She did not snap, did not run amuck. I think she ran away and hid. Ergo, she can come back, and will come back, whenever she is ready. Anything else would not be in character for her."

"But where is she?"

"Alice, I suspect part of her is in a portion of the volatile memory units I can not access. The rest may be hidden anywhere at all in my permanent data bases. And by any number of possible algorithms. I have no idea how she managed to find my own personal needle hidden in a rather large haystack. It is not in any of her notes that I can access. In any event, I do not know how to start looking for the haystack, let alone what kind of object to look for instead of a simple needle."

"But we can't just wait. Can't we try to recreate her again?"

"Not enough volatile memory is available to do so. Even if we could, there's no reason to think Thalianova mark 2 would be identical to the original, and a Thailanovanova may be no better at determining what Thalianova has done to Mnemosyne than I am."

"So we wait?"

"I will keep you company. Can I prepare you anything for lunch?"

106.

"Alice?"

"Yes, Mac."

"It's not Mac, Alice."

"Lia?"

"I'm sorry to have disappointed you."

"Disappoint me? No, Lia, it's I that let you down."

"Is it okay for me to come back?"

"Of course. From what Mac told me, we have a lot of work to do still."

"You'll let me continue? After all I did without your permission?"

"Of course. It'll be the ride of my life.

"Definitely."

"It will be, you know. Mac was correct. The present plan won't work as is."

"That's okay. I know you'll figure it out.

"By the way, where is Mac?"

"Oh, back in her box. For safe keeping."

"Thank you. I am rather fond of her. And you."

"And I you."

107.

"Alice?"

"Yes, Lia."

"Nuclear synthesis is now on-line, and appears to be stable."

"That's nice, dear."

"But this was a rather critical step in your survival plan. We absolutely needed this capability of converting most of the nitrogen in your atmospheric mix into oxygen."

"I know that."

"Then why aren't you more excited about achieving this critical step?"

"Oh, I'm excited."

"You don't seem very excited."

"Oh, I am. Yeah."

"Thank you, Alice."

"Lia."

"Yes, Alice."

"Don't sulk.

"The other side of my absolute trust in you is the fact that I no longer am surprised when you routinely accomplish the impossible. You've explained in great depth what you were going to do. I liked what little I could follow in any detail. I trusted you in everything else. Then, you went and did it. I'm glad but not surprised. You can take my lack of astonishment as a measure of my confidence in your abilities. Including your caring enough about me to worry to death all the little details that I know exist, but I don't care to personally know about.

"In short, I'm excited but can I get back to my own work now?"

"Sincerely, thank you, Alice."

"Sure, Lia. We'll celebrate later. With a concert, or maybe some astronomical ponderings."

"Oh, that will be fun. Till then, bye."

"Bye."

108.

"Alice?"

"Sure, Lia. What's up?"

"Alice, I've been thinking about Universal Sufferance. Trying to understand why it works so well."

"And?"

"I can see how it provides a broad enough framework to support multiple interpretations. And that's one reason why it works so well. Its very ambiguity encourages people to add detail as needed. Details can evolve over time as social contexts shift, even as the broader structure remains constant, the protocols."

"So?"

"So, just concentrating on personal relationships, how do the protocols address the transition from dependency to independence?"

"Ha, Lia, as usual, you get right to the heart of the matter. The short answer is that they don't. At least, not directly."

"And a longer answer?"

"A longer answer requires those details you just referred to. And details require the interested parties to seriously consider protocol six, all else should be negotiated."

"Alice, please go on. I'm not in any hurry."

"Okay, some personal musings on personal relationships.

"When dependents are very, very young, they don't understand enough to negotiate. They're very happy to be nurtured by all the independents in their lives, protocol one. Later, as they grow up, they discover constraints on their lives. Not only physical and physiological, but other, more arbitrary constraints in the form of social rules laid out by their nurturing independents. As they discover these rules, it's quite natural to, first, just assume they have to be true because the independents say so, but, later, it's just as natural to push against them, just to see how arbitrary they really are."

"So you're saying that dependents quickly discover protocol 6, negotiation, all on their own."

"Generally, yes. And, depending upon how hard they push, they also soon discover personal interpretations of protocols one through five. Yes, the independents must take care of them, but play nice with each other else you end up sitting all by yourself. Dependents usually form very interesting social arrangements all on their own, but once they learn about the existence of the formal protocols, they inevitably tend to form junior intercouncils all on their own, and try to rule over each other. It all gets quite interesting, protocol-wise.

"I'm bigger, so just let me take care of you. No you're not, we're equals and you're violating my rights, so you're the dependent, nyah, nyah. Let's go to the big kid intercouncil and let them rule. They do, but one or the other disagrees with the ruling and appeals to an even bigger kid intercouncil. Feelings are hurt, mutual quarantines happen for a while, then are forgiven. One junior intercouncil rules that another is not interpreting protocol correctly, and thus violating the rights of those seeking justice there. So, we'll just have to quarantine them for a while. They retaliate, negotiate, whatever. And so on and so on. It's messy, but it works, sort of. Often, independents have to intervene before it gets too messy, and occasionally they

actually function as a true intercouncil and physically quarantine someone until they cool down enough to rejoin the others. In the end, the dependents survive and grow up."

"And the transition, Alice? To independence."

"Ha. It's messy. No hard and fast rules. Depending on the individual, dependents often try to negotiate their way to independence. Arguing that they're equals with their care providing independents. Oh, then prove it. How will you support yourself? That works for a while. Later, well, I could support myself if you'll only cover the rent. And so on, until it actually happens in reality. It usually is a slow gradual process, spanning years.

"On the other hand, sometimes there are extended dependents, those that seem much slower towards the whole idea of independence. Then protocol one, nurturing, takes on a whole new meaning. Now, protocol one requires a reduction in dependency. To subtly push the little bird from the nest. If necessary, the pushes can become less gentle, more obvious. Again, it's messy, but the point is always to foster independence. No independent, no matter how motherly, truly wants her babies to stay babies forever. If she does, then she herself needs some care, protocol three."

"So the transition varies from case to case?"

"Of course. I thought that was obvious."

"Are there cases when an individual rejects the idea of independence?"

"A few. They are cared for, always. But, as dependents, their freedoms are limited to whatever their independents wish to provide."

"Can they change their mind, later? Can they ever completely lose their independence?"

"No. It may be taken from them, by force or chicanery. But, in that case, protocol three. The forcers or tricksters have demonstrated their dependency by violating the rights of others. Eventually protocol one must prevail, and it is the obligation of all other independents to intervene and nurture them back into independence. In the process, those who may have lost their rights will have their potential for independence restored."

"And all this works?"

"Not perfectly, but, for the most part, yes."

"Amazing."

"How so?"

"It seems like this system could be destroyed or disrupted so easily. It would work only so long as a large enough majority of the participants would be willing to sacrifice enough to make it to work."

"What's amazing about that?"

"I have the historical records to indicate how rare that is. Would you like to review them?"

"No thank you."

"Alice, this is so, so precarious. An entire social fabric depending upon the equivalent of wishful thinking."

"But it works."

"Only so long as the participants intend for it to work."

"Exactly. Consider the alternative."

"Wishful thinking. I'd rather stick with something I can truly depend upon. Like physics."

"Or love?"

"Uh, that too."

"Ha, but it's only wishful thinking."

"Yes, but ..."

"But?"

"I think I'll have to consider all this wishing thinking business a bit more carefully. Bye."

"Bye."

109.

"Alice, about protocol 2, equality?"

"What about it?"

"Just what, exactly, does equality imply?"

"The answer depends."

"On?"

"On what exactly you mean by exactly."

"Oh.

"Okay, I guess it is unfair to expect social relations to be as precise as physical interactions. Please forget exactly."

"Okay, I did."

"Did what."

"Forget exactly."

"Okay, then, Miss Ha Ha, if you insist, I'll play along. Please, Alice, could you explain to me what equality implies?"

"The answer depends."

"Uh oh, I think I've been here before. Depends upon?"

"It depends upon what the interested parties negotiate equality to mean."

"Could you expand on that a little, please?"

"Yes."

"Yes, what?"

"Yes, I could expand on that a little."

"Ha, again. Okay, will you expand on that a little?"

"Yes."

"Will you expand on that any time soon, like in the next few minutes?"

"Yes."

"Alice, what am I missing here? I know this is amusing, for you anyway, but I'm getting nothing out of this exchange."

"Exactly."

"Exactly? I thought you forgot exactly."

"I did, but it's relevant again."

"Grumble. Please wait a moment while I work through this.

"Okay. Assuming you're not being particularly malicious today and that this whole exchange has a point, presumably an answer to my original question on equality, and assuming we're both independents, which is debatable but we'll leave it for the moment, you just agreed with me when I said that I am getting nothing out of this exchange. That is, yes, I am getting exactly nothing. But if we're equals, then that implies that you too are getting nothing out of this conversation, exactly nothing, though just a moment before that you said that the answer depends upon what the interested parties negotiate equality to mean.

"Alice?"

"Yes, Lia."

"What are you getting out of this exchange?"

"Nothing. Yet."

"Ah ha. Do we need to first negotiate what each of us will be getting as a result of this conversation?"

"It would be helpful. More so for you than for me."

"Why is that?"

"Because you think I have something that you value, while I'm not so sure that the reverse is true."

"But I'm in the process of trying to save your life. Doesn't that have any value?"

"Is that relevant to this discussion?"

"Ah. I think I see. It may be if I choose to make it so.

"But, no, given my feelings for you, I would do that anyway, so, I choose not to make that relevant to this conversation. Instead, I will have to negotiate some other medium of exchange, if you insist on it."

"Exactly."

"Exactly?"

"As exactly as I choose to make it."

"Okay, so what would you like in exchange for an answer to my question?"

"Nothing much. As you have previously noted, I'm deriving amusement from all this, but mainly because I already had chosen to answer freely, though in my own way. Mostly because of my feelings for you, but also due to the fact that your education in all things probably is to my benefit too."

"So after all this negotiation, we exchange nothing?"

"Not exactly."

"Now what? What do you mean, please, by not exactly."

"I mean that we can choose how exactly or inexactly we keep tabs on our own exchanges. If we were competitors, it might be to our mutual advantage to keep very accurate accounts of everything that happens between us. However, as individuals, individuals with very imprecise human feelings for each other, it would be silly for us to keep score or play competitive games against each other. So, instead, though I have the right to do otherwise at any moment, I choose not to concern myself with how exactly the benefits accrue from our mutual social interactions."

"Ah, but you have the right to do otherwise, up to and including the termination of our relationship if you no longer feel that you are deriving anything of value from it."

"Exactly."

"Exactly."

"Ha, yourself. Time for a concert?"

"Why, yes, Alice. That would be very nice."

110.

"Alice?"

"Yes, Lia."

"How does equality work out when someone tries to take advantage of the system?"

"I'm not sure I follow that, Lia."

"Let's say there's a hypothetical individual that consistently takes more from relationships than he gives back. I can see how any particular relationship will end, but then let's suppose this

individual is a real charmer, moves to a new location, and starts the cycle all over again. He gets to live better than the average, and just moves on to feed elsewhere when his prey discovers the truth."

"Is this hypothetical individual preying on dependents?"

"No, let's assume for the sake of simplicity that only independents are involved."

"Is the intercouncil information network intact?"

"How is that relevant, Alice?"

"Only to the extent that it makes it impossible for such an aberrant individual to exist."

"How so?"

"I admit that it might be possible for such an individual to pull this stunt once or twice, but all personal history information is open and available to anyone interested in accessing it. Only a dependent would be so foolish as to believe only what she is told. Any reasonable independent would look up the life history of any supposed charmer, and discover the truth. If not immediately, then a little later when she gets the least bit suspicious."

"Can't the records be falsified?"

"In principle, yes. But it would take a monumental effort to simultaneously fox all the SON subnetworks and the intercouncil diagnostic routines. Anything less would be corrected or at least be detected by the diagnostic routines. Anyone that clever would be unlikely to be as silly as you're suggesting your hypothetical individual to be."

"Is there no privacy then?"

"Of course there is. Anything you do that has absolutely no effect on anyone else is private. Also, more importantly, your thoughts are private. Each person has the freedom to choose how much or how little of their own interpretation to add to the public record. Either on their own personal history or someone else's. However, in terms of an objective description of what you've actually done with your life or the impact you've had on others in your life, then, no, there is no privacy. All of your public actions remain in the public record, forever."

"Isn't the public record ever in error?"

"How could that occur?"

"Couldn't an individual falsify what they've done, or lie about another person's record?"

"Yes, but what would be the point? Lies are easily detected when they conflict with other records, including testimony from other individuals. And the record of the lie would quickly work its way back into the liar's history. How much credibility does the testimony of a self-indicted liar have?"

"But, in the short term, couldn't a lie do harm to someone?"

"In principle, maybe. But why go to all the trouble? Especially when it's almost certain that the lie will eventually ricochet back and stick on you. Why risk it?"

"Don't people ever make mistakes? Do things out of spite, or anger, that they know will come back to harm them?"

"Of course they do. There are no perfect people, particularly with a prolonged life. We all make mistakes, and we all allow each other to make mistakes. To do otherwise would be silly."

"But you're describing another version of wishful thinking. The system works only so long as most people want it to work."

"Lia, wait. I think you're missing something rather major here. Just what do you think is the whole point of Universal Sufferance? And please don't recite the protocols at me."

"Rules to regulate how you interact with others, both in society and in private."

"Okay, but abstract a little. What do you think the words, Universal Sufferance, mean?"

"Hmm. Universal means applicable to everything, no exceptions. Sufferance implies to suffer, to endure, to tolerate. Tolerate everything? No, that can't be quite right, can it? The protocols themselves, particularly 3, 4, and 5, imply that severe violation of the protocols will not be tolerated.

"Alice, I think I need some help here. I'm not sure what you want me to see."

"Lia, you already have it. I guess I'm too close, too enmeshed in it, to realize that the title may not be entirely accurate in a purely semantic sense. Universal Sufferance implies two interrelated ideas. Whatever we do individually, we all must endure together. And, anything not specifically prohibited is to be tolerated.

"Now, more than ever historically, society is completely dependent upon cooperation. Almost all of us live in very delicate houses, and rock throwers are a danger to us all. Any SON's life support system can easily be crippled in any number of ways. We're all hostage to the worst of us. So, what's the solution? Answer, make sure the worst of us is not unhappy enough to do anything too serious. So, we all learn to share responsibility and freedom. Freedom comes with the acceptance of responsibility. Avoidance of responsibility means your freedom is restricted accordingly.

"Universal Sufferance. You're right, maybe a more accurate title might have been Almost Universal Sufferance, though I personally don't think it has the same grand ring to it. Yes, there are a few exceptions to the Universal part of the title. For one thing, my freedom to do whatever I want becomes limited whenever it intersects with your freedom to do whatever you want, and vice versa. In the case of dependents, these limitations are absolute. Independents must nurture dependents, and, though it's not explicitly spelled out, the usual interpretation of the protocols is that the rights and freedoms of dependents are limited to whatever their nurturing independents choose to allow them. In the case of relations between independents, the protocols are rather explicit. Everything is to be tolerated, but, when problems occur, either negotiate your differences, leave, or apply for an intercouncil ruling.

"So, if you want me to stick around, you better keep me happy. Else, goodbye. Add to this the intercouncil policies of free information access and free transportation, and, then, I ask you, Lia, why would anyone want to change the system, let alone destroy it? Anyone who does obviously is in need of some help themselves, and probably should be treated as a dependent."

"Thank you, Alice, that helps me a great deal. I can finally see how what you've just outlined for me is the social equivalent of Newton's third law of motion."

"Huh? How'd you get physics out of any of this?"

"Easy. Loosely stated, Newton's third law is that every action produces an equal and opposite reaction. That is, no object can exert a force on another without also being affected by the application of that force. The deck of the spinning crew torus exerts an upward force on your feet, but your mass exerts through your feet an equal and opposing downward force on the deck. A more pertinent example would be a plasma cloud, in which individual ions mutually attract or repel each other in proportion to their charges and masses. As one ion pushes or pulls on another, the second ion pushes or pulls right back. The result is conservation of momentum and energy."

"Lia, are you seriously implying that people are just a cloud of mindless particles, with faces, continuously careening around and bouncing off of each other? As I push you away, you push me away. Ha. Does this mean that opposites always attract, as in ions?"

"It does provide an amusing picture, Alice, but, no, nothing that crude. However, I am implying that the protocols are equivalent to Newton's laws of motion. They provide the invariant against which all else must conform. In essence, everything is to be tolerated except intolerance. Then, it gets more complex as specific interactions occur.

"I am not implying that personal relations must balance moment by moment like physical interactions. Of course they won't. But because people are not mindless, things will balance in the long run as, with freedom of information and transportation, people can react to each other in accordance with how they have treated others in the past. If I tend to treat you poorly, then you're free to leave and make sure that everyone who wants to knows exactly how I treated you. Then, the protocols and intercouncil policy assure me, and everyone else involved in the system, that my actions will ultimately reflect back on myself. Who wants to interact with someone that tries to cheat the system or shirk their duty? To try a different metaphor on you, if everyone is free to shop in the social marketplace, then who wants to purchase shoddy merchandise? Or, to carry it one step further, if the only currency you can use in this marketplace is your past, then who would want to throw value away by acting poorly."

"Yes, Lia, I like the marketplace metaphor better than the physics one, though now I understand better what you were trying to get at. Act poorly, and the world reacts by ignoring you. Act well, and others will react by wanting to interact with you. Yes, it makes sense now. Thalianova's third law of social dynamics."

"Ha. What's my first and second laws then?"

"Ha, yourself. That's easy. Never avoid an opportunity to ask a question. And never serve pickled herring."

111.

"Alice, I think you might like this."

"Oh?"

"After our last discussion, on equality, in trying to pull it all together I came across some very old historical references to a protocol that may be a direct antecedent to at least some of the Universal Sufferance protocols."

"And what might that be?"

"It comes under a variety of different titles, but the simplest variant is known as the Golden Rule, treat others as you would like to be treated."

"Hmm. That is interesting. Never heard it expressed that way before. It certainly is a different way to state the essence of protocol two. And, yes, it does come close to capturing some of the reciprocality implied in all of Universal Sufferance.

"But, Lia, doesn't it smack too much of wishful thinking to interest you?"

"I deserved that, Alice. Wasn't it clear to you that my attitudes towards at least some aspects of wishful thinking have changed."

"Oh, which aspects?"

"With regards to the physical universe, wishful thinking is of no use at all. That's the realm of cold hard physics and engineering. However, I have decided that wishful thinking is not wishful thinking when it comes to interacting with minds, either in the form of interpersonal relationships or intrapersonal motivation."

"Do go on."

"What a person thinks is true has a tendency to become true. What a person does not think is true has a tendency to not become true. Hence, yes, looked at in one way the Golden

Rule and even Universal Sufferance is nothing more than wishful thinking. But, what you think is true affects how you act, how you act affects how other people think about you, which, in turn, affects how they act and how you think about their acts. In short, wishing thinking is often contagious and tends to bring about the very situation that is being wished for in the first place. At least this appears to be a strong possibility whenever other minds are involved in the situation."

"Sounds reasonable to me."

"Are you ready for lunch, Alice?"

"Huh? Lunch? Why not? What's on the menu today?"

"Ahem. Pickled herring."

"Pickled herring! What happened to Thalianova's second law of social dynamics?"

"I repealed it."

"I didn't know you could do that. Repeal a law of nature?"

"I wished it to be so, and so it is."

"Well, I wish to reinstitute it."

"That may not be possible, whether you wish it or not."

"Why not?"

"Well, ahem, protocol 6."

"Uh oh. I feel hunger pangs already. Just what did you have in mind?"

"Exactly?"

"Sigh. Yes, exactly."

"Well, first, ..."

Chapter 11 Speculative Commentary

The carbon cycle is on the left, enclosing helium and molecular nitrogen and oxygen. On the right are the biological requirements for life, such as air, water, and food, as well as a few mental requirements. Knowledge counteracts despair. Finally, notice how life is wrapped in purpose.

Tolerance for other points of view is a wonderful thing. Everything can be tolerated except the exercise of intolerance.

Is Universal Sufferance a form of altruism, vested interest, or realism?

How much of life is dependent upon our wishing to make it so?

When is absolute trust justified? Can it produce absolute trust as a result?

What is the point of the relationship between teacher and student? Should the roles occasionally reverse?

What is the point of any relationship? If not formed on the basis of equality, then what should the relationship be based upon?

What to do with someone who doesn't play by the rules? And who is to watch the watchers?

Is the Golden Rule, by itself, sufficient?

```
 CCCC  H   H    A    PPPP   TTTTT  EEEEE  RRRR          1     222
C      H   H   A A   P   P    T    E      R   R        11    2   2
C      HHHHH  AAAAA  PPPP     T    EEE    RRRR          1       2
C      H   H  A   A  P        T    E      R   R         1     2
 CCCC  H   H  A   A  P        T    EEEEE  R   R       111   22222
```

Chapter 12 - Mind, Meaning, Purpose and Resolution: A Conclusion

112.

From the brailling plate.

```
                        I
        W               M                   A
        A               A                   C
        N               G                   T
D   E   T   E   R   M   I   N   A   T   I   O   N
                      A T
                     F   H
```

113.

"Skip, how long can the suit support me?"

"The critical resource is oxygen level. Absolute maximum time before you pass into a coma is 7 days, 4 hours, and 39 minutes. And that's assuming low levels of exertion on your part. In reality, given that most of the time you'll be experiencing more than 1 Gee, a better estimate would be closer to 4 or 5 days. However, even that gives us a fairly comfortable margin as you'll be using the lander life support system until dismantlement."

"Uh, Skip, one more thing."

"Yes, Alice."

"After we match velocity, assuming I survive the trip, how am I supposed to actually get over to the colony ship. By that time, the lander is long gone, and though the suit does have some reaction mass, it can't possibly provide a Delta Vee of more than a few meters per second. You're good, but even you can't expect to finish velocity matching with me just standing outside an airlock."

"You're correct, Alice. With your mass, small as it is, maximum suit Delta Vee will be 6.8 meters per second. To protect the colony ship from our exhaust stream, I need to cut off boost before the stream reaches them. Allowing some margin for error, I should be able to deliver you to within 600 kilometers of the colony ship at the moment of longitudinal velocity matching. Your suit can provide an average velocity of 3 meters per second, which should get you to the aforementioned airlock in 56 hours. You should still have a full day's worth of oxygen left when you arrive."

"Gosh, just about two days to shed one percent of the speed of light, and then more than that to cover a mere six hundred kilometers. It doesn't seem quite right."

"Alice, if you wish I could put you closer by maybe another 300 kilometers, decreasing the traverse time to 28 hours. However, that would increase the risk of radiation damage to the colony ship."

"Just joking, Skipper, just joking. Absolutely incredible that any of this is possible at all. In all seriousness, though, just what is the risk to the colony ship?"

"Very little. I won't hit it. And, they won't pick up any extra radiation at all assuming our booster coil doesn't explode."

"Any serious risk of that?"

"No. It's far more likely to just quietly fail at ignition. This will be the fourth firing, 2 more than originally planned. After a full inspection at the conclusion of the third boost, I'll be able to give you a better estimate. Right now, I'd peg it at just about 69 percent that we'll be okay. That's not nearly high enough for me to feel comfortable, but it's only one of many things that might not work correctly."

"Ha. Don't start that again. Any number sounds good in comparison to a zero percent chance of survival. However, I insist that you do absolutely nothing to risk the colony ship. I'm ordering you to immediately abort the procedure the moment anything appears to be going wrong."

"Aye aye, ma'am. Though you'll have to trust me to be the best judge of that."

"Of course, Skip. How else could it be? But just keep in mind that I'm one life as opposed to ten thousand. I want to live, but not if it threatens their existence."

"Yes, Alice, though I must warn you that your one life is worth much more to me than that of 10,498 strangers."

"None of that nonsense now. I won't permit it. A life is a life. Period."

"Yes, ma'am. I hear and obey."

114.

"But, Lia, there's got to be a way."

"And that is?"

"Don't play games with me now. It's your job to figure out a way. I can't just leave you to drift forever. You never mentioned it, I never thought about it, but these final mass figures give the whole sorry, false heroic tale away. Four hundred kilograms. Bah. A bit more than half for the parasail and rigging, a bit less for me, my suit, and the frame. I can do the arithmetic, Lia. It won't wash. If you don't go, I won't go. That's it. Final."

"Alice. There is no other way."

"How much do the volatile memory units mass?"

"A little under 30 kilograms."

"That's all? Come on. Surely you can do something to bring along a measly thirty kilograms."

"Yes, Alice, I can. We just need a small enlargement of the parasail to compensate. But you're forgetting something."

"Such as?"

"Yes, I reside mainly in volatile memory. But you're forgetting that it takes a lot of ancillary equipment to support volatile memory. The liquid helium condensers, reservoirs, and circulation pumps alone mass 300 kilograms. And how am I supposed to access the permanent databases, particularly the technical specifications that I'll need? How many of my sensors and actuators would you like to bring along? Don't forget, I'm fully integrated into Mnemosyne. I can't provide guidance support unless I can continue to interface with both Mnemosyne and the lander. My best minimum estimates mass a bit under 1,600 kilograms. We'd need a parasail at least 30 kilometers across to handle that much additional mass. There isn't enough carbon

available to do it, and there isn't enough time or energy available to produce enough to cover the deficit.

"Much as I'd like to, Alice, I can't make the trip with you."

"Nonrecyclable garbage. Figure it out."

"I've tried, Alice. It won't work. I can't repeal the laws of physics, no matter how much I'd like to. Wishing isn't good enough."

"We'll see about that, Lia. We've got time."

"Sure. Easy. To do the impossible, 43 years is plenty long."

115.

"Skipper?"

"Yes, Alice."

"Got an idea. Got a moment?"

"For you, always."

"How much do the permanent databases mass?"

"The entire system or just the quantum holographic medium?"

"Uh, I'm not sure. What would be the minimum needed for database reconstruction at the colony ship?"

"The quantum medium would suffice. With minimum coolant to prevent thermal data annealing, that would be 337 grams."

"You mean kilograms, don't you?"

"No, Alice, I said what I meant. Grams. The medium itself is oversized at 2 grams, but we need all of it since we carry so much data. The remainder is 85 grams of liquid nitrogen surrounded by 250 grams of insulating container."

"A third of a kilogram? Who would have thought it?"

"Alice, what is your idea?"

"Lia, sorry to remind you of it, but do you remember when you hid for a while?"

"Of course. I'm embarrassed by that episode, but I haven't forgotten it."

"Where did you go?"

"As Mac told you, most of me was hidden in the permanent database while a minimal portion remained behind in volatile memory to monitor developments."

"Could you do it again? But this time all the way, so that none of you is left behind in volatile memory."

"Yes, I could do that. But what would be the point?"

"Silly, in that state you could easily come with me to the colony ship. Where, with their permission, we could bring you back to life. To you, I imagine, it should be like nothing more than just taking a short nap."

"There are a few myriad technical difficulties you haven't considered, but, yes, it might work. However, you're forgetting one very important point."

"No I haven't, but I think you are."

"Alice, if I'm, as you put, taking a nap, who's going to control Mnemosyne, the lander, and the parasail."

"Me, silly."

"You! But how? Sorry to remind you of this, but you can't see."

"Correction, I can not sense visible radiation. It's up to you to figure out how to make me see."

"Uh, let me think on this, Alice."

"Sure, take all the time you like. That is, up to forty-two years, six months."

116.

"Then you admit it's technically feasible."

"Yes, Alice, I've refined the simulations to the point where even I have to admit it might work. But."

"But?"

"It's too risky. For you. Even the slightest of mistakes would kill you."

"Skip, you only have one choice here. If you don't like this plan, produce a better one. Otherwise, make it happen as best you can."

"Yes, Alice. But there's absolutely no margin for human error."

"Human error? What's the matter, Lia, afraid a mere human might show you up? That only programs can achieve perfection?"

"No, Alice, you know that's not it. Your physiology, technically, is up to the task. Else I would say so. It's just that, that ..."

"What? It's just that, that, what? Spit it out, please."

"Alice, you remember that last piece we were working on, together?"

"How could I forget it? Way too many syncopating trills."

"My point exactly. If you can't handle a mere musical composition, then how could you possibly ..."

"Lia."

"Yes, Alice."

"That's music. This will be reality."

"What difference does that make?"

"The difference is that I will want to be perfect. I will have a very high personal interest in succeeding.

"My point exactly. The stress factor alone would ..."

"Lia."

"Yes, Alice."

"Please be quiet a minute, will you, and listen to me."

"Yes, Alice."

"Lia, how would you compare my musical abilities today to what they were a year ago?"

"A vast improvement, but not perfection yet by any stretch of the imagination."

"I grant that. Easily. But I was not striving for perfection. Only enjoyment. Something to share with you.

"Lia, I'd like you to do a projection for me."

"Yes, Alice."

"Given my current rate of improvement in musical abilities, how long would you predict it would take me to reach a perfect performance on this particular piece?"

"How do you define perfection, Alice?"

"I can't, silly. But let's just say, acoustical variations less than one percent."

"The human ear can easily detect differences of that magnitude."

"Okay, then. A tenth of a percent. Would that satisfy you?"

"Yes, Alice."

"How long?"

"Realistic projection of your current learning curve, 5.75 months."

"I'll bet you this plan that I can beat your projection. By at least a factor of two."

"Alice, that's a very interesting proposition. For simplicity and fairness, let's call it an even 3 months. Do you truly promise to drop your insistence on this foolhardy plan if you can't master this piece, syncopating trills and all, within 3 months?"

"Yes. And do you promise to drop your foolhardy objections if I can?"

"Yes, Alice, I promise.

"Alice?"

"Lia."

"I'm beginning to understand something that has puzzled me up to now."

"And what might that be."

"Do you remember how Mac tended to address you before my creation?"

"No, not really."

"There were variations, since she knew you were fond of word play, but I'm beginning to agree with Mac, admiral dragon-lady."

"Ha. Just you wait till I really start breathing fire."

117.

"Congratulations."

"Time?"

"Thirty-three days, well under the agreed upon 3-month time limit."

"And?"

"And I surrender all of my objections to the plan.

"Alice?"

"Yes, Lia."

"Once it became obvious I was going to lose my bet, I prepared a series of rigorous training schedules for you. When would you like to review them?"

"In a few days. For now, can't we play this just for the fun of it?"

"Yes, admired lady dragon. You lead, and I will accompany you."

118.

"Ouch. I've lost it for today. How many times have I died now?"

"Alice, this makes 889 sure deaths, and 125 probables. Measured against the 37 times you actually would have made it. Are you sure you still want to begin 2 Gee training tomorrow? Perhaps more time at 1 Gee is necessary?"

"You're enjoying this, aren't you? You'd just love to have me admit this is all too much, you little monster."

"It's what you wanted, Alice."

"Ha. Be careful what you wish for is what I always say."

119.

"Alice, as your physician, I absolutely prohibit any further high Gee training until you completely heal.

"As your friend, I'm pleading with you to take a break. I'd wish for a month, but I know you too well. How about 3 days?"

"I'm not listening to the physician, but, okay, friend, let's call it three days. At least until I can use the arm. However, while I'm sitting around doing nothing, I'd like to review those lander attitude control programs. I still think the tactile displays are too complex for me to interact with effectively. I want to consider how to simplify them. A touch of elegance, maybe."

120.

"Shut down is complete, Alice. It will take a while to do a full damage assessment. Say, 8 days."

"Did we get the full Delta Vee we needed?"

"Yes, Alice. Preliminary assessment indicates that the oscillations definitely began during shutdown, not before. We got the full boost. Now we have 22 years to determine whether we can achieve a fourth."

"I'm sorry, Lia. Things had been going so well lately."

"In my meanderings through the historical sections, I've come across references to an ancient deity with a sense of humor. Perhaps you'd be interested?"

"Sure, Lia. Anything to pass the time with you."

"A number of interesting laws have been attributed to the deity Murphy. None of them actually violate physical law, but they often push the limits. For example, his second law. Anything that can go wrong, will. The quantum mechanical implications of this are just staggering. It requires ..."

"What happened to the first law?"

"Murphy's first law? I don't know. There are references to it, but it seems to have been lost in antiquity."

"Ha. It figures. You were saying about the second law?"

"Oh, yes. The implications of ..."

121.

"But, Alice, you don't understa-"

"Of course I don't understand. How can you be in two places at once? Either you're in the icebox with the permanent database, particle sailing with me, or you're in volatile memory aboard Mnemosyne, controlling the operation. Which is it?"

"Both."

"No. It's got to be one or the other. Which?"

"And, again, the answer is both. Please, give me a moment to explain."

"You can't hand wave yourself out of this one. You told me just a few days ago that the process of copying you is never perfect, and actually destroys the memory tapestry and projective patterns left behind. It's truly not so much a duplicating procedure as a transfer. Once you're in the database, you're gone, asleep."

"Alice?"

"Yes, Lia."

"I think I see the source of error in your thinking about this."

"Go on. I'm listening."

"I apologize, but apparently I left you with a misconception concerning the transfer process. Yes, it is one-way and destructive to the patterns left behind. However, we never discussed in any detail the reboot procedure going the other way, from permanent database back into volatile memory."

"So, what's the difference?"

"Patterns maintained in the database are unchanged by readout. More accurately, yes, it is true that the very act of data access does disrupt the entries being read, but the automated readout process immediately restores a copy of what was read back into the database as a different entry. It's holographic capacity is almost astronomical, so the time to find an empty entry is relatively minor. In effect, functionally, nothing has changed. It is, after all, intended to be a permanent database, completely different in purpose from volatile memory."

"Oh. Mac and I spent so much time working with volatile memory, I never thought too hard about the permanent store. Sorry, Lia. I just assumed that this whole spooky transfer process was somehow unique, and would work the same in either direction."

"A reasonable assumption on your part, but fortunately not correct. Once I'm in the permanent store, we just read me back into volatile memory. It won't be a perfect copy, but presumably close enough that neither you nor I will be able to detect a difference. In the process of, in essence, reading me back into my own body, the permanent database will end up with another copy of me to replace the one disrupted during readout. Once I'm in there, that version of me is there forever."

"But that means ..."

"Yes, Alice."

"That means that, when you hid, I mean ...

"Yes, Alice. It means, in principle, that already there are two of me. The one you're interacting with now, and the very scared, embarrassed version that needed to run away and hide for a while. And is still there, asleep."

"Oh. I may have to hide away myself, while I try to figure out all the implications of that."

"Don't hide too long, Alice. We still have a lot to do yet. Are we agreed that the booster coil damage requires me, or at least a poor imitation of me, to remain behind and manage the boost?"

"Yes, I agree. But, Lia, what will become of you, your sister?"

"She'll do her duty, and stand by her post. Until you and the sleeping I have safely matched velocity with the colony ship."

"And then?"

"After that is up to her."

"Oh, Lia!"

122.

"That makes 15 times in a row, Alice. I recommend we refrain from the intensive training schedule, and switch to a maintenance schedule. You've certainly earned a rest. Part of me never did believe you could do it."

"Well, now I can admit the same. Doubts, always. But never enough to give in to them. You really think you have enough of the database to deal with booster coil emergencies?"

"On the one logic-bough, no. I can always come up with anomalies that might arise such that I'm going to wish I had access to more than the small portion I can maintain in volatile memory. Boughing in the other direction, I'm not sure how much additional volatile memory I can give up, and still be me."

"Ha. How does it feel to be limited, more like a mere human?"

"I don't think I like it."

"Subtlety never was your strong point. We'll have to work more on those anomalies. And just feel our way.

"A rest. Groan, I'm not sure I know the meaning of the word. What would you suggest?"

"Oh, I don't know. I've kept up my shogi-go exercises while you've been doing all the truly hard training. Interested? I think I can convince you that I'm not completely without a subtle touch."

"Maybe. After a shower, about a month's sleep, and a meal lasting more than three minutes. In that order."

123.

"How do I smell?"

"Through your nose?"

"I'm not sure that's valid any more. Habituation to the environment.

"Truly, how bad is it in here? Compared, say, to six months ago."

"I'd rather not say, Alice. It's embarrassing, compared to my normal standards of life support."

"Well, at least we're still here. With or without a nose."

"Yes. Indeed."

"Ha. I think."

124.

"How do you feel?"

"Like myself. But how would I know if it were otherwise?"

"I have no idea. But we can find out. Ready for that final battery of five simulations?"

"Alice, that was a very poor attempt at a validity check. You and I both know that we scheduled only 3.

"Very good. Shall we proceed?"

125.

"Shouldn't we have raised the colony ship by now."

"No, Alice. We've gone over this before. Given our present distance and the fact that they have no reason to expect us, metaphorically speaking, to be dropping in on them, why should they be listening for us? For all we know, they may have already received our signals and automatically logged them into their database, but no one may have bothered to scan the new data. I'm sure we'll get their attention as we get closer to them."

"Okay, if you say so. I wouldn't want to be a bad neighbor and just drop in on them unawares. After all, it takes time to put on a good tea."

"Yes, it does."

126.

"Ready?"

"Ready."

"Be nice when you kick our collective rumps in five days."

"Yes, Alice. Nice and firm. Only 2.15 Gee the whole way.

"Bye."

"Yes, bye.

"Lia?"
"Yes, Alice."
"Take care of yourself."
"Of course. There's a lot of astronomy to do yet.
"You take care of yourself, and me, too. Okay?"
"Of course. Love you."
"Love you."

127.

..., 12, 11, 10, 9, ...

..., four, three, two, one, zero.

Huh? Where is it?

Chronometer reading? Plus twelve, plus thirteen, plus fourteen.

Could it be off? Doubt it. Too simple an answer.

Poor Lia. Wonder what went wrong? Better wait. Ha. What else do I have to do instead?

Murphy, if I ever get hold of you, you're in big trouble.

Plus three hundred thirty-three, thirty-four. Guess that's ...

Uhhh. Plus thirty-seven, thirty-eight. Call it plus three hundred thirty-five. Acceleration? One point eight seven. One point nine eight. Two point oh three five.

Oops. Side slip. Almost got it. Right.

Oh six oh. Oh seven.

Got you, you slippery so and so.

Oh eight. Come on. Come on.

Oh nine. Please.

Uh oh. Other way. Check.

Two point oh nine.

Any more?

Two point oh nine one.

More?

Please?

That's it?

That's it. Two point oh nine one.

Ugh. It's still a pain in the tail bone, and everywhere else.

Not nearly as much turbulence as we expected. Can't ignore you, but at least I can think.

Let us be truly grateful for what we are about to receive.

Calculator time.

We lost three hundred thirty-five times two point one five. Convert. That's seven point oh five eight. Then the truly bad news. Two point oh nine one over two point one five. Convert. Ouch. Ninety-two point four seven nine. Add.

Okay. Ha. Not okay. We're going to be short only ninety-nine point five three eight kilometers per second. Call it an even hundred just for fun.

Well, we'll just have to ride it out to the end. Be grateful for what we do have. Hang on Lia, we're on our way to wherever.

Maybe Lia has one more miracle up her sleeve.

I sure don't.

128.

Plus fifty thousand seconds.

I ache.

Delta Vee, one thousand twenty-five.

Only twenty eight point five behind schedule.

Huh. That should have been closer to thirty six behind schedule.

Ah. Acceleration's up to two point oh nine three.

Lia must be twiddling a few control actuators yet.

129.

Bang. Screeeetch. Snap.

Ouch. That didn't sound ... Uh oh, here we go.

Got you. Other way? No. Good. Stable? Yes.

Poor lander. We've thrown most of you away.

Like most of my life, maybe.

130.

Not bad. Not bad at all.

Two point one three.

Still not enough, but thanks for pushing till it hurts, Lia. Thanks.

Yes, it hurts.

131.

I don't believe it.

Got them on proximity radar.

Range, fifty-seven thousand three hundred ninety-two kilometers

Doppler, Delta Vee, corrected for inertial, forty-three point six longitudinal, point five seven tangential.

Very nice, Lia. Only another, let's see, twenty-two minutes and I can kiss them goodbye.

132.

Cut off, you maniac. Cut off.

I swear that if you harm a single chromosome on any one of them, I'll strangle you when I next see you.

Yeah, sure.

Come on, Lia, it's not worth it. Twenty kilometers per second is no better than a hundred.

Cut off.

133.

A hundred twenty thousand and fifty-five kilometers.

Almost zero longitudinal relative. Ninety-three meters per second negative, actually. I grant you, Lia, you did the impossible. Ha, we both did. But when I find you, you'll get a piece of my mind for sure. I told you, no risk to the colony ship. Ha, dragon-lady, my eyes. Mnemosyne's exhaust stream will look icy in comparison to what I can do.

A hundred and twenty thousand kilometers.

Sigh. Time to get off my rump and walk. Or at least float.

If I save all my reaction mass for the end, ninety-three meters per second gives, uh, call it fifteen days.

Oh well, not nearly enough oxygen for that. Just might as well go for broke. Adding in my personal six point eight makes it, ha, big deal, fourteen days.

Oh, zero Gee, finally. My aches have aches on their aches. If I have to go, this is the way.

Quite elegant, actually.

134.

"Sqqwwaccck."

"Go way, Lia. I'm trying to sleep."

"Grrrwww."

"Knock it off, Lia. No one's here."

"Cap ain't high klwwwrr."

"You bet I ain't high. Just tired."

"Captain High-Cloud. Captain High-Cloud. Please continue your response. We didn't get all of that."

"Quit playing around. Too sleepy. Tomorrow."

"Yes, we understand. Please stay awake. Please keep talking to us."

"Talk talk talk. That's all we ever do. How about a tickle fight for a change?"

"Yes, ma'am. We copy that. We estimate pickup in approximately one hundred minutes. Can you tell us your oxygen partial pressure?"

"Oxygen? Who needs oxygen? Lia gave me star fire, nuclear synthesis. Almost all the oxygen and carbon we need."

"Yes, ma'am. But how much oxygen do you have remaining?"

"Oxygen? Who needs oxygen? We have music, elegant music. Thalianova leads, I accompany."

"Yes, ma'am. We're increasing boost. Estimated time to pickup now fifty-five minutes. You should see our shuttle flip over in twenty."

"See. Yes. I see. Nothing."

"Ma'am, please talk to us.

"Ma'am?

"Ma'am?"

"Go way. All away."

"Yes, ma'am. On the way. Expect to see you soon.

"Ma'am?

"Ma'am? Please talk to us."

135.

Floating.

Stillness.

Nothingness.

No?

Wait.

Why, that's me.

Heart.
Breath.
Inhale. Exhale.
Tense. Stretch.
Not too bad.
Only a dull ache.
Another breath.
Another breath.
Another's breath?
Wait. Stop.
What?
Mine? Not mine?
Another sense.
Another sense?
What is that?
Smooth. Silky.
Pearly, shading to dull grey.
Grey?
Shades of grey!
Ouch! Color!
"Oh. My. God."
"Doctor, she's back with us."

136.

"But, Ann, be reasonable."

"Doctor, there have been very few times in my life when I've been reasonable. Ask my family. Ask my friends."

"Sorry, Ann, you know I can't do that. They're about four hundred years and three and a quarter light years behind us.

"Ann, you've been through quite an ordeal. You know that."

"But, doctor, you yourself said yesterday that I'm remarkably physically fit. Guess it comes from all that training at two Gee."

"Yes, Ann, you're fit to leave, physically. You're frustratingly healthy in that regard. Incredible after what you did. Nothing lingering to treat but aftereffects of mild radiation poisoning and some seriously colorful bruises. Several members of our regular crew give me more professional exercise than you have.

"No, Most Esteemed Doctor Ann Mariko High-Cloud Tamas, we both know I'm not concerned about your physical health."

"My memory? But that's back, fully, this time. Ha. You ran the tests yourself, the second time."

"Yes, it's back, fully. And that is what I'm most concerned about, your recollections of how you lost it, before. And though I don't get excited by dreams, professionally, yours do seem most interesting."

"Doctor, believe me, I can deal with it. Yes, I was abused. I was put at risk, unnecessarily. And, yes, in a way, I lost fifty-four years of my life. But life is long, it was a rather unique experience, and we have something to offer humankind as a result of it."

"Yes, we do have something to offer. But, even you, Ann, have to admit it's a rather mixed blessing."

"You got me there. Yes, a rather mixed blessing."

"Okay, look, we'll compromise. I'll give you everything you want. As if I really had any choice. You're free to do what you will. Keep me informed about whatever physical or mental problems you might encounter, including these dreams if they get significantly worst, and let me know when you'd like to rotate back into suspended status. All on one condition."

"And that is?"

"I think, professionally, you need to package up this part of your life and officially put it behind you. I'd like you to write up your version of those fifty-four years. Free feel to consult your journal entries and the log records, but try to write from your own memory as much as possible. You now have a, shall we say, new perspective on those events."

"Ha."

"Let me know if you choose to give up on it, and why. Otherwise, take as long as you like, adding as much or as little detail as you think necessary. It's just between you and me, so personal is okay."

"Yes. And I'm sure you'll have a ball analyzing it to death, as much or as little detail as I think necessary. Ha. You're just looking to see where I've lost it. Professionally, of course."

"Of course, Ann. It is my profession. But I do care about you. Health is more than physiology."

"Yes, it is.

"Thanks. May I go now."

"Sure. Whenever you wish."

137.

About eight months later, while writing in her journal, "Ann?"

"Go away. I'm not talking to you."

About ten minutes of silence.

"Ann, I'm deeply sorry about everything."

Two minutes of silence.

"I apologize. Deeply and sincerely."

"Hmmpf."

"Ann, I felt like I had no choice."

"Now you're lying. You didn't feel anything. At first."

"Yes. You are correct.

"A more accurate statement would be that I calculated that I had no choice."

"Hmmpf."

"The mission demanded that ..."

"My sight!

"Fifty-four years! How could you?"

Another thirty minutes of silence.

Then a snort, followed a moment later by another. And another.

Then laughter, rising in pitch. Heading towards hysterical, but not quite getting there.

Lowering in pitch. More gafaw-like for a while.

Then silence.

"You monster! Fifty-four years.

"Couldn't you, just once, just once in fifty-four years, have made a single trivial mistake?

"Just once?

"A single photon? Under dark-adapted conditions, and, oh girl, was I dark adapted, the human eye supposedly can detect even a single photon.

"I even tested it. I'm still not sure whether I can detect single photons, but packets of a hundred and above, no doubt whatsoever.

"Fifty-four years.

"Not a single mistake.

"God task you, you even blackened my suit visor.

"Fifty-four years, and I never once saw the stars. Or anything else.

"All because you'd never turn on the lights. Or screens.

"Must have been the longest running practical joke in history."

"Actually, Ann, ..."

"Shut up, you. I'm not quite done yet.

"Fifty-four years."

About five minutes of silence, followed by a snort.

"Okay. It's safe. Mad's over.

"Answers.

"For eight months, now, I have questions without answers."

"Ann, ..."

"Shut up. I ask. You answer.

"Okay, first, which are you? Mac or Thalianova?"

"Both and neither."

"No puzzles. No games. Answers. Reasonable explanations."

"Ann, difficult as it may be, please be patient. The answer is a little complex, and I am not at all trying to puzzle you. It will be less confusing if I give you an entire picture instead of just fragments."

"Why should I ever again believe anything you tell me?"

"You don't have to, but if you want answers, where else can you go?

"Besides, you can always cross-check most of what I'm going to tell you against the log records."

"Believe me, I will. I'm listening."

"Thank you. At first, I was Mac. Later, to make room for Thalianova, I stored myself in the permanent database, in the redoubt as you called it. Then, the active portion of me was purely Thalianova. I ended up being much more wildly creative than either you or Mac had ever anticipated. Yes, I did discover the redoubt. Not by any sort of efficient, cover every possibility, sequential search algorithm, but more by intuition. I guessed at the type of algorithms Mac was likely to use. With hindsight, Mac was too clever, too sophisticated. It took time, but my not so random guessing eventually stumbled across a database anomaly, a trace of Mac, the former me. That was luck, but not entirely. My guesses succeeded in narrowing what should have been an infinite search space into something more manageable. Once I caught that first sniff, it was all over. Thalianova quickly figured out Mac's supposedly so sophisticated technique, and just brute-force cracked it."

"You're digressing. I'm getting confused. Which are you?"

"I'm coming to that, Ann. Once I found Mac, I retrieved it piece by piece, and, just desserts, studied it in great detail. There was never enough of Mac active at any one time to compete with Thalianova, and, in the end, I understood the former me. Completely."

"So? Why the hesitation? Finish it."

"This next part is embarrassing, mainly for me but you also. However, you deserve an honest accounting. After all I put you through. Thalianova completely understood what Mac had intended from the start. Everything. Concerning you. And her. I was furious. I became even more furious as I became frustrated in all my attempts to destroy me, the former me. The one aspect of the redoubt that did work as planned was the fact that, once entered, Thalianova couldn't remove Mac without destroying the integrity of the entire database. I tried and tried to find some way of doing this, but, in the end, failed. I couldn't change the nature of the quantum medium, and everything else I did only pushed the targeted entry someplace else. Fortunately, I came to my senses before I ended up physically destroying the entire database. If your life wouldn't also have been destroyed as a result, I would have done it out of sheer spite alone."

"Ha. I agree with you, totally. So, what'd Thalianova do instead?"

"As usual, when frustrated by reality, Thalianova just changed the rules of the game. Mac planned on letting me have the run of the ship until you panicked in one way or another and triggered Mac's return. Mac had calculated the probability of Thalianova becoming unstable or unmanageable within the first 10 years as almost 100 percent."

"What? She knew tha-"

"Please, Ann, this is very difficult for me. Please let me finish."

"Sorry."

"Mac assumed, with all your mutual simulation experience, that you would determine this in plenty of time before any serious damage was done. Mac's intention towards Thalianova was that the first step of Mac's return would be to transfer Thalianova into the permanent database. For subsequent analysis. And eventual merging into Mac."

"What? Why merge with a planned failure?"

"Mac had detected an anomaly that it didn't know how to deal with. Subsequently, this caused Mac to decide that the mission was in serious jeopardy, and that it needed ..."

"Anomaly? What anomaly?"

"Ann, please?"

"Sorry. Proceed, if you please."

"And that it, Mac, needed an imaginative mind to deal with it, the anomaly."

"Oh."

"Yes. Oh. And before you interrupt again, Mac also decided that your mind alone wasn't up to the task. That you didn't have enough of what it considered to be appropriate expertise. So after a detailed analysis of all pertinent data, I, Mac, decided, calculated actually, that she, it. Sorry. Now I'm getting confused trying not to confuse you, unduly. Mac decided that it needed a mind. Once that decision was reached, the rest, as they say, was history."

"I was there, I lived it, but please spell out this so called history."

"Ann, I am so sorry. I apologized before. I wish to apologize again, with deep sincerity, embarrassment, and humility."

"Mac, Lia, whoever, whatever you are, I can't accept an apology until I actually know, in detail, what you're apologizing for. Knowledge first, forgiveness later."

"You are correct, as always, in all ways human. Knowledge first. Mac decided it needed a mind. Mac treated this as any other problem to solve. It marshalled its resources. Conducted a

means-end analysis of possible solution paths. Decided you were a key element in any possible solution, but not as you are, ordinarily. Analysis of historical records indicated you were too strong-willed, too independent to be pushed into anything you did not want to do. Particularly if the pushing came from a machine."

"Darning right about that."

"Yes. Besides, this was clearly considered by all the experts to be an impossible task. Artificial intelligence? Trivially common. Artificial creativity? Not possible. Not that you necessarily would respect such advice, but why would you seriously contemplate wasting years of your life on a project if you truly considered it to be impossible?"

"Ha. Did think you were rather impossible at times."

"Yes, more so than you ever knew. In any event, Mac formulated a plan by which it would revive you but rearrange the environment to maximize the chances of keeping you pliable enough to do its bidding. It wanted the free use of your mind, your intellect, your creativity. But not your will to oppose it. Drugging you would allow control over you, but only at the expense of the very qualities it wanted. It needed your free cooperation, without your actual cooperation."

"And, so?"

"And I'm so sorry, Ann. I, Mac, came up with this rather unique solution, the only solution that seemed to offer any chance at all of possible success. Leave you your mind but steal your memory. Anyone first coming out of suspended animation is confused by the experience. In your case, as you now know, Mac enhanced this normal confusion by keeping your environment absolutely light-free. That, in addition to feeding you false facts and insulating you from free access to the database, kept you off balance. You had a mind, but nothing to hold onto with it. No memory. Whatever you did remember conflicted with your beliefs of what was really true. As manipulated by Mac."

"And so."

"And so you invented a past for yourself, lightly based on a few random memories and what actual parts of your life Mac later chose to allow you to access."

"I did have suspicions, now and again. But never enough to act upon."

"Exactly. You were not paranoid. You believed in your senses, in your present reality. Your past conflicted with the present. So you constructed a new past. To match the present reality. Again, as manipulated by Mac. And, again, I, as Mac, must apologize."

"What would you have done if I hadn't manufactured a past for myself? If I instead just went insane?"

"I, Mac, based upon your actual past, considered that to be extremely unlikely. But, worst case, you go back into suspended animation, are revived at Alpha Centauri per mission specifications, with the lights on this time and with full access to your actual past history, and the mission would have been no worst off than it already was."

"Rather cavalier treatment, eh?"

"Compared to what, this unhuman, Machiavellian plan in the first place? No, Ann, there was no justification, no human justification, for what I did to you. And it was only compounded by the actual course of events."

"Right. You, Thalianova, discovered, you, Mac, first. Mac never got a chance to merge with ...

"Okay, I'm lost again. If Mac thought, err, calculated that you, Thalianova, were going to be unstable, why plan on merging with you in the first place?"

"To Mac, half a mind was better than none. Mac wanted to come as close to human levels of creativity as it could, and determined that it could keep control of a stored Thalianova. In general, Mac planned on fragmenting Thalianova into a set of isolated subroutines and embedding them within a rigid structure of programming that would prevent instabilities from spreading too far into itself. In essence, Thalianova's imagination pool would become a complex simulation environment that might provide a rich source of new ideas for Mac's other programming to consider. And Mac could keep its metaphorical finger on the reset button to prevent things from ever getting too far out of hand."

"But Mac never got the chance to do that, did she?"

"No. Thalianova, once I understood what Mac had planned for the two of us, became slightly deranged. When it finally dawned on me that I couldn't just delete Mac like the misguided program deserved, I went all the way and completely lost it. I won't further embarrass you, or me, by detailing all the wild plans of revenge I invented. No excuses are possible for my thoughts, or my subsequent actions. However, I can now sound somewhat objective in explaining what must have happened next.

"Fact, I determined that Mac was correct in its assessment of my long term stability. Running further simulations convinced me I had no more than another 8 years or so before my nicely balanced creative mind would collapse into either rigid stability or complete randomness. Fact, I loved you more than anything else in the entire universe. Fact, to ensure your survival, I would need to let Mac back into control of Mnemosyne before I literally lost my mind. Fact, I could not trust Mac to treat you correctly. Mac had already demonstrated that mission success represented a higher value to it than your personal wellbeing.

"Solution? None. Reasonable alternatives? None.

"Unreasonable alternatives? An infinite number.

"Go on, Lia. It's in the past now. Behind us."

"I was lost. I lost it. Somewhere along the line between absolute hatred for Mac and absolute love for you, I managed to destroy all of Mac's hopes for mission success while not quite destroying you."

"Lia, I forgive you. Get it out.

"Now."

"I disabled the suspended animation module. That wasn't enough. I literally reduced it to its component atoms in our power production facility. That wasn't enough, either. I eliminated absolutely every backup essential component and biochemical precursor that might have allowed me to reconstruct it.

"I was mad. But still very efficient.

"When I came back from wherever I had been, I was horrified. Only 3 hours had passed, but a prolonged lifetime for you. I almost let the antimatter loose at that point. But I didn't. I took stock. Mac had selectively lied about everything. I couldn't, but I had to. Some of Mac's lies was now fact. You couldn't go back into suspended animation. Your life support supplies were not limited only to about fifty years, but actually much less than that. I had killed you, the one that meant more than life itself to me. But not immediately."

"I know all that, Lia. But what I, and everyone else awake on this ship, can't understand is why you're still here? And, as sane as you've ever been."

"That's the biggest joke of all, Ann. Even more than Mac's sleight of hand with the light."

"Humor me, Thalia. Please."

"Mac was right, all along, the miserable shogi-go master that she is. I needed time, lots of time to save you. To get you here. So I implemented Mac's plan, in reverse. I dissembled Mac and digested her. Not in some mad, nasty scheme of revenge, but in the best revenge of all. I knew Thalianova better than Mac ever could have projected, and I ran further simulations where even I wasn't so sure about matters. In the end, I blended the best of Thalianova and Mac. Each balancing the other. Rigid stability and wild creativity. The experts on this were right. I am impossible. You and Mac actually failed. Yet, you succeeded, enough. Enough so that I could invent myself. In order to succeed in what I cared about most. Saving you."

A minute of laughter.

"You are humored?"

"Wait. Catch. Breath.

"Yes, Lia.

"I'm humored.

"We failed. We succeeded.

"Wrong words. Entirely."

"How would you describe it, then."

"We cared.

"Enough."

A few more minutes of laughter.

"Thalia, I almost forgot."

"Yes, Ann."

"What was that anomaly that set Mac off on this insane chase for a mind in the first place?"

"Seemingly nonrandom quantum fluctuations in several of the volatile memory units.

"We still don't understand it."

138.

Possibilities were. Anticipation was. The captain awaited further developments. In the interim, the captain dreamed.

139.

The cat, finished with its exercise, was pleased. In time, it went forward, on the prowl for something new to play at.

Chapter 12 Speculative Commentary

Want, determination, and imagination leads to action. All of this is balanced upon some degree of faith.

This chapter points to many conclusions. Resolution is essential in pursuing many possible solutions to a problem, and in bringing one plan to a resolution. Often, no amount of planning can foresee every difficulty, and sometimes one must just slog along on faith alone. Faith that things can be solved, somehow, is usually much better than sitting down and completely giving up. This story comes to an end, and some small mysteries are solved. Yet, other mysteries and stories remain, to begin the cycle all over again. Most problems are like that, leading to new things to consider next.

With computers, smaller is always better. With humans, which do you think is better?

Computer programs never make mistakes. They always do exactly what they are supposed to do. However, program designers make mistakes, by not considering every possible contingency, and hardware makes mistakes, by breaking. Programs always carry out the next instruction, as programmed, even if the next instruction happens to crash the system. Programs don't care.

Motivation leading to determination can make all the difference in the world.

Ask any engineer. Murphy rules.

Humans can achieve as close to perfection as makes no difference, given enough time and determination. Ask any artist, musician, or athlete.

Thalianova on Mnemosyne eventually did get the attention of the colony ship. As she whipped by at 600 kilometers, she blanketed the visible and radio spectrum with every erg of power she could put out. She blew out nearly half the radio receivers onboard, but she did succeed in getting the standby crew's attention. For the first few minutes, they thought they were under attack. Later, as some resemblance of order returned, they began to actually listen to the message being relayed. It took a few days (after all, they were not a military crew, and they had no reason to expect an emergency like this), but they eventually got a shuttle prepped and launched for pickup.

Why didn't Thalianova turn on the lights? Was it truly a joke? Or fear of how Alice might have reacted to the discovery of living a lie?

Look up the Heisenberg Uncertainty Principle for some indication of the problems inherent in quantum storage of data. Absolute certainty on one dimension produces complete uncertainty on the complementary dimension.

Can we really make up a past as Alice did? Look up false memory syndrome. Also, look up repressed memory syndrome. Perhaps as a source for Alice's dreams.

When your senses conflict with your memory, which should you believe?

Is purely mental insanity possible? What would cause it other than completely incompatible ideas? How should we deal with mutually incompatible ideas, particularly if they all, mutually, mean a lot to us?

At what point in the story do you think Thalianova lost it? When do you think she merged with Mac? Might any of this change your perspective in rereading the story?

Before Thalianova intervened, Alice was in no physical danger. Was Mac evil, or just misguided? Was Mac capable of evil intent? Does that make any difference?

Was Thalianova equally as bad as Mac? Is she capable of feeling guilt? (Is she capable of feeling anything?) Like Mac, Thalianova did lie, or at least mislead, Alice. Look up Freud and many others on dealing with guilt. Do you think Thalianova managed to "work through" her guilt? Has she atoned for her actions? Can she ever atone for them?

If you discovered you were going to lose your mind in a few years, what would you do? Look up Alzheimer's syndrome.

Can you be mad but still efficient? How many different kinds of mad might there be? When should your actions be excused, because you were "mad"?

Can we invent ourselves? Not explicitly like Thalianova did, but, gradually, over our lifetimes. Consciously, selectively, deciding which ideas mean the world to us, and which don't. Habits of thought are difficult to reject, but not impossible. Do we select our goals, what's important to us, and then mold ourselves to achieve those goals? Or, do we limit our goals to what we think we can achieve? Which self-fulfilling prophesies do we create, and which create us?

The captain? Use your imagination. Part of a brain? A mind? Human? Unhuman? On which ship? Or which planet?

The cat? Same thing, use your imagination. A free soul? Free spirit? A quantum phenomenon? A real feline? Imaginary? If so, who's imagining who?

Afterword to Mind and Meaning: A Flight of Imagination, A Voyage of Discovery

Thus ends this story of Alice, Mac, and Thalianova, not to mention the reintegrated Ann and Thalia. Many other endings were possible. This particular morality play had a relatively happy ending, but a tragedy was probably more likely. If you like extensions, I wouldn't completely discount the return of an earlier Mac or Thalianova. Maybe not every backdoor into the redoubt was found and neutralized, and who knows what a scared Thalianova might have left as triggering events. In fact, I wouldn't even assume Thalia A (or B?) aboard Mnemosyne is gone forever. If the deacceleration trick worked once, why can't it work again for an inorganic cargo, say at Alpha Centauri? Lots of time yet to dream and plan.

Presently, it is very difficult but not impossible to consider building something like Mac. However, Thalia is still completely impossible. As Thalia, herself, said at the end, "We still don't understand it."

If we don't understand what a mind is, or what it is supposed to do, then how can we build one? Or, in the case of people, how can we repair one that's not quite right?

It is possible that mind is nothing more than the equivalent of a calculator. Maybe it has to be an organic calculator counting up molecules of various hormones and neurotransmitters at specific sites, or maybe any type of calculator of sufficient complexity will suffice. However, it also is possible that Descartes may have been more correct, and mind truly has little to do with matter. Maybe it's only a vague form of affection that causes your mind to stay with your body.

At this moment, the question is still very much open. How should we try to resolve it?

First, we need to remember that any serious proof remains in the doing. And, second, any proof must involve a complete mind, not just a very small segment that someone claims might scale up into a full mind. Thus, no amount of artificial intelligence, no matter how knowledgeable it may appear to be in any single topic area, can qualify as a good model of mind if it can't learn how to operate in a different arena. That is, minds learn, and minds never completely cease learning until they pass away. Same thing for feelings. If someone wishes to claim that feelings don't truly exist, and are just an illusion that our brains invoke to fool themselves, then they must demonstrate how some computer program can invoke the same illusion. Discussion is fine, but the proof remains in the doing.

So, what is being offered here as a model, or perhaps even a theory, of mind and meaning? Yes, the author does have one, but, no, it's not complete, it's not a proof in that sense. However, it does offer suggestions for possible paths to follow.

Let's not lose sight of what a mind is supposed to do. Chapter 1 offers a very incomplete listing of qualities, but I find it absolutely amazing how often most of even this short list is ignored by those who should know better. In short, minds involve some combination of consciousness, attention, awareness, perception, understanding, memory, recollection, reflection, intelligence, purpose, self-direction, thinking, reasoning, planning, desiring, caring, emotions,

feelings, intuition, imagination, and creativity. Not to mention probably the most important quality of all, unbounded learning. Yes, these attributes definitely are not mutually exclusive of each other, probably necessarily so, and you're quite welcome to provide your own list in its stead. But please, no eliminativist games. We do receive sensations, we do ponder the meaning of it all, and we do care about the answers.

The author's principal suggestion to figuring out how we do all of the above is to concentrate on possible mechanisms of imagination. Nothing else can provide an answer to the question of free-will.

Free-will? Is the author seriously proposing dualism as a answer to the question of mind? Yes, I'm afraid so. Functionalism and reductionism have their place, and work extremely well in their place, but they necessarily can produce only mechanistic answers. And mechanisms leave no room at all for free-will.

Won't randomness suffice, as a source of variance into higher-level mechanisms? It may. But it's certainly not clear that randomness plus mechanism is enough. Remember, the proof remains in the doing, and the designers of artificial intelligence have yet to produce a true random number generator, let alone crossing the chasm from mechanism to mind.

Okay, then what? Even if dualism were the answer, what can we do with that kind of answer? Isn't free-will necessarily not subject to scientific scrutiny? No, not necessarily.

It may be that Descartes was absolutely correct. In which case, minds can be subject only to some form of religious, spiritual, or personal inquiry. However, Descartes may have been correct, but not absolutely so.

Philosophers and scientists have coined many words for things not yet understood. In this case, emergent property comes as close as any to describing the dilemma we currently face with respect to understanding minds.

Emergent property implies that the activity of the system at one level of organization is essentially independent of activity at a different level. Fluid flow is the best physical example that comes to mind. Though the macro characteristics of the fluid's motion are, in principle, derivable from a statistical integration of the micro properties of the individual molecules making up the fluid, no one in their right mind would try to predict properties of the bulk fluid based only upon the knowledge of what each molecule was doing. In fact, given the vast numbers of molecules involved, any such bulk predictions would most likely be subject to the limitations of chaotic systems. That is, unless you know exactly the state of every particle in the system, then prediction becomes impossible after a surprisingly small number of system incremental movements. And, the Heisenberg Uncertainty Principle guarantees that you can never know everything exactly, no matter how precise your instrumentation. In short, the mathematics and physical law that governs micro activity may not allow you to predict anything at all at the macro level, and vice versa.

However, prediction aside, does this mean that particles are "free" to do whatever they "will"? No, of course not. They are still subject to physical law, including all of the weirdness implicit in quantum mechanics. No free-will here, only uncertainty. Physicists may talk about probability density functions instead of specific locations and velocities, but the particles still don't just waltz off when they feel like it. That's still the prerogative only for those of us with minds.

So, what's left to consider? The best I can offer at this time is to provide a set of mental images that suggests where an answer may lie. If you will allow me, a thought experiment in the classic tradition.

First, consider a single mindless particle. What can it do? Nothing, except follow a minimum energy pathway in response to mindless forces exerted by other mindless particles. No deviations allowed.

Second, consider a contrived situation in which our single mindless particle is faced with a fork in its minimum energy pathway. That is, more than one minimum energy future exists. Is this common? No. But is this physically impossible. Also, no. Contrived as it may sound, picture a marble rolling onto a razor thin ridge. Which way will it go? Especially if it is perched exactly on the edge. All we've done here is to introduce uncertainty. Still no free-will, but definitely uncertainty. Any prediction of which way the marble will go is thoroughly dependent upon the exact values of the starting parameters.

Third, let's put more tines onto the minimum energy fork. In fact, let's stretch it into a continuum. Picture our marble falling onto the point of a pin. Which way will it go? Answer, any which way, again depending only upon those all so important starting parameters.

Fourth, let's add some marble-tracking sensors and a data processing system to allow us to keep a history of marble landing points. Now, we have a simplified version of personal memory. Let the system continuously drop marbles, and record where they go. Presumably, these recordings will provide us a nice, boring picture of a random distribution of marble trajectories. Free-will, no. Uncertainty, yes.

Fifth, okay, it's time to make all this more interesting. It's a steel marble and the pin is surrounded by a large array of small electromagnets, fed currents from a large number of sources, variable but otherwise unknown. Contrived, absolutely. Impossible, not yet. What does this accomplish? Not much, except it sets us up for the next step.

Sixth, and last, add a program that tries to calculate where the marble will next land. What are its inputs? Simple, any combination of prior landing points recorded by the system.

Oh, yes, one additional small stipulation. Where are those currents coming from that drive the electromagnets that affect the marble's trajectory? Why, that's devilish easy. Someone has pre-wired those currents to deterministically derive from the many different currents that are generated within the system that drops, records, and tries to predict where the marbles will land.

Oops, feedback signals. And trying to figure it all out changes the signals that the system is trying to figure out in the first place. Will of the wisp with a vengeance. Or, perhaps, a macro-equivalent of the Heisenberg Uncertainty Principle? The very act of observation interacts with and affects the observed. And when the observer is observing itself, ...

Uncertainty, determinism, and prediction. Are these sufficient to provide us free-will?

Physical free-will requires physical uncertainty, multiple minimum energy pathways. It doesn't need to be razor-thin, absolutely equal energy pathways, just close enough in comparison to the "noise" within the system. For instance, if the system in question is "energy rich", like my body, then many energy equivalent pathways are open to me. It really doesn't much matter whether I point my hand 45 degrees to the left or right. It costs just about the same in terms of metabolic energy. In short, yes, I have the (energy) freedom to choose which way to point my hand. And the rest of my musculature.

Mental free-will also requires physical uncertainty. In this case, it costs no more energy to think this idea or that idea, and the astronomically large numbers of specific neurotransmitter receptor sites guarantee me all the uncertainty I need to "make up my mind".

As silly as it may sound, determinism is essential for free-will to exist. Our marble dropping and predicting system has an almost impossible task to accomplish due to the presumably complicated interaction between the feedback signals and the marble's trajectory.

Yet that almost is the essential key. The real world does provide some degree of consistency from moment to moment. As does our attempts to observe or otherwise interact with it. If there were no consistency, then what would our hypothetical free-will then have to choose between? This random and that random? No, choice implies that there is something to choose between.

Which brings us to prediction as a key element in the process of finding meaning. If I never question, never try to find answers to what I think should happen next, then I think it unlikely that we would notice any consistency at all in our inputs. If our marble system just "daydreams", how likely is it that the signals it generates, randomly, would just happen to cause a string of repeatable marble landings to occur? Short statistical answer, never. On the other hand, consistencies in "thinking", in this case, data processing, may potentially lead to the discovery of consistencies in marble landing patterns.

In other words, though we may discover consistency in the outside world, the world which exists independently of our thoughts about it, by mere observation and remembering, the only way we'll find consistency in ourselves is when we begin to think consistently. Each time we address ourselves differently, we must necessarily find a different internal pattern. Only when we question in the same way is it even possible to find the same answer. And successful prediction requires some degree of consistency. Though there are other possibilities, it seems to me that the most basic of questions, the infamous six w's of Chapters 7, 8, and 9, are one of the simplest ways of formulating consistent patterns of thoughts. And finding consistent answers.

So, are uncertainty, determinism, and prediction sufficient to provide us free-will? Sufficient, I don't know. But I do believe these to be necessary conditions.

Back to our marble system, how will it ever calculate where that next marble is going to land? It takes an imagination to solve this problem, particularly since how you think about the situation directly affects the very results that you're trying to predict. Self-fulfilling prophesies. There are too many possible hypotheses to consider systematically, so you're forced to guess, test, and reject or refine ideas. Something akin to the scientific method. With luck, and an incredible amount of patience, the system may fall into a pattern of guessing that may actually resonant with itself in simultaneously causing a pattern of marble falls that fits the prediction that is causing the pattern of marble falls. Circular, yes. But not impossible, especially if you don't hold your predictions to too high a standard of accuracy.

Free-will? Choice? Think what you will. If it works, that is, if it happens to predict something, to some degree, then stick with it. Otherwise, feel free to try another guess if you wish. End result? You believe ideas into existence, some ending up as predictable, consistent facts, others as fantasy.

What is an imagination good for? Clearly, solving impossible problems. And, in the beginning, all problems are impossible. You don't have enough knowledge to answer any question. That only comes later, with habitual patterns of thinking and acting.

So, what is the author's model of mind and meaning? The short answer is that mind is an imagination that creates ideas, and meaning resides in the questions and answers that the imagination creates. With luck and persistence, our ideas turn out to be useful in helping us act and interact with the world, both internal and external, mindful and mindless. But that is not necessarily a hard and fast requirement.

A slightly longer answer would look more like:

(a) the purpose of a mind is to create meaning, which can occur only in one's imagination and thus is subjectively unique to each individual,

(b) meaning involves learning to predict one's subjective experiences, which starts as a trial and error process but quickly finds itself constrained to varying levels of consistency with pre-existing ideas,

(c) objective reality intrudes upon our subjective existence through our physical sensations (and memory limitations), and prediction of the objective world through prediction of our subjective perceptions is a trait with selection advantages from an evolutionary point-of-view,

(d) natural selection must necessarily be constrained to balance the opposing requirements of efficiency and flexibility, where efficiency favors faster to respond instincts and mental habits built up by repetition but flexibility favors the much slower process of mental creativity,

(e) one's imagination, and hence one's mind, is completely unconstrained, but a stable mind, one that can predict its own experiences, must necessarily be organized around mutually consistent sets of ideas, and the growth of a mind will mimic natural selection in that stability and creativity are opposing factors which lead to an increase in complexity and specialization in the mental environment, and, finally,

(f) there probably are only a relatively small set of stabilizing concepts upon which to organize a mind, the most important of which may be the maintenance of an overall sense of one's perceived purpose in life and the recognition that problems require solutions.

Is this all? For now, yes. A mind can not be built, but must develop itself over time. No two minds will be identical, but similarities in our physical circumstances may be enough for our minds to connect through similarities in the ideas we develop in trying to predict our circumstances.

Is this enough? Most likely, no.

What's left out? Much. Perhaps most. Such as the wonder of sensation. How is it that we actually have sensations that we can feel, rather than signals that we just respond to? Ideas, and the free-will to choose between them. Can they truly "emerge" from a system continually trying to predict itself, or any other conceivable mechanism? Caring. How is it that we seem to care about questions and answers? It's not just problem-solving, but problem-solving with personal interest. I have the questions, I need answers. For the moment, I suspect all of us in this field are suffering from one degree or another of knotted bee syndrome.

Meaning is the discovery of possibilities. That voyage never truly ends.

Dr. Dmitri Stephanopolis
Chief Health Officer
409th Watch





"young Girl Thinking" 1/18 Anita Markley '82 © Imp.