

NARRATIVE THINKING AND DECISION MAKING:

How the stories we tell ourselves shape our decisions,
and vice versa

Lee Roy Beach



Cover art: "Risk"/Enamel on board/Lee Roy Beach/2009.

Text and cover art copyright © 2009 by L. R. Beach: All rights reserved. No part of this book may be reproduced or utilized in any form or by any means, electronic or mechanical, including photocopying, recording, or by any information storage and retrieval system, without permission in writing from the author.

Citation Format: Beach, L. R. (2009). *Narrative Thinking and Decision Making: How the Stories We Tell Ourselves Shape our Decisions, and Vice Versa*. Online publication: www.LeeRoyBeach.com.

This book is dedicated to
Jerome Bruner and Paul Falzer

with special thanks to
Linda Miller and Clark Johnson

Contents

Preface	6
Chapter 1: The Fundamental Puzzle	9
Chapter 2: Narratives and Forecasts	31
Chapter 3: Memory, Rules, and Expectations	58
Chapter 4: Plans	77
Chapter 5: Decisions	95
Chapter 6: Decision Aids	128
Chapter 7: Organizational Decisions	151
Chapter 8: Commentary	177
Appendix 1: Research on the Discrepancy Test	216
Appendix 2: Interpreting D'	225
Subject Index	230
Reference Index	235

Boxes and Figures

Box 1a: A brief history of cognitive studies	10
Box 1b: Pascal's argument	15
Box 1c: Expected value math	18
Box 1d: Where things are now	21
Box 2a: Components of experience	33
Box 2b: Narratives in psychotherapy	39
Box 2c: Sources of the desired future	46
Box 2d: Demonstrating forecasting	52
Box 3a: Cues for inferring values	70
Box 4a: Tools	80
Box 4b: Recycled plans	83
Box 4c: Incremental and master plans	90
Box 5a: The primacy of discrepancies	98
Box 5b: Thresholds	103
Box 5c: Recognition and identification	118
Box 8a: Playing the lottery	207
Figure 4.1: A diagram of a plan...	81
Figure 4.2: A tree diagram of Ed's plan	89
Figure 6.1: Example of a decision sheet	131
Figure 6.2: Decision sheet for commuting to work	137
Figure 6.3: Decision sheet for having another child	144
Figure 7.1: Decision sheet for Steller Art Frames	169

Preface

Cognitive narratives are the stories we tell ourselves to make sense of our experience, by providing continuity between the past and present. Narratives also allow us to make plausible forecasts about what the future will be if we do not intervene to change it. Decision making is the act of evaluating the desirability of the forecasted future and, when it falls short of our values and preferences, choosing appropriate interventions to ensure that the actual future is more desirable than the forecasted future. In short, decision making is the process through which we manage the development of our narratives and, in doing so, manage the progress of our lives.

This view of decision making is in sharp contrast to the view that has dominated the psychological study of decision making since its inception. In the 1960's, psychologists provisionally adopted economics' rational choice theory as a descriptive psychological theory of decision making—so called Economic Man was the model for how decisions should be made. Rational choice theory is fundamentally about the best way to make bets, but by assuming that all risky decisions are essentially bets—the gamble analogy—the theory could be extended to all risky decisions, which is most decisions. Doing so permitted rational choice to be broadly incorporated into economic theory as the model of what consumers and other economic actors do when making risky decisions.

Whatever its prescriptive validity for betting, research soon produced ample experimental and observational evidence that the theory does not accurately describe what decision makers actually do when they make financial decisions or any other kind of risky decision, which means that it is an inadequate psychological theory of decision making.

There were three responses to these research results:

- First, having declared decision makers to be irrational and “mental cripples” because they are either unable or unwilling to follow rational choice theory's prescriptions, there was a continuing effort to catalogue the irrationalities and, in some cases, to devise methods aimed at ameliorating them so decision makers' behavior might better fit the theory.
- Second, the theory has been altered to make it more descriptive, culminating in a new discipline called Behavioral Economics, the goal of which appears to be to “psychologize” economics and turn it into a more experimental social science.
- Third, most psychologists appear to have abandon hope for a general psychological theory of decision making and turned instead to mini-theories tailored to specific areas; for example, the effects of emotions on decisions.

Common to all three responses, however, is the explicit or implicit retention of the gamble analogy.

There are numerous reasons for the lack of correspondence between rational choice theory's prescriptions and what decision makers actually do. The primary one, however, is that decision makers don't frame their decisions as bets; they frame them as choices of courses of action to which they expect to devote time, energy, and good judgment in an effort to make sure they turn out the way they want them to. In short, they view decisions as tools for actively managing the future so it conforms to their values and preferences; they simply aren't trying to do what rational choice theory does.

Note the difference between active management of the future and the primary requirement for betting: gamblers must make their bets and then wait passively to see if they won or lost—intervention is cheating.

In contrast, decision makers seldom are passive and almost always intervene to make sure things come out the way they want them to. The fact is, the gamble analogy is irrelevant to real-life decision making. (Even gamblers rely on a form of intervention, luck; perhaps the gamble analogy doesn't even hold for gambling.)

If the gamble analogy is irrelevant to a psychology of decision making, the theory of rational choice is irrelevant, leaving psychology without a theory. My purpose in writing this book is to provide the needed theory.

Chapter 1: The Fundamental Puzzle

The scientific study of the human thought began with three questions that comprise "the fundamental puzzle:"

First, how do floods of minute electrical discharges in the neurons of your auditory, visual, gustatory, and olfactory systems, as well as in your skin and muscles, become sound, vision, taste, and smell, as well as touch and kinesthesia,¹ and how do these sensations come together to form perceptions of discrete objects that have an independent existence in a world that you regard as external to you and that seem completely real, not illusory?

Second, how do these perceptually discrete external events get sorted into meaningful sequences of related events and how do those sequences get blended into a meaningful flow of subjective experience that gives rise to an understanding of the past, experience of the present, and expectations about the future?

Third, how do expectations about the future give rise to actions; in particular, actions that prevent undesirable expectations from materializing—(and what does "undesirable" mean)? In short, how do we manage the future by anticipating it and changing it when it needs changing?

The first of the three questions is studied by sensory and perceptual psychologists and neuroscientists; the second and third questions are studied by cognitive psychologists and are the subject of this book.

Origins of the Puzzle

Although the mysteries of consciousness, subjective experience, and thought have long been the subject of philosophy, religion, and what

¹ Kinesthesia is the sense of the position of your body's core relative to gravity and of your head and limbs relative to your core.

are sometimes called Esoteric Studies—which includes a broad range of beliefs and techniques for obtaining understanding—it took a long time for Western science to address them. This is largely because the modern scientific movement actually isn't very old and, until recently, its focus has been on the physical world. As a result, scientists didn't regard the fundamental puzzle as a puzzle. Like non-scientists they simply regarded subjective experience and mental activity as evidence of The Divine Spark—the will, the spirit, the soul, or the mind (which often was likened to a little man in the head, called a *homunculus*). It was only in the 17th and 18th centuries, when scientists began to learn more about sensory and brain physiology that the fundamental puzzle became a scientific puzzle, eventually leading to the birth of psychology as a scientific discipline.

Box 1a: A Brief History of Cognitive Studies

It was attempts by physiologists to solve the fundamental puzzle using scientific methods that gave birth to psychology as a scientific discipline, within which cognitive psychology developed as a sub-discipline that eventually grew into a new, vibrant inter-disciplinary area of study called cognitive science.

In the 19th century, the German scientists, Ernst Weber (1795-1878), a physician-physiologist, and Gustav Fechner (1801-1887), a physicist-physiologist, pioneered research on part of the fundamental puzzle of experience, how perception changes as a function of changes in sensory stimulation. In the 20th century the Russian physiologist Ivan Pavlov (1849-1936), of drooling dog fame, pioneered research on another part of the puzzle, how expectations arise and how actions derive from them.² Both lines of research were taken up by other scientists throughout Europe and were brought to the United States, most notably by William James (1842-1910), a physician-physiologist-philosopher who established the first psychological laboratory in North America, at Harvard, and who wrote an influential textbook, *Principles of Psychology* (1890), that is still worth reading for its astute insights.

Throughout the 19th, and early 20th centuries, philosophers studied the nature of mind using trained “introspectionists”—people who were adept at analyzing their subjective experience and describing its various nuances. This way of studying experience continued and evolved in the hands of people like the Austrian physician and brain specialist Sigmund Freud (1856-1939), whose psychoanalytic theory shaped how the public in general thought about the human mind throughout most of the 20th century. In scientific circles, however, introspection and related methods, as well as psychoanalytic theory,

² Pavlov won the Nobel Prize for medicine in 1904 for his work on digestion and he did notable work on brain physiology in addition to his work on conditioned reflexes.

came into increasing disrepute because of their imprecision and the concept of "mind," in all of its manifestations, was increasingly regarded as fundamentally unscientific.

In 1913, an American, John Watson (1878-1958), one of a new breed of psychologists whose doctorates were in psychology rather than medicine, physiology or philosophy, issued "The Behaviorist Manifesto," in which he declared that psychology "...is a purely objective experimental branch of natural science. Its theoretical goal is the prediction and control of behavior. Introspection forms no essential part of its methods, nor is the scientific value of its data dependent on the readiness with which they lend themselves to interpretation in terms of consciousness."

Watson's manifesto launched a revolution, called Behaviorism, which eschewed reference to conscious experience and distained the concept of mind in any form, attempting instead to build a science of observable behavior. The building blocks of this science were Pavlov's "classical conditioning" and, later, B.F. Skinner's (1904-1990) "operant conditioning," known to most of us by their respective signature experiments in which dogs salivated in anticipation of being fed when a bell sounded and hungry rats learned to press a lever in order to obtain food.

Behaviorism was the prevailing viewpoint in most English-speaking countries from about 1920 through the 1950's. Interest began to wane as it became clear that it required impossibly convoluted reasoning to account for complex animal and human behavior in terms of chains of simple conditioned reflexes or even complex operant learning. Fortunately, a new invention had come along that prompted what has been called the First Cognitive Revolution.

Although he probably didn't recognize it as such, the underlying metaphor for Sigmund Freud's psychoanalytic theory of the human mind mirrored the premier technology of his time, the steam engine. The id boils away, building up steam that threatens to blow the top off the superego, and the ego provides a safety valve by directing the pent-up energy toward less hazardous ends. The technological metaphor for Behaviorism was the telephone switchboard, in which connections are made between stimuli and responses. The technology of the 1950's, the computer, provided a metaphor for psychologists who were wary of Freudian theory and weary of connectionist Behaviorism. The pivotal notion was that "information" was "processed" using programs made up of specific steps that allowed the computer to perform tasks that previously had been the preserve of human thinking. Moreover, the steps in these programs were more sophisticated than conditioned and operant reflexes, becoming increasingly complex as computer technology grew.

Psychologists began to talk once again about what had in earlier times been called mind and mental activity. Only now they talked about them in computer terms, which seemed to remove the mystery and provide an aura of scientific respectability. I don't want to oversimplify, because there were other trends in play, but it isn't wrong to trace the first cognitive revolution, the birth of what was to become mainstream cognitive psychology, to the advent of the computer.

Once the first cognitive revolution got solidly underway, the computer metaphor itself became obsolete and the new cognitive psychology began to stand on its own. Before long, allied disciplines, particularly linguistics, and new ways of studying cognitive functioning joined forces with cognitive psychologists to form a new interdisciplinary field, cognitive science. Today, cognitive scientists freely use the word "mind," but its meaning is much

different than it was in times past. It no longer is viewed as the mysterious entity that it was in the 17th and 18th centuries or even the mere by-product of brain activity that it was in the 19th and 20th centuries. Instead, it is defined by capacities and functions, the products of which can be clearly observed (memory, perception, language, reasoning, etc.) and, in some cases, can be specifically localized in the brain using modern brain scanning techniques—a far richer and more concrete definition than in former times. The awe that cognitive scientists have for the mind is not the awe due the mysterious but the awe due something intricate and challenging to study.

Every metaphor has its drawbacks. The computer metaphor's was its necessary emphasis upon information processing to the exclusion of information meaning and upon computability to the exclusion of any other kind of operation. The necessary result was a focus on paradigmatic, procedural reasoning that has never satisfied anyone who is aware of the vagaries of their own and others' thinking. Some psychologists believe that we are at the beginning of a second cognitive revolution. This one focuses, not on elegant, contentless paradigms for the processing of "information," in the generic sense, but on how humans use narratives to understand, and communicate to others, the very specific meanings and implications of the specific events and people that are at the core of how they experience their lives.

As described in Box 1a, scientific psychology began in the 19th century, starting in Germany and Russia, spreading to Britain and throughout Europe, and then moving quickly to North America and beyond. As is so often the case in science, research revealed that the fundamental puzzle was far more complex than it seemed at first. This shouldn't be surprising. After all, thinking scientifically about thinking is somewhat paradoxical—the subject of investigation is doing the investigating—which means that the task is likely to be more difficult than it might otherwise be. However, over the years enough progress was made to show that an understanding of the fundamental puzzle requires, at the very least, an understanding of four universal human cognitive abilities:

- *Memory*; the ability to retain, retrieve, and use information about the past.
- *Perception*; the ability to combine current sensory information and information from memory to produce a coherent understanding of the present.

- *Imagination*; the ability to use information about the past and present to...
 - *Forecast* the future,
 - *Envision* desirable alternatives to the forecasted future.
- *Decision Making*; the ability to detect an undesirable forecasted future, to select actions that will promote a desirable alternative future, and to monitor the actions' progress toward achieving the alternative future.

These four abilities are the answers to the four questions that underlie experience: Memory = what led up to this moment? Perception = what is happening now? Imagination = what can be expected to happen in the future? Decision Making = what can be done about it? Note that answering the last question requires a value judgment about the desirability of the forecasted future as well as about the desirability of envisioned alternatives to that future, a value judgment which involves feelings and emotions. Note too that the answers to the four questions comprise a story that is rooted in the past, includes the present, and extends into the future which will be shaped, at least in part, by purposeful action.

Of the four universal human cognitive abilities, memory and perception have long, distinguished research histories and a great deal has been learned. Imagination has received less attention than it deserves but, even so, a good deal is known and more is currently being learned. In contrast, empirical research on decision making only began in earnest in the 1960's. Until then, economists were the only scholars who gave much careful thought to how decisions are made. As a result, when psychological research on decision making began, the first order of business was to test how closely economists' ideas about decision making correspond to real decision making.

In the remainder of this introductory chapter, we will examine (briefly, and therefore somewhat simplistically) the basics of Economic Decision Theory and what happened when psychologists examined how closely it matches real decision making.

Economic Decision Theory

Ask people to recall decisions they've recently made and they'll come up with all sorts of things—decisions to get a haircut, go to a movie, take dance lessons, have pizza for lunch, go to bed early, and so on. But, ask them to recall only difficult decisions and they'll almost always describe circumstances in which they had to choose the best from among two or more risky options. That is, they had to choose but there was a risk that their choice would not turn out to be as good as they hoped it would be.

Economic Man

One of the earliest discourses on decision making appeared in the 17th century, when the geometrician, physicist, and mathematician Blaise Pascal (1623-1662) formulated his famous argument for choosing to believe in God (Box 1b). His argument translated faith into a bet and the choice to believe or not was determined by the expected payoffs for winning or losing the bet. In the 18th century, Daniel Bernoulli (1700-1782) advanced a similar logic to help gamblers choose between actual bets; which to play and which to avoid. (He also proposed the concept of utility, which is the subjective worth of a payoff to the bettor rather than its "objective" market value, and he demonstrated that insurance is a form of betting in which you bet you'll meet with disaster and the insurance company bets you won't.)

Box 1b: Pascal's Argument

Blaise Pascal was born in France in 1623. At an early age he showed a talent for mathematics and geometry. At the age of 16 he formulated Pascal's Theorem on conic sections and at the age of 19 he constructed a mechanical calculator that could add and subtract. Throughout his career, he did important work on, among other things, the basics of probability theory and the principles of hydraulics and vacuums. He invented the syringe as well as the hydraulic press, which amplifies force.

Pascal became religious in his later years, writing philosophical and theological works, among which is his famous argument for believing in God. I will phrase the argument in modern terms:

Think of whether or not to believe in God as a choice between two bets, between which you must choose: (1) Don't believe or (2) Believe. You will learn whether you won or lost your bet when you die: (A) He doesn't exist or (B) He exists. The question is, in light of the consequences of winning or losing, which bet you should choose?

Consider the following diagram: The bets from which you must choose are listed on the left side (1 and 2), the possible outcomes are listed on the top (A and B), and the consequences for winning or losing each bet are in the cells of the diagram.

		<i>Possible Outcomes</i>	
		(A) God Doesn't Exist	(B) God Exists
<i>Bets</i>	(1) Don't Believe	No Loss	Eternity in Hell
	(2) Believe	Wasted Worship	Eternity in Heaven

Your bets (options) can be summarized as:

- *Bet 1:* If you bet that God doesn't exist and it turns out that He doesn't, then you win nothing other than the satisfaction of being right. But, if it turns out that He exists, you lose and you'll spend eternity in Hell.
 - *Bet 2:* If you bet that God exists but it turns out that He doesn't, then you lose the time and energy you've invested in worship. But, if it turns out that He exists, you win and you'll spend eternity in Heaven.
- Because the consequences of betting that God doesn't exist (bet # 1) and winning are minimal but the consequences of losing are Hell, and because the consequences of betting that God exists (bet # 2) and winning are Heaven but the consequences of losing are merely inconvenience, Pascal argued that rational persons should bet that God exists (bet #2) and behave accordingly.

Time passed with various degrees of development of these basic ideas until, in the 1940's, two mathematicians, John von Neumann (1903-1957) and Oskar Morgenstern (1902-1977), adopted Pascal's characterization of risky choices as bets and elaborated upon Bernoulli's rules. In doing so, they developed an axiomatic, mathematically precise "theory of rational choice," which prescribed how to make risky decisions. Their formulations were broadly accepted by economists and soon were viewed as the basic paradigm for making any decision involving

uncertainty—what so-called Economic Man would do when faced with a risky choice.³ For convenience, we will call this paradigm Economic Decision Theory, or Economic Man for short.

The paradigm. In its most basic form, Economic Man's paradigm for evaluating a risky choice starts by treating it as a bet (see Box 1c for the math). Then

- For each bet on offer, he multiplies the probability of winning by the value of what he will gain if he wins the bet (after having subtracted the stake he will have to pay to play the bet).
- Next, he multiplies the probability of losing by the value of what he will lose if he loses the bet (usually just his stake, but sometimes his stake plus additional losses).
- Finally, he subtracts the second product from the first and makes his decision:
 - If the bet is the only one on offer, he plays it if the difference is positive (the first product, for winning, is larger than the second, for losing).
 - If more than one bet is on offer, he evaluates each of them in the manner described and plays the one that has the largest positive difference.⁴
- Having decided, he pays the stakes and waits to see if he wins or loses.

³ Economic Man (aka, *Homo Economicus*) is a convenient fiction used by economic theorists. He is completely "rational," in that he consistently behaves the way the theorists think he should: He knows his preferences and can attach numbers to them to indicate their magnitude—usually their market value; he knows the probabilities of future events occurring—usually based on actuarial data; he is capable of accurate complex computations; he is capable of entertaining multiple options simultaneously; and he is capable of precisely appraising and comparing these multiple options in order to select the best option.

⁴ If there are no bets with a positive difference, he refuses to play because none of the bets is favorable to him—which usually is the case for every bet offered in a casino and is why both you and Economic Man should stay out of casinos.

Notice that when he multiplies the amount he can win or lose by a probability, Economic Man is, in fact, discounting the amount to be won or lost because a probability is a decimal number that is a measurement of risk inherent in the bet. (For example, the probability of a tossed coin landing heads up is .50, so if the payoff for the coin turning up heads is \$1, multiplying the two numbers yields \$.50 or fifty cents—the dollar was discounted by 50%). The result of this is that he ends up comparing the discounted amount he might win versus the discounted amount he might lose—subtracting the latter from the former makes it easy to compare how much one amount surpasses the other, called the bet's *expected value*. Even with a positive expected value he will sometimes lose because the outcome is uncertain (hence the inclusion of probability), but it can be shown mathematically that, over a large number of opportunities, if he only plays bets with positive expected values, he will make a profit. When there are multiple bets on offer, always playing the one with the greatest positive expected value will yield the maximum profit in the long run. This strategy is called *maximizing expected value*.

Although you probably never used this precise paradigm to make a decision, the general idea has some appeal. We all know we have to strike a balance between attractive outcomes and the odds of ever getting them. Pursuing unattainable goals may be romantic but it isn't

Box 1c: Expected Value Math

Formally, the expected value for a simple win-lose bet is computed using this equation:

$$EV_i = P_i W_i - (1 - P_i) L_i,$$

where

- P_i = the probability of winning bet i ,
- $1 - P_i$ = the probability of losing bet i ,
- W_i = the value of winning bet i (after subtracting your stake),

L = the negative value of losing bet i (the loss of your stake and any additional losses).

That is, the expected value of bet i (EV_i) is the difference between the positive value of winning the bet (W_i) discounted by the probability of winning (P_i) and the negative value of losing the bet (L_i) discounted by the probability of losing the bet ($1 - P_i$). (Probabilities discount payoffs because probabilities are decimal numbers; multiplying a whole number, W or L , by a decimal diminishes the whole number—the larger the probability the less W or L is discounted.) A bet that has a positive EV_i is in your favor; if multiple bets are offered, the one with the largest EV_i is the one that is most in your favor.

Equation 1.1 is merely the basic equation; it can be elaborated in many ways to reflect more complicated bets, including multiple-outcome bets and multi-stage bets in which winning at one stage allows you to move to yet another stage and another bet. The complications are virtually endless and have given rise to very sophisticated tools for analyzing complex bets.

very practical. Pursuing things you have a shot at getting is a far wiser strategy for most day-to-day decisions. Somehow, Economic Man's strategy of increasingly discounting the attractiveness (and unattractiveness) of payoffs when they're increasingly unlikely to happen seems reasonable. Furthermore, choosing to do things that have a greater prospect of ending well than ending badly seems imminently reasonable. Therefore, you can see why psychologists, and economists before them, were willing to at least consider the Economic Man paradigm as appropriate for decision making; it isn't at all far-fetched.

Real People

In the 1960's psychologists began to investigate the correspondence between Economic Decision Theory and actual decision behavior—between Economic Man and the rest of us. They started with the working hypothesis that our normal way of making choices approximates Economic Man's paradigm for choosing bets—and, by extension, that it is the way we make every kind of decision involving uncertainty; which is to say most decisions. Put another way, the researchers' hypothesis was that Economic Man's decision paradigm is a logical and mathematical formalization of what most of us do naturally and intuitively when we make any kind of decision. They tested the

hypothesis by examining how closely our decision behavior actually resembles Economic Man's in a wide variety of laboratory and real-world situations. Note that in virtually all of these experiments, decision makers had to rely on their own resources. They weren't given pencil and paper, calculators, or instructed about how Economic Man would go about the tasks that they were given. (This is called *unaided* decision making as opposed to *aided* decision making, which we will discuss later.)

Research Results. When the researchers began making their comparisons, they weren't expecting us to behave exactly like Economic Man. After all, Economic Man is superhuman. But, even at that, they were surprised by how little resemblance there actually was (surprised enough to publish hundreds of research articles about it).

Not only do our decision strategies fail to resemble Economic Man's, even people who know about his strategy tend not to use it and when they do, they don't do it very well. Everyone, it seems, has trouble appreciating the value of the payoffs associated with options, even when told the market values. Moreover, we don't seem to know how to handle information about probabilities; in fact we don't understand much about probabilities at all. We aren't very good at mental arithmetic. We oversimplify in order to avoid complex calculations, yet we complicate things by factoring in our aversion to risk as well as, in some cases, our belief in luck—both of which would be irrelevant for Economic Man. As a result, from Economic Man's standpoint, our evaluations and procedures are full of error, making it an absolute certainty that the resulting decisions will be severely flawed.

Common Sense. But, common sense and simple observation don't wholly square with the research results. If it were the case that Economic Man's paradigm, is in fact the single best way to make all decisions, you'd think that our disinclination to do it, not to mention our inability to do it properly,

would doom us to making consistently rotten decisions. But, look around. Most people survive and prosper in spite of their seeming unwillingness and inability to behave like Economic Man, which suggests that Economic Man isn't the right criterion for good decision making.⁵

The reason Economic Man isn't the right criterion is that what he was designed to do isn't what most of us are trying to do when we make decisions. Economic Man was designed to pursue an optimal strategy for gambling. But, gambling requires you to place your bet and wait passively to see if you won or lost—any attempt to influence the outcome is regarded as cheating. In contrast, when most of us make decisions, we fully intend to do what it takes to make things turn out well—and it isn't cheating, it's initiative. Even though we sometimes talk like we regard our decisions as bets, our intention to intervene in the outcome means that they really aren't; therefore Economic Man's paradigm isn't appropriate. Therefore, we must look elsewhere if we want to develop a useful psychology of decision making (see Box 1d)

Box 1d: Where Things Are Now

When research revealed that humans don't much resemble Economic Man, researchers went off in either of two directions. Some sought to "psychologize" economics by introducing the behavioral research methods pioneered in psychology, work that culminated in the 2002 Nobel Prize in Economics for Vernon Smith, an economist, and Daniel Kahneman, a psychologist whose work with the late Amos Tversky provided a bridge between economics and cognitive psychology. This, together with earlier work by Herbert Simon (Nobel Prize in Economics, 1978), who linked economics and organizational psychology, has created a new discipline called Behavioral Economics.

⁵ We all have made decisions we came to regret, largely because they didn't work out the way we thought they would, not necessarily because the decision process was flawed. Without the benefit of hindsight, I suspect we'd make most of those bad decisions all over again.

Other researchers, retaining the name Behavioral Decision Theory, which was given to the field in the 1960's, remained within psychology. Perhaps disillusioned by the failure of Economic Decision Theory, they appear to have given up on crafting a general psychological theory of decision making. Instead, they have created a multiplicity of semi-independent research areas. The area receiving the most attention at the moment is how emotions impact decision making: the influence of mood on risk aversion and the influence of anticipated or experienced regret on choice. (The general impression one gets from this work is that emotion continues to be viewed, as through the eyes of Economic Man, as deleterious to rational decision making.) Other areas study sunk costs (throwing good money after bad), endowment and status quo effects (resistance to change), overconfidence (being more certain than is warranted about what you think you know), and risk assessment (e.g., driving seems safer than flying although statistics show it isn't)—all of which, in one way or another, are about decision makers' "biases" relative to economic reasoning or statistical information. In short, Economic Man is alive and well, he merely is a little stealthier than he used to be.

A New Look at Decision Making

Economic Decision Theory, Economic Man, was adopted as a viewpoint about the psychology of decision making early in the First Cognitive Revolution in psychology (Box 1a), just as computers began to have an impact. The idea was that if computers could do tasks previously doable only by humans, it might be profitable to think of thinking in terms of how computers function. It quickly became fashionable (if that isn't too trivial a term for something that was taken very seriously) to use the language of computer programming to describe the sequence of cognitive procedures that it was assumed made up human problem solving.⁶ Doing so both introduced and reinforced the view of thought as fundamentally procedural; how "information" was "processed" in order to solve a problem or perform a task. Behavioral Decision Theory fit the spirit of the times because it is procedurally prescriptive; how to process

⁶ This view led naturally to questions about what in cognition is analogous to the computer programmer, the person who designed, instigated, and controlled the cognitive programs as they were executed. This problem was solved by introducing the concept of an Executive Function; which, many observers thought, looked suspiciously like the long-repudiated homunculus—the little man in your head. Thus, after being on hiatus during the Behaviorist regime, the mind quietly returned to psychology.

information about probability and value in order to evaluate the worth of a bet. We've seen, however, that things didn't work out well when it was compared to what people actually do.

The problem with using information processing and computation as basic concepts of cognitive theory (and, similarly, with using Economic Decision Theory as a psychological model of decision making), is that it gets the cart before the horse. This error wasn't apparent at the time, but it became so with the advent of what some have called the Second Cognitive Revolution. This revolution started in 1968 with a book by Jerome Bruner, a major figure in psychology for more than half a century. Bruner distinguished between two forms of cognition, narrative thought and paradigmatic thought. Roughly speaking, narrative thought is discursive and paradigmatic thought is procedural. Theories that derive from the analogy between human cognition and computers, and between human decision making and Economic Decision Theory, implicitly assume that paradigmatic thought is the "natural" way for people to approach problems. From there it is just a short jump to assuming that people always are trying to think paradigmatically, but they aren't very good at it. Bruner challenged this assumption, arguing instead that paradigmatic thought grew out of—and is supplemental to—narrative thought, which the primary mode of thinking.

Narrative Thinking

The next chapter (Chapter 2) is about the nature of narrative thought and the role it plays in our lives. But, before we get there, you at least need to know what a narrative is:

The elements of narratives are symbols that stand for real or imagined events and actors, where the latter are animate beings or inanimate forces. The glue that binds the events and actors are causality and implied purpose. A narrative consists of a temporal arrangement of

events that are purposefully caused by animate beings or are the result of inanimate forces (although we often superstitiously attribute purpose to inanimate forces). The narrative's storyline is the emergent meaning created by a coherent arrangement of the events and actors; where coherence is provided by time, purpose, and causality. That is, just as arranging words into sentences creates emergent meaning that the unarranged words do not have, and arranging sentences into larger units creates even more emergent meaning, arranging events, actors, time, purpose, and causality into a narrative creates the emergent meaning that is the narrative's storyline or plot.

A good narrative is plausible, in that none of the component events, actors, and actions is outlandish or improbable, and it is coherent, in that everything is neatly encompassed by the narrative, leaving few or no loose ends. And you easily can tell a good narrative from a bad one. Studies in which participants are shown various arrangements of elements find a great deal of agreement about which arrangements "make sense" and which are meaningless or ambiguous. This isn't so surprising when you consider that the ability to create narratives begins early in life, soon after children learn to talk, and improves with time. Between the ages of two and eight, most children learn to tell stories that listeners can understand and respond to, and by the age of 10 they have learned to tell well formed, elaborate stories.

Narratives need not be true, in the sense of being objectively factual. They can be about the past as it was or as it might have been, about the future as it might be or as you wish it would be, about real or imaginary actors, known or imagined motives, and so on. We usually try to make sure the narratives we use to make sense of our own lives are reasonably accurate but, as Sigmund Freud tried to tell us, we don't always succeed. Of course, we often indulge in untrue narratives, either

our own (as in daydreaming and wishful thinking) or others' (novels, TV shows, gossip), because they provide entertainment and, sometimes, unexpected insights into our more serious narratives that allow us to reorganize them or create wholly new ones. And, of course, we lie to ourselves and to others, and we often believe our own lies.

Because the rest of this book is about the uses of narratives, we needn't say more here. Instead, let us turn the other mode of thought, paradigmatic thinking.

Paradigmatic Thinking

A paradigm is a set of prescribed procedures for performing a specific class of tasks. Both informal and formal education is, in large part, aimed at acquisition of paradigms—how to do things, among which is solving various kinds of intellectual and manual problems. Thus, for example, as an infant you informally learned the prescriptions of grammar as you learned how to talk. Later, as a school child, you were formally taught the prescriptions of standard grammar (which may have differed to some degree from what you learned earlier) as well as the procedures for doing arithmetic, geometry, and writing book reports. As your education progressed, you were taught the prescribed procedures for doing algebra, trigonometry, calculus, and writing term papers, theses, and, perhaps, dissertations. And, somewhere along the way, you may have been taught the procedures for bookkeeping, computer programming, calculating the probabilities of events, performing statistical tests, formal logic, and the scientific method. In addition, you learned paradigms for manual tasks: how to drive a car, how to play baseball, how to raise roses, how to iron shirts, how to fry eggs, how to operate a computer or an electric drill, and so on. In short, in the course of your life you have learned many paradigms for performing many kinds of intellectual and manual tasks, as well as how to use the tools that they

require for their execution (pencils, calculators, computers, tape measures, hammers, screwdrivers, electric drills, etc.).

All paradigms, intellectual or manual, have three things in common. The first is that they each grow out of a human need to think precisely and economically about some class of tasks. Precision requires definitions and sequences of clear-cut steps. Economy requires that solutions be general, repeatable, and communicable to others faced with the same tasks (and to you, when you are learning the paradigm).

The second commonality is that paradigms are contentless. It is only when they are brought to bear on a specific task that content enters into things, and that content comes from the task itself and why it is being done, not the paradigm itself. Grammar tells you how to structure a sentence, but not what to say, and certainly not what you mean. Probability theory tells you how to calculate the probability of an event, but not what the event is, why it is of interest, and certainly not why it may or may not happen.

The third commonality is that paradigms are general. Grammar allows you to construct a variety of sentences to communicate a variety of meanings (even if you can't consciously describe the grammatical rules you use to do it). You learn the algebra so you can use it on a variety of problems. You learn how to play baseball so you can play repeatedly. It is this generality, this "portability," that makes paradigms so valuable—once learned, the paradigm is available whenever it is needed (even if, like grammar, it is so overlearned that you can't describe it or how you used it). Portability is neatly embodied in the saying, "Give a man a fish and he will eat today; teach a man to fish and he will never go hungry."

Paradigms have their origin in narrative thought. Their purpose is to provide a way of thinking more precisely and rigorously about some topic

than we can with narratives. However, once created, paradigms become *things* that can themselves be subjects of both narrative and paradigmatic thought. For example, a physicist can give you a broad narrative about the movement of the heavens, but he or she must resort to a mathematical model, a paradigm, to provide either a precise description or to make precise predictions. That mathematical model can itself become the subject of a narrative and it might be tested using formal logic or a higher order of mathematics. The point is, once the leap is made from narrative to paradigm, and as the paradigm becomes increasingly sophisticated, the link between it and the parent narrative may become decreasingly apparent.

Probability theory is a good example of a paradigm that is somewhat estranged from its parent narratives. The general concept of probability is rooted in a narrative need to express uncertainty about truth of a proposition ("It probably will rain). A few hundred years of effort to add precision has resulted in a self-contained, precise, mathematical theory in which the concept of probability has become so esoteric that it is unrecognizable as the subjective uncertainty that started it all. For probability theory, probability is the likelihood that a specified kind of event will occur in a random sample of events, often measured as a relative frequency (more precisely, as the limit of a relative frequency over an infinite series of random samples from a population of events). This is very different from what most of us mean when we say, "It probably will rain," but it has produced a theory that is of immense practical value—for things like quality control, insurance, opinion polling, agricultural and social research, and, of course, gambling.

Economic Decision Theory is another estranged paradigm. It began as an answer to gamblers' need for procedures to reliably evaluate bets, but its prescriptions now reach far beyond such petty pastimes. Once the

analogy between bets and risky decisions was made, it was but a short step to the assumption that market participants are rational decision makers (in the manner of Economic Man) who look out for their own self interest and, in doing so, create markets that are similarly rational—a conclusion that is sorely tried each time the economy runs into trouble. Indeed, economic rationality is at the core of the debate about modern economic theory—and way down deep, so deep that it usually goes unnoticed, economic rationality turns on the analogy between gambling and all risky decision making. If that analogy is unsound, if market participants' decisions do not conform to Economic Man's procedures for evaluating gambles, then the edifice is in danger of collapsing—or, at least, of needing some substantial shoring up.

Returning to the 1960's, given the spirit of the times it was reasonable for psychologists to use paradigmatic Economic Decision Theory as the starting point in their search for a descriptive theory of decision making. What wasn't reasonable was their persistence in assuming that all decision making is paradigmatic and that Economic Decision Theory is the correct paradigm, even though the data showed that decision makers don't think that way. The usual rule in behavioral science is that when behavior and a theory don't agree, it is the theory that is questioned, not the behavior. Had the researchers applied this rule, they would have rejected the theory as an inaccurate description of human decision making and looked elsewhere for a descriptive theory. Then, armed with knowledge about how decisions actually are made, they could, if they wished, have developed paradigms for helping decision makers do a better job of it—after all, that is what paradigms are for. Those paradigms might or might not resemble Economic Decision Theory, depending on what is needed to help decision makers do what it is that they are trying to do when they make decisions.

This Book

In the interest of full disclosure, I have to say that I played an active role in the attempt to study decision making as paradigmatic thought. Like everyone else, I was loath to give up the endeavor. (If nothing else, Economic Decision Theory is so elegant it exudes the seductive scent of science.) But experimental data and personal observation eventually wore me down. I then set out to find an alternative approach, which led to a theory that was less paradigmatic in tone, but that failed to recognize narrative thought as applicable to decision making.⁷ Of late, I have gone the full distance, and am convinced that narrative thought is the bedrock of cognition and, therefore, is the bedrock of decision making.

I also have concluded that narrative-based decisions have their limitations, especially for complex decisions, and at some point it is helpful to introduce compatible paradigmatic procedures to mitigate these limitations. For example, it is difficult to keep enough in mind at one time to do justice to complex decisions. Paradigmatic procedures, together with tools like pencil and paper and a calculator, can help us overcome this, while at the same time making the decision more transparent so we can spot inconsistencies and similar difficulties before the decision is made. But, paradigmatic thinking is not an end in itself; it always plays a supporting role. Narratives are where the action is.

The purpose of this book is to elaborate on the idea of decision making as narrative thought supported by paradigmatic thought, and to suggest some useful paradigmatic processes that can help us keep our narratives straight when things get complicated. The first part of the book (Chapters 2 through 5) will examine four propositions:

⁷ The theory, which was developed with my good friend and colleague Terrance Mitchell, was called Image Theory.

- Cognitive narratives provide meaning and continuity by linking the past with the present and allowing us to make educated guesses about the future (called forecasts). We will address this in Chapter 2.
- Based on our past experience and our values, we create rules, some of which are used to make forecasts and some of which are used to evaluate the desirability of the future offered by those forecasts. We will address this in Chapter 3.
- When the forecasted future is undesirable, we construct plans that guide actions designed to promote a more desirable future. We will address this in Chapter 4.
- All decisions are about discrepancies and what to do about them. We will address this in Chapter 5.

The latter part of the book (Chapters 6 and 7) will provide practical guidance, based on the foregoing four propositions, about how to make more satisfying personal and organizational decisions as well as providing paradigms for doing so. As we will see, these paradigms do not much resemble Economic Decision Theory.

The final chapter (Chapter 8) will summarize, compare the present view with other views, and expand on some of the earlier themes.

Sources and Further Reading

- Basu, K. (2007). The Traveler's Dilemma: When Playing This Simple Game, People Consistently Reject the Rational Choice. *Scientific American*, (June), 90-95
- Beach, L. R. (Ed) (1996). *Decision making in the workplace: A unified perspective*. Mahwah, NJ: Lawrence Earlbaum.
- Beach, L. R., & Connolly, T. (2005). *The psychology of decision making: People in organizations*. Thousand Oaks, CA: Sage

- Bruner, J. S. (1986). *Actual Minds, Possible Worlds*. Cambridge, MA: Harvard University Press.
- Hacking, I (1975). *The emergence of probability*. NY: Cambridge University Press.
- Herman, D. (Ed.) (2003). *Narrative theory and the cognitive sciences*. Stanford, CA: Center for Language and Information.
- Kahneman, D., & Tversky, A. (1979). Prospect Theory: An Analysis of Decision Under Risk. *Econometrica*, 47, 263-291.
- Lowenstein, G. (2001). The creative destruction of decision research. *Journal of Consumer Research*, 28, 499-505.
- Polkinghorne, D. E. (1988). *Narrative knowing and the human sciences*. Albany, NY: State University of New York Press.

Chapter 2: Narratives and Forecasts

Most of the decisions you make don't seem much like decisions, they are just the normal way of going about your life. But their importance is revealed by the trail they leave, which, over the long run, is the history of your life. The points at which the storyline zigs or zags, which are what make your history different from anyone else's, are the points at which you made a decision, acted on that decision, and something, good or bad, happened as a result. The subtext of your history is the emotions that accompanied and influenced those decisions and actions.

To understand how your decisions and actions create your unique history, we must go back to the fundamental puzzle of how you make sense of your ongoing experience and how this helps you forecast and shape your future.

A little introspection will reveal that your ongoing experience consists of a flow of perceptions: visual, auditory, touch, taste, smell, and kinesthetic. You can, at will, direct your attention to each of these; focusing, for example, on what you are seeing and then shifting to what you are hearing and then to what you smell or the taste in your mouth; memory tells you what each of them are. Moreover, when prompted to do so, like right now, you can focus your attention on the pressure of your buttocks against your chair as you read this book, then on the tension in your fingers as you hold the pages open, then on the blackness of the print, and finally on the movement of your eyes as you read this sentence. But, when you step back from introspection into your normal state, these different perceptions are no longer separate. Instead, they meld into an integrated, meaningful flow of conscious experience which, if asked, you are likely to summarize as "reading."

Consciousness

The glue that binds perceptual experiences into a meaningful stream of consciousness comes from within you. First, your brain is wired to experience a continuity of consciousness rather than a set of discrete perceptions. As a result, consciousness is experienced as a stream that flows out of the past (memory), through the present (awareness), and into the future (foresight). Second, what has happened in the past influences your interpretation of what you are experiencing right now as well as what you expect to experience in the future, supplying context and meaning for your stream of consciousness.

We all are preoccupied with the flow of our own and other's experience, which we think about in the form of stories. We tell stories to each other when we're together and we tell stories to ourselves when we're alone. In the morning we consult newspapers, TV, the internet, and radio to get updated stories about what's going on in the world. Throughout the day we e-mail, text-message, telephone, or meet with each other to swap stories about what we're doing, and why, what we're going to do and what we expect the result to be. In the evening we relax with stories in books, magazines, and on TV. Or we go to sporting events which offer stories of striving and victory or defeat. Or we go to the movies, which offer drama, comedy, and romance. Or we attend social events at which we swap stories in the course of conversation. Later, we read bedtime stories to our children and then read a few pages of a novel to lull ourselves to sleep. And as we sleep, we experience more stories in the form of dreams.

Some stories are simply entertaining, others are serious. All are instructive because they help us make sense of things that go on around and within us. Stories may not be the only way of making sense—after all, other species seem to get by without them or with only very simple

ones—but without stories to tie things together, we would find our world far more baffling, fragmented, and unpredictable than it generally seems to be. Moreover, stories tie us to others around us, as we engage in conversation, and to the broader culture, as we read magazines, newspapers, and information on the web. The general term used to refer to the stories that we tell ourselves and others is *narratives*

Box 2a: Components of Experience

Conscious, unconscious, and attention are related, but different, concepts. To help differentiate among them, at home tonight, turn off the lights and shine a flashlight on a wall some seven or eight feet in front of you. The dark beyond the lighted part of the wall, as well as the shadowy edges of the lighted area, is analogous to the unconscious—only indistinctly perceptible but real nevertheless. The more clearly lighted area around the bright center of the light is analogous to consciousness—details are distinct but peripheral to whatever is illuminated by the bright center of the light. The bright center is analogous to attention—details are vivid.

Just as moving the flashlight's beam will move the bright center and make previously vague details vivid and distinct, shifting your attention allows you to more closely examine details of which you were previously conscious but which were not distinct and vivid. The contents of unconsciousness can be brought to consciousness and be made the focus of attention, but it is not done easily. As Sigmund Freud taught, there often are psychodynamic reasons that the contents of unconsciousness are labeled "Dangerous" and resist being brought into the bright light. On the other hand, some things are unconscious merely because they aren't pertinent to the moment, not because they are deep dark secrets you're hiding from yourself.

Narratives

Narratives are more complicated than the rather simple, paired-down stories we tell others or that we are told by others. Our private narratives aren't made merely of words, like a novel or newspaper article. Neither are they simply combinations of words and pictures, like a movie or TV. Nor are they simply words and gestures, like when a friend tells you something. Nor are they simply your interior monologue—the little voice in

your head.⁸ Instead, they are a rich mixture of memories, of visual, auditory, and other cognitive images, all laced together by emotions to form a mixture that far surpasses mere words and visual images in their ability to capture context and meaning. Indeed, it is only when you are called upon to describe them that they become heavily encoded in words; but that isn't their natural state and translation into words always impoverishes them.

Origin of Narratives

Private narratives derive from instruction or from your own experience. Instruction can be either formal, such as the curriculum of a history or management or physics course, or informal, such as what you glean from magazines or novels or learn by searching the web. Private narratives that derive from your own experience are unique to you because your experience is unique to you. Private narratives that derive from instruction will be less unique to you; others who received the same instruction will have similar private narratives but, even at that, no two of them will be exactly alike.⁹ After all, you bring your own background, priorities, and interests when you take a course or read a book. As a result, the private narrative you come away with won't be exactly the same as the ones that others take away.

A Hierarchy of Narratives

Think for a moment about that most familiar of public narratives, a novel. The main plot of a novel describes either a segment or the entire life of the novel's central character or characters. The main plot

⁸ The next time you find yourself lost in thought, note how much your interior monologue resembles a one-sided conversation with another person, which implies that it is comparable in its poverty to the versions of your narratives that you tell other people.

⁹ This is brilliantly illustrated by the classic literary work, *The Alexandria Quartet*, by Lawrence Durrell. Published between 1957 and 1960, it consists of four novels about the same events, each recounting what happened from the startlingly different perspectives of each of four principle participants.

frequently is revealed through descriptions of two or more intertwined and mutually supporting subplots, each of which is a reasonably coherent story in itself. The subplots are revealed through descriptions of episodes; fairly circumscribed vignettes that further the development of the subplot, thereby furthering the development of the main plot. The integrity of the novel as a whole relies on episodes contributing coherently to subplots and subplots contributing coherently to the main plot.

But, a novel is merely ink and paper until someone reads it. Then it becomes something alive and vibrant as the reader constructs his or her own imaginary world from the material provided by the text. This imaginary world is a private narrative the reader creates as he or she proceeds through the novel; a narrative that parallels the story as it unfolds. It focuses most clearly on the part of the novel being read at the moment, but it includes what already has happened leading up to this moment and what the reader anticipates might happen as the story continues to unfold. This private narrative has a main plot, subplots, and episodes that resemble the story in the text but that nevertheless are unique to each reader. We'll call this reader-supplied element of the novel, the reader's narrative.

This hierarchy of a novel's main plot, subplots, episodes, and the reader's narrative has a counterpart in the hierarchy of your own private, real-life narratives

- The novel's main plot is analogous to your autobiography, if you were to write it, and is the way in which you think about the history of your life, your present circumstances, and your dreams for the future. This is your *private autobiographic narrative*.
- The novel's subplots are analogous to your private narratives for the various sectors of your life; work, romance, health, and so on. These are your *private sector narratives*.

- The novel's episodes are analogous to the events that make up your private narratives for the various sectors. You can think of episodes as the incidents that come to mind when you are reminded of something that happened to you in the past. These are your *private episodic narratives*.
- The reader's narrative is analogous to your flow of consciousness as you go about your life. You can think of it as the story you are telling yourself, about what is going on, how things came to be this way, and what you expect to happen if things continue in the same vein—where all of the information derives from all your other, previously listed, narratives. This is your *private current narrative*.

(To simplify things, I am going to drop the word “private” in future references to private narratives, leaving it up to you to remember that the narratives you construct for your own use are unique to you and far richer than the versions of those same narratives that you use publicly.)

The general idea is that your memory contains your autobiographical narrative, the big picture of your life, which is made up of your sector narratives for work, relationships, health, current events, sports, politics, religion, history, and so on. Nested within your sector narratives are your episodic narratives for specific events. All of these narratives are retrievable from memory, but episodic narratives are retrieved most readily. They all influence your current narrative.

The Current Narrative

Your current narrative is the story you're telling yourself to make sense of what just happened to you, what is happening right now, and what will happen next. That is, it is partly memory, partly current awareness, and partly expectations for the future. As each second passes, as the present becomes the past, that part of your current

narrative that was the present just a moment ago becomes the past and is moved to memory where it is stored as an episode.

Perhaps an analogy would be useful: You've probably seen the "crawl" across the bottom of your TV screen, on CNN or Fox News for example. The crawl is a single line of print that scrolls from right to left under the picture. The portion of the message that you're currently reading is, of course, visible on the screen. The portion you've already read has passed out of sight on the left. The portion you're going to read has yet to come into view on the right. Messages that appear in the crawl seldom are very long, but there is no reason why they couldn't be longer—perhaps, a standard-length news story, an extended editorial, or even a whole book.

Now, think of the message on the crawl as analogous to your current narrative. The visible portion is analogous to your immediate conscious experience—what's happening right now. The portion that you've read and that has passed out of sight on the left is analogous to the part of your current narrative that has passed into your memory in the form of episodes. The portion of the crawl that you haven't read yet is analogous to the part of your current narrative about what hasn't happened yet, the future. You don't know exactly what the as-yet-unread portion of the crawl's message will be, but in the context of what you've already read and what you are reading now, you could make an educated guess. Analogously, you can use the memory and immediate experience portions of your current narrative to make an educated guess about what will happen next.

Just as what you currently are reading on the crawl only makes sense in light of what you've already read, your immediate experience only makes sense in light of what has led up to it, which is the part of the current narrative that has passed into memory. Moreover, as we shall see

in a moment, the part of the current narrative that is stored as episodes in memory provides access to similar episodic memories throughout your memory. This means that at any moment your immediate conscious experience is informed by ties to a large knowledge base of episodic and sector narratives, which expands the meaning of what is happening to you and enriches the current narrative you are constructing.

The Nature of Episodic Narratives

Episodic narratives are, in effect, short stories. They're about specific incidents and they are "stored" in epistemic memory.¹⁰ They contain information about the times, places, and emotions related to the incident, what led up to it, what happened, what you did as a result, and how things worked out. As with any narrative, episodic narratives are far simpler than the incident itself, and, perhaps to make storage more economical, they tend to become even simpler the longer they are stored in memory.

Good Narratives

In Chapter 1, the elements of narratives were identified as symbols that stand for real or imagined events and actors, where the latter are animate beings or inanimate forces. The underlying dynamic is causality and implied purpose. A good narrative is plausible and has a coherent storyline; plausible in that none of the component events, actors, and actions is outlandish or improbable, and coherent in that they are all neatly encompassed by the narrative, leaving few or no loose ends. Moreover, the more coherent they are, the more resistant they are to change. Because autobiographical, sector narratives, and older episodic narratives have withstood the test of time, they tend to be more resistant to change than the current narrative, which is still being composed.

¹⁰ To say that something is "stored" in epistemic memory isn't really accurate because epistemic memory refers to a category of memories, not to a location.

Conflicting information tends to be ignored or distorted to fit the older narratives rather than prompting their revision. This can lead to problems when change is desirable, but it also shields us from the bewildering instability that greater pliancy would inflict.

Box 2b: Narratives in Psychotherapy

In the 1960's, as the Behaviorist revolution was being replaced by the first cognitive revolution (see Box 1a), many clinical psychologists began to question the effectiveness of psychotherapeutic strategies derived from Sigmund Freud's psychoanalytic theory. From this reassessment emerged a new form of therapy, called cognitive behavioral therapy, which focuses on changing problem behaviors rather than trying to "cure" underlying psychodynamic ills. The primary methods of treatment are biofeedback, desensitization, and talk therapy. The latter consists, in large part, of helping the client understand the erroneous or contradictory beliefs that manifest themselves in their problem behavior.

Implicit in the cognitive behavioral approach is the assumption that the source of neurotic behavior, and perhaps some psychotic behavior (although this is debatable), is not an underlying illness. Instead, the problem lies in the beliefs upon which we base our decisions. When these are erroneous or contradictory, the resulting behavior fails to produce desirable results. But, trapped by our beliefs, we persist in doing things that are at best unprofitable and at worst self destructive.

If one examines what happens during a Cognitive Behavioral therapy session, it is clear that it isn't just the client's underlying beliefs that are being examined; it is the narratives that reflect these beliefs. Indeed, it is fair to say that the beliefs have to be inferred from the narratives, and this sometimes takes a good deal of time as various narratives reveal a common underlying belief or set of beliefs. The usual procedure is for the client to recount troubling incidents and other stories about his or her life, as well as fears and hopes and all the rest. Sometimes the therapist guides the monologue; other times he or she just lets it flow. In either case, the goal is for the client to recognize emergent patterns of beliefs and behaviors and to understand how they fail to deliver desirable results.

In the text, we have talked about the need for narratives to be plausible and coherent, and that is true. But, thinking that they are doesn't make it so. In fact, most narratives harbor more vagueness, wishful thinking, and contradiction than we realize. And, even those that are reasonably plausible and coherent may not be accurate. Consider a young man whose narrative tells him that he has always been fat and that he will always be fat and that being fat means that he is worthless and unworthy of anyone's regard. Even if this narrative is false—he isn't particularly fat and he certainly isn't worthless or unworthy of regard—his narrative prevails; merely telling him he's wrong isn't going to change anything. Even if he tells himself he's wrong, his underlying narrative about being fat will win; every time he looks in the mirror, his narrative will color his interpretation of what he sees and, by coloring it, reaffirm it. No surprise then that his behavior reflects his narrative: his posture, how he dresses, his relationships with women, his expectations about

how others will treat him—perhaps eventuating in eating disorders such as anorexia or bulimia. He is defeated before he begins.

We all possess erroneous narratives, some of which actually help us—slightly exaggerated self confidence or optimism, even if unwarranted, may actually be helpful. But, when erroneous narratives lead to self-destructive or futile behavior, it is reasonable to seek professional help. And, having discovered the errors in our narratives, it behooves us to attempt to correct them, which isn't always easy but in the long run it is easier than continuing to engage in self-defeating behavior.

The Impact of Others' Narratives

By the time you become an adult, you have an extensive library of narratives, all related to each other to one degree or other. A few of these narratives are rooted in your personal experience, but most of them derive from other people; the things they've told you or intentionally taught you. A notable characteristic of human beings is that we are dependent for a long time on our parents as we learn from them and from the community that surrounds us. And, what we learn is largely in the form of narratives—what is right and wrong, why things happen, what the consequences are of doing this or that. But, most important, the narratives we acquire from others tell us who we are, where we fit in the scheme of things, and what our rights and responsibilities are. By borrowing and honing and reinterpreting the contents of these narratives, we make them into our own private narratives; similar to, but different from, the narratives of those around us.

Even when their narratives have been reshaped into our narratives, we continue to be influenced by what others have to say—we are, after all, intensely social creatures. So, we listen and we revise our narratives in light of what others tell us about how the world works and what has happened and what might happen, often accepting what they say unquestioningly, other times rejecting it or accepting it provisionally or accepting only parts. As we change our narratives, we change our view

of reality, because narratives *are* our truth, they're all we've got to make sense of things.

Plausibility, Coherence, and Consistency

Before we look more closely at the impact of others' narratives on our narratives (and vice versa), recall that a good narrative is plausible and coherent. Narratives that meet these standards are satisfying and resistant to change. Narratives that fail to meet these standards arouse uncertainty and stress, both of which encourage revision of the narrative to make it into a good narrative.

Just as an implausible, incoherent narrative arouses uncertainty and stress, inconsistencies among your closely related narratives arouse similar discomfort, if only because they can lead to conflicting expectations about the future. In the course of even a few moments you switch back and forth among the narratives for the different sectors of your life—you note a spot on your arm and switch to your health narrative, in the middle of which your spouse reminds you that you have a meeting with your boss, causing you to switch to your relations-with-my-boss narrative, and so on. As a result, you constantly are juggling different narratives and correspondingly different expected futures. These expected futures have to be reasonably similar or you may find yourself utterly baffled about what to do about any of them.

Let's take this another step: Inconsistency between your own narratives and the narratives of the people around you also arouses discomfort, if only because inconsistent narratives are likely to dictate inconsistent expected futures and resulting inconsistent actions, which is likely to result in conflict. In contrast, being among like-minded people assures you that conflict is less likely. But, perhaps more important, being with people who share your narratives tends to legitimize your narratives, reassuring you that those narratives must be correct. Of course, consensus

is no assurance of validity, but we seem to think that it is. In the right company—a hate group, for example—consensus can validate a narrative that people who don't share it would think both wrong and dangerous.

Truth

For our purposes, there are two kinds of truth, empirical and conceptual, both of which inform your narratives by affirming them or by guiding their revision.

Empirical truth is the truth of science; theories are true insofar as they successfully predict events, including the results of experiments that are designed to test them. Similarly, your private narratives are true insofar as they successfully predict the future, including the results of actions you take based upon them.

Conceptual truth is the truth of literature, drama, poetry, and of social reality. It derives from the underlying narrative's plausibility and coherence; a good narrative "feels right" and, therefore, is believable. Similarly, you regard other people's plausible and coherent narratives as believable—as long as they aren't significantly discrepant from your own narratives and/or if integrating them into your own narratives increases the latter's plausibility and coherence.

Not every good narrative is regarded as true; you don't think a well written fairy tale or a well crafted movie is true because you know up front that it is fiction. But, if you think it is about actual events, a good story is very convincing. Perhaps you've heard about what happened in 1938 when the actor Orson Welles presented a dramatization of H. G. Wells' *The War of The Worlds* on the popular radio program, *The Mercury Theater*. Approximately 6 million people heard it and a substantial number believed that the earth was, in fact, being attacked by aliens. There was widespread panic and Orson Welles subsequently was sued by people

who believed he intentionally tricked them. When people are so inclined, even a patently absurd story, will seem believable.

Empirical and conceptual truths often come in the same package. For example, suppose you are trying to make some sort of change in the organization for which you work and your actions prompt a protest from your co-workers. Their protest speaks to both kinds of truth. The existence of the protest is itself empirically true; evidence that you may not be able to achieve what you want to achieve if you continue doing what you're doing. The content of the protest is a narrative about what your co-workers think of your goals and/or your tactics and why they're wrong; a narrative that may or may not be compatible with your own narrative. If it isn't compatible, you'll probably reject it as false—which doesn't change the empirical truth that it is a barrier to your plans. If it is consistent with your narrative, you'll probably accept it and incorporate it into your own narrative. Either way, you are likely to revise your plan for changing the organization or give it up altogether.

Persuasion

Instructive narratives exert their persuasive influence in two ways, immersion and teaching. Immersion means that the narrative is part of the milieu in which you find yourself, and you learn about it without any real effort. Teaching means that the narrative is conveyed to you with the intention of either informing you about its content or convincing you of its relevance to you, both of which are designed to persuade you to revise one or more of your narratives or add this one to your library of narratives.

For example, most of us learn our culture's narratives through immersion—although newcomers may have to be taught. Simply by hearing them repeatedly, we all learn the plausible sustaining myths that every culture cherishes. These myths are narratives that may or may not be factually accurate, but that's not the point. Their function is to

persuade members of the group to define themselves and the group in particular ways, to accept the principles that the group values, and to share the group's viewpoint about its place in the world.

Cultural narratives cover a large swath; religion, politics, governance, popular culture, regional identity, racial and ethnic identity, attitudes toward other members of the culture and toward minority members, attitudes toward outsiders, and so on. Any group with which you are affiliated has some sort of culture, and you either share the narratives that each of them use to explain their existence or, at the least, you don't find their narratives to be seriously inconsistent with yours. If their narratives clash with your narratives, you are likely to be uncomfortable and elect to leave—or you may be asked to leave by others who find your dissenting views unpalatable.

There is a large literature on methods of persuasion which we don't need to review here. As was stated above, most of the narratives that others convey to you are designed to persuade you to incorporate them into one or more of your narratives or to add them to your library of narratives. Advertising is a prime example, followed closely by political speeches, sermons, talk radio, newspaper and magazine editorials, and so on. But more subtle, perhaps, is everyday conversation, of which persuasion is an integral part. Sometimes the intent of conversation is merely to share information, but more often the intent is to shape your narratives and your subsequent behavior. Conversations can be thought of as negotiations between two or more people in which various narratives are introduced and their merits advanced in an attempt, with various degrees of insistence, to persuade the other participants to incorporate them into their own narratives.

Forecasts

The current narrative isn't merely about understanding the past and present, it also allows you to make educated guesses about the future.¹¹

- An educated guess about how the future might unfold if you do not make an effort to change it is called an *extrapolated forecast*.
- An educated guess about how the future might unfold if you make an effort to change it is called an *action forecast*.

Actually, both of these forecasts are projections of the current narrative. The first, the extrapolated forecast, is based on the current narrative alone and the second, the action forecast, is based on the current narrative coupled with a plan of action (Chapter 4). To keep things straight, we will use separate labels for the two, but they both derive from your ability to use your current narrative to make educated guesses about the future.

The main reason for differentiating between extrapolated and action forecasts is that they play different roles. The extrapolated forecast tells you what the future is likely to be if you let events unfold as a natural extension of the past and present. The action forecast tells you what the future is likely to be if you intervene to direct the course of events away from the extrapolated forecast's future and toward an alternative future.

¹¹ For ease of exposition, I'll talk about narratives and forecasts as though they were relatively static, but they aren't, they're quite dynamic. They constantly change as you learn more about the environment in which your future is playing out, as you come to anticipate the actions of other people, as unexpected events occur or look like they might occur, as your preferences wax and wane over time, as your actions take effect and move you toward your desired future (or fail to do so), and so on. Everything doesn't simply stop so you can update your narrative and make nice clean forecasts; you must do it on the run. The result is more like a movie than a photograph, as you update everything to incorporate changing circumstances (as well as re-evaluations of what you thought you knew a little while ago but now know you were wrong about). Throughout the following chapters, your job is to keep in mind that narratives and forecasts are far messier than I make them sound. On the other hand, remember that, messy as they may be, they are all you've got to tie your past and your present to the headlong rush of events that is your future; without them your life would be chaos.

The alternative future that your intervention is designed to achieve is called your *desired future*. Usually, your desired future isn't very different from your extrapolated forecast because the latter consists in large part of your normal, ongoing life, most of which is satisfactory (you will wake up, brush your teeth, and eat breakfast, and so on, much as you have always done). It is only when features of the extrapolated forecast are at odds with your values and preferences (discussed in Chapter 3) that you are prompted you to consider what you could do to make things turn out better (see Box 2c). Which is to say, how can you make the actual future conform more closely to your values and preferences than the extrapolated forecast's future does?

Of course, you seldom are wholly satisfied with every aspect of your extrapolated forecast. But, if only a few features of it are undesirable, and they aren't too undesirable, you usually are content to leave well enough alone. That is, you settle for the forecasted future; you continue doing whatever it is that you're doing and let the future unfold without attempting to change it much. On the other hand, if the extrapolated future has too many undesirable features, or those features are too undesirable, you're prompted to try to fix things.

Box 2c: Sources of the Desired Future

There are four sources of desired futures:
Repaired Future

This is what your extrapolated future would be if its undesirable features were repaired to make them conform to your values; which would make them desirable, thereby making the future desirable. Its advantage is that you don't have to know what you desire, you just have to know what it is about the forecasted future that you don't like—that conflicts with your values. Then you can focus your efforts on changing things so those undesirable features of the forecast are transformed into desirable features by the time the actual future arrives. The downside of this source of the desired future is that it is conservative; the future your intervention is designed to produce is seldom substantially different from the future you expect to happen anyway.

Suggested Future

Even if the extrapolated forecast's future is desirable, someone or some event, or even your imagination, may suggest an altogether different future: you're offered a new job, you come across a travel brochure extolling the virtues of Marrakech or Katmandu, you see a TV documentary that sparks a dream of being an astronaut, an actor, a stand-up comic. Whatever its origin, you treat the suggested future as a tentative forecasted future and check its features against your values and preferences. If it turns out to be more desirable than your extrapolated forecast's future, you can try to figure out how to make it happen. If you can do so, the suggested future replaces your extrapolated forecast's future; if you're not, you drop the whole idea and settle for the extrapolated forecast's future—"I'd love to go to Marrakech, but I just can't see how, so I'll just stick with my summer job at the library".

Envisioned Future

Sometimes an undesirable extrapolated forecast is unrepairable, signaling that something daring needs to be done before things get bad—or worse. When businesses are faced with this kind of need for substantial change, they try to come up with a revolutionary new vision to guide them through the substantial changes they'll have to make to break away from the trajectory they've been on before. Individuals do the same sort of thing, calling it "a clean break with the past," or "taking charge of my future," or similar words to indicate that their current narrative is about to take an abrupt and lasting change.

A envisioned future is tentative only in the sense you check to make sure it conforms to your (or the organization's) values; it isn't in competition with the extrapolated future because the latter already has been rejected.

Imposed Future.

Whenever you are subordinate to someone else, they have the right to tell you what the future is to be, frequently under the rubric of goal setting. Your job is to find a way to achieve the future that has been imposed on you. Thus, children have their futures imposed on them by their parents ("No, you can't backpack in Europe, you are going to summer school to bring up your grades!") and employees have futures imposed on them by their bosses ("I want you to hire 35 new sales clerks by the end of the month and have them trained and ready to go by the time the new store opens.") It usually is irrelevant to the person who imposes the future if it doesn't conform to your values because your role is to unquestioningly accept it and devise a plan to achieve it—that's what it means to be subordinate.

Forecast Plausibility

Forecasts are narratives and, like other narratives, the more plausible and coherent they are the more believable they are. Plausibility means that it is easy to imagine the extrapolated future evolving from the current situation in a natural, lifelike manner. Or, for action forecasts, that it is easy to imagine a potential plan of action successfully changing the extrapolated future into the desired future. That is, forecasts tend to be reasonably straightforward; they don't require an undue stretch of

imagination, they don't rely on improbable coincidences, and they don't presume the occurrence of unlikely or exotic events. In both cases, coherence means that the forecasts don't leave loose ends; they include everything you need to know to assess the desirability of the futures they offer.

Your desired future also must be plausible and coherent; straightforward and free from improbable events and complete enough to represent all of your primary values, as well as any relevant secondary values (discussed in Chapter 3). True, from time to time we all entertain thoughts about futures that are patently implausible or grossly at odds with our values. But, we usually don't take them seriously and we don't waste time or effort in pursuit of them.

Imagination

The key to forecasting is imagination. Because you have the ability to imagine, you can entertain all kinds of forecasts: What would the future be like if I didn't do anything to change it? What would it be like if I didn't do anything and then something happened (he or she marries someone else)? What would it be like if I did such and such (if I proposed and was accepted)? What would it be like if I did such and such and it didn't work (my proposal was refused)? What would it be like if I did something else entirely (ran away to join a circus)? In short, your ability to imagine alternative futures allows you to consider their desirability and how you might achieve them.

Unfortunately, less is known about imagination than ought to be—there has been a tendency to get it confused with fantasy, creativity, and innovation or to romanticize it as something mysterious and otherworldly. But the fact is, imagination is a fundamental, everyday cognitive tool that has considerable practical value—it allows you to think constructively about the future so you can intervene if you need to.

One of the most remarkable things about imagination is that you usually can discriminate between what you imagine and what you know to be real, which is what allows you to entertain alternatives as conjecture rather than fact. The importance of this discrimination becomes evident when for some reason it fails. A common example is awakening from a dream unsure whether it was real or not. Less commonly, some forms of mental illnesses interfere with the ability to distinguish reality from unreality; the patient's attempts to deal with imaginary experiences (hallucinations and delusions) as though they were real can result in behavior that seems bizarre and irrational to other people.

Scientists have studied the normal development of humans' ability to discriminate between what's real and what's imaginary. From about the age of two years onward, you were able to tell the difference. As you matured, your ability to discrimination enabled you to pretend, to make and understand jokes, to create and enjoy theatrical performances and literary fiction, to understand metaphor and analogy, and to tell lies. And, of course, your ability to distinguish between what is real and what is imagined allows you to entertain all sorts of variations on the past and future without losing your bearings and slipping into a wholly imaginary world. This in turn permits you to use your imagination productively, rather than merely as a diverting pastime:

- You can use it to consider the implications of errors in your current interpretation of events that occurred in the past (called counterfactual reasoning), thereby providing new insights about the nature of the present situation as well as revisions of forecasts of the future. "If things happened in this way rather than the way I thought they did, then this is what actually is going on and this is what I can expect to happen in the future." In short, you can

reconsider elements of your current narrative and revise accordingly.

- You also can use it to consider whether you like what your extrapolated forecast offers and, if you don't, you can imagine things you could do and the results of doing them. "If I don't do anything, I expect A will happen and I'm likely to feel this way about it. But, I imagine that if I were to do such-and-such I could make B happen instead, which would make me feel much happier."

When you imagine the future, you seldom construct an elaborate story. Instead, you see little vignettes in your mind's eye—fragmentary scenes in which you see yourself doing this or experiencing that or having such and such happen to you—all accompanied by emotions that tell you whether what you're imagining is good or bad, desirable or undesirable. These vignettes sometimes form a rough narrative, if only because they seem to arrange themselves along a time line, but they don't necessarily form a very coherent story. After all, a well-formed story would be premature; these imagined scenes are merely educated guesses about the future, which hasn't happened yet.

- Your vignettes of the extrapolated future are little scenes about what things would be like if you just keep doing whatever it is you're doing and let the future happen as it looks like its going to happen; the scenes are detailed just enough to give you the gist of what the extrapolated future might be.
- Your vignettes of your desired future are similarly incomplete. With notable exceptions, you don't usually know precisely what you want from the future, but you almost always know what you don't want, which is perhaps more valuable.

- Your vignettes of the future you'll attain by intervening in the course of events are equally incomplete; you can imagine yourself doing this or that and how it might move you toward your desired future. But, even while imagining it, you know that when the time comes many of those imagined actions may no longer be warranted and you'll probably end up doing something else entirely.

As you can see, forecasting is pretty rough and ready. But, for the most part, rough and ready is precisely the right thing, given that flux and change are all around you. Outside forces will make things happen that you never anticipated, making your extrapolated forecasts obsolete. Even your own actions change things, often in unanticipated ways, making your action forecast, if not wholly obsolete, at least imprecise. Demanding flawless forecasts and being unwilling to settle for anything less than exactly what you've forecasted would be foolish because you'd never be happy. Besides, such inflexibility would bar serendipity, forestall happy accidents, and deny you the pleasure of stumbling into new and more desirable futures than you could have ever imagined.

Time Horizons

Just as weather forecasts cover different time intervals—this afternoon, tomorrow, this week—your extrapolated and action forecasts have different time horizons. They can be short range (“If I keep spending money like I did today, I’m going to be broke by Wednesday”) or long-range (“If I keep getting mediocre grades, I won’t be accepted to med school and I’ll never become a surgeon”).

Subjectively, time and space are similar in that they both extend outward from a central point, you. Space is three dimensional and you can move attention from one location to another within it. Thus, you can look around the room, focusing your attention on your computer keyboard or a picture on the wall or a shadow on the ceiling. But

subjective space is not just what you can see or reach out and touch, it also includes Paris and Mongolia and Mars. And, just as you can move your attention around a room, you can move your attention to these distant, unseen places, confident that they exist and that their inhabitants (excepting Mars, perhaps) are going about their daily lives just as you are going about yours. In short, subjective space is huge and largely imaginary, but not fantastical.

Similarly, subjective time is huge, largely imaginary, and not fantastical. You can think about the near past (3 minutes ago) or the distant past (when you were a child), even a distant past you never knew (the signing of the Magna Carta or the fall of the Roman Empire), and you can think about the near and distant future with equal facility. This ability to focus attention on different points in time allows you to think about your forecast as an interval having its near boundary in the present moment and its far boundary at some point in the future. The far boundary, the forecast's time horizon, is determined by your estimate of when the forecasted events will have occurred. And, within the interval between now and the forecast's time horizon, you can mentally examine the details—how things will look in a moment, tomorrow, or next week as the forecasted events unfold.

Box 2d: Demonstrating Forecasting

Forecasting rests on the assumption that the past is prologue to the future, so that you can use your understanding of what has happened in the past, together with what is happening now, to forecast what will happen in the future. To demonstrate your own ability to do this, begin reading a short story (a narrative), but stop about three quarters of the way through and briefly describe how you think the story will unfold from the point at which you stopped reading. Your ending constitutes a forecast of the future for the short story, based on the assumption that the future (how the story will proceed toward its conclusion) is a logical extension of the past (the part you read). Of course, your forecast may turn out to be wrong—if only because authors usually try to keep their endings from being too obvious. And, just as you may be wrong about how the short story actually unfolds, in real life you may make incorrect forecasts, which is likely to have more serious repercussions than

being wrong about a short story. Moreover, just as you probably can think of more than one plausible ending to the short story, in real life you often can think of more than one plausible forecast of the real-life future.

To further demonstrate your skill for building narratives, return to the short story, skip forward a few pages from where you stopped reading and resume reading again. After reading a while, stop and describe (make inferences about) what happened in the skipped pages (then go back and see how accurate you were).

A Warning

In the preceding, and throughout what follows, we will talk as though you self-consciously construct forecasts and carefully compare to them to an entity called your desired future. But this is merely a manner of speaking about what is going on. As you surely know, real life isn't as clear-cut and deliberative. In fact, real experience is more like: "I don't know, something just doesn't feel right..... I just don't like the way this is heading..... its just going to get worse unless I phone them and straighten things out.....but I've got to pick up the kids at school and then stop and pick up something for dinner.....maybe I can phone after I get back from the gym or, better yet, maybe I can get Ralph to call for me; he knows guys there who might be willing to help if he asked them to.....etc."

It's difficult to discern a forecast or a desired future in this stream of consciousness, but its there if you look. The speaker apparently senses that things aren't going well and foresees them getting worse unless he or she does something, which seems to be phoning somebody about something. But, as is so often the case, timing is a problem—kids, groceries, and the gym must be dealt with first, and even then it might be best to see if someone named Ralph can do it instead. The extrapolated forecast is of a future that falls short of the speaker's desired future. The plan for correcting things consists of phoning someone to get something straightened out, but there are obstacles to implementing the plan—in this case, kids, groceries, and the gym. The revised plan consists of having

Ralph make the necessary phone call and the action forecast is that by using his influence he will produce the future that is desired.

So, in what follows, particularly when we discuss decision making in Chapter 5, keep in mind that this abstract way of speaking is simply to help us think more clearly about things, but at the price of oversimplifying them and imposing greater order and precision than actually exists. Of course, we'll be in good company because it is the nature of descriptions to be simpler and more orderly than the events or phenomena they describe.

Progress Report

If you still are a little confused about the precise nature of narratives, forecasts, and their uses, don't be discouraged, all will be revealed. The remainder of the book constitutes an extended explanation designed to give you a satisfying intuitive and intellectual understanding of both concepts. For the moment, simply note that this chapter has examined the first of the four propositions listed at the end of Chapter 1:

- *Cognitive narratives provide meaning and continuity by linking the past with the present and allowing us to make educated guesses about the future (called forecasts). We have addressed this in this chapter, Chapter 2.*
- Based on our past experience and our values, we create rules, some of which are used to make forecasts and some of which are used to evaluate the desirability of the future offered by those forecasts. We will address this in Chapter 3.
- When the forecasted future is undesirable, we construct plans that guide actions designed to promote a more desirable future. We will address this in Chapter 4.
- All decisions are about discrepancies and what to do about them. We will address this in Chapter 5.

In the next chapter we will examine the second proposition.

Sources and Further Reading

- Beach, L. R. (1990). *Image theory: Decision making in personal and organizational contexts*. Chichester, UK: Wiley.
- Beach, L. R., & Connolly, T. (2005). *The psychology of decision making*. Thousand Oaks, CA: Sage.
- Bruner, J. (1990). *Acts of meaning*. Cambridge, MA: Harvard University Press.
- Denning, S. (2007). *The secret language of leadership: How leaders inspire action through narrative*. San Francisco, CA: Jossey-Bass.
- Gerrig, R. J. (1993). *Experiencing narrative worlds: On the psychological activities of reading*. New Haven, CT: Yale University Press.
- Jungermann, H. (1985). Inferential processes in the construction of scenarios. *Journal of Forecasting*, 4, 321-327.
- Jungermann, H, & Thüring, M. (1987). The use of causal knowledge in inferential reasoning. In J. L. Mumpower, O. Renn, L. D. Phillips, & V. R. R. Uppuluri (Eds.), *Expert judgment and expert systems*. Berlin: Springer-Verlag.
- Klein, G. A. (1989). Recognition-primed decisions. *Advances in Man-Machine Systems Research*, 5, 47-92.
- Klein, G. A. (1996). *Sources of power: The study of naturalistic decision making*. Mahwah, NJ: Lawrence Erlbaum.
- Klein, G. A., Moon, B, & Hoffman, R. R. (2006). Making sense of sensemaking 1: Alternative perspectives. *IEEE Intelligent Systems*, 21, 70-73.
- Klein, G. A., Moon, B, & Hoffman, R. R. (2006). Making sense of sensemaking 2: A macrocognitive model. *IEE Intelligent Systems*, 21, 88-92.

- Klein, G., Snowden, D., & Pin, C. L. (June 2007). Anticipatory Thinking. In K. Mosier & U. Fischer (Eds.), *Proceedings of the Eighth International NDM Conference*, Pacific Grove, CA.
- Klein, G., & Crandall, B. W. (1995). The role of mental simulation in problem solving and decision making. In P. Hancock, J. Flach, J. Caird, & K. Vicente (Eds.), *Local Applications of the ecological approach to human-machine systems*. Hillsdale, NJ: Lawrence Erlbaum Associates.
- Klein, G., Snowden, D., & Pin, C. L. (2007). Anticipatory thinking. In K. Moiser, & U. Fischer (Eds.). *Proceedings of he Eighth International NDM Conference*. Pacific Grove, CA, June.
- Lipshitz, R. (1993). Decision making as argument-driven action. In G. A. Klein, J. Orasanu, R. Calderwood, & C. E. Zsombok (Eds.), *Decision making in action: Models and methods*. Norwood, NJ: Ablex.
- Mar, R. A. (2004). The neuropsychology of narrative: story comprehension, story production and their interrelation. *Neuropsychologia*, 42, 1414-1434.
- McAdams, D. P. (2006). The role of narrative in personality psychology today. *Narrative Inquiry*, 16, 11-18.
- Mitchell, W. J. T. (1981). *On narrative*. Chicago: University of Chicago Press.
- Ochs, E., & Capps, L. (1996). Narrating the Self. *Annual Review of Anthropology*, 25, 19-43.
- Pennington, N., & Hastie, R. (1988). Explanation-based decision making: Effects of memory structure on judgment. *Journal of Experimental Psychology: Learning, Memory and Cognition*, 14, 521-533.
- Pennington, N., & Hastie, R. (1992). Explaining the evidence: Tests of the story model for juror decision making. *Journal of Personality and Social Psychology*, 2, 189-206.

Polkinghorne, D. E. (1988). *Narrative knowing and the human sciences*.

Albany: State University of New York Press.

Thüring, M., & Jungermann, H. (1986). Constructing and running mental

models for inferences about the future. In B. Brehmer, H.

Jungermann, P. Lourens, & G. S. Sevòn (Eds.), *New directions in*

research in decision making. Amsterdam: North Holland.

Chapter 3: Memory, Rules, and Expectations

In Chapter 2, we discussed narratives and how they enable us to forecast the future and how the desirability of that forecasted future leads us to accept it or to try to change it. In this chapter, we will discuss where forecasts come from as well as what determines their desirability.

Memory

The building blocks for forecasts and the criteria for desirability all reside in your memory. Therefore, it is instructive to briefly review the structure and function of memory.

Heirs to more than 100 years of research results, cognitive scientists have determined that memory as a whole is divided into immediate memory and retentive memory. Retentive memory is in turn divided into procedural memory and declarative memory. And, declarative memory is further divided into episodic memory and semantic memory. Each of these divisions represents a difference in the kind of stored information. Hence:

Structure

- Immediate memory—the present.
- Retentive memory—the past.
 - Procedural memory—skills.
 - Declarative memory—information.
 - ✓ Episodic memory—specific information.
 - ✓ Semantic memory—general information.

Function

Immediate memory is closely related to focus of attention (Box 2a in Chapter 2). It is, in effect, a limited-capacity buffer that holds information about that infinitesimal instant that is “now” and the brief time preceding it. As time progresses, “now” is merely a flickering moment that is

immediately replaced by a new “now,” and then another and another, with the result that “now” actually lasts only a moment. Immediate memory extends that moment by retaining the few moments before it. This creates the brief interval that you experience as “the present.”

Because time doesn't stand still, the older (milliseconds) information in immediate memory is pushed out as new information comes in. Normally, this older information moves to retentive memory and is experienced as something that already has happened, although very recently. However, you can retain older information in immediate memory by focusing your attention on it, often by rehearsing it (as when you repeat a telephone number over and over to yourself until you have successfully made the call). The mechanism by which this is done is called *working memory*, which bridges between immediate memory and retentive memory—between the present and the past. In doing so, working memory melds information from your past with what is contained in the interval you experience as the present, thereby giving meaning to what is happening to you.

Retentive memory is about the things you have learned, either from past experience or by instruction. It consists of procedural memories and declarative memories.

Procedural memories are about skills; how to use a screwdriver, operate a computer, drive a car, read a map, handle a knife and fork, and all the other skills one acquires in the course of being alive and to which one pays little conscious thought once they are well learned. Procedural memories run the gamut from basic motor skills, such as walking, to abstract intellectual skills, such as doing long division or algebra.

Declarative memories are specific episodes that you have experienced and narratives that link multiple episodes into a cohesive, meaningful story.

Episodic memories are information about specific persons, objects, and events (“this cow is”) as well as the times, places, and emotions associated with them. They are the episodes discussed in Chapter 2—the sorts of things you can recount to someone else as an incident you experienced or that you can see clearly in your “mind’s eye” when you try to remember what happened.

Semantic memories are narratives that link your specific episodic knowledge with your general knowledge (“all cows are”) to form coherent stories of the kind described in Chapter 2. Semantic memory also contains information about the relationships among narratives; the hierarchical relationships among narratives described in Chapter 2 as well as conceptual, structural, metaphorical and emotional similarities between component elements of different narratives.

So here’s the general idea—things that are happening “now” become part of the current narrative in your immediate/working memory. Then they are copied to your episodic memory where they are stored as facts about specific incidents--episodes. After a time in episodic memory they are copied to your semantic memory where they are merged into your constantly evolving narratives, although the parts of them that involve skills are also copied into your procedural memory to add to your store of skills.

Memories as Knowledge

We tend to think of memory as a mechanism for remembering past events. But, other than perhaps for nostalgia, remembering for its own sake would be pointless. Presuming that something as highly developed and so central to our experience (to say nothing of all the brain tissue

devoted to it) isn't pointless, it is reasonable to assume that memory is a mechanism for accumulating useful knowledge. For our purposes, knowledge is useful when it allows you to formulate rules that will:

- Predict the likely results of your (others', nature's) potential actions.
- Predict how desirable you will find those results to be, should they occur.

Rules for predicting the likely results of actions you perform are called *contingent rules*. Rules for predicting the potential results of actions that other people or nature performs are called *noncontingent rules*. Predictions are necessarily conjectural, hence these Action _ Result, cause and effect rules allow you to conjecture about the future (forecasts) and to conjecture about how you might change that future (plans).

Rules for predicting the desirability of the likely results of your, others', or nature's potential actions are called *normative rules*. They predict how happy you'll be if the actual future turns out the way you've forecasted or planned. These rules set standards for what constitutes a desirable future and, in doing so, set standards for what constitutes acceptable actions.

Contingent and Noncontingent Rules

Contingent rules are for predicting the results of actions that you are considering taking. They sometimes derive from an episode in episodic memory—"I did this last time and it worked, I'll try it again and expect it to work again," but more generally they derive from knowledge in procedural memory about how to make things happen. In either case the knowledge is based on previously successful actions that produced particular results and, in the case of procedural knowledge, on instruction about what to do to produce particular results or on your ability to

imagine plausible ways in which you can intervene to produce particular results.

Contingent rules tell you *what to expect* as a result of something you do: If you're nice to people you expect that they'll be nice to you, if you work hard you expect to be promoted; if you exercise you expect to get healthier; if you pay your taxes on time, you expect to avoid paying a penalty. Contingent rules are of the "If _ Then" form; "If I do X, then Y will happen." The dependability of the rule, together with the amount of effort you intend to devote to its implementation, tells you how confident you can be that Y will happen if you do X.

Because there always is a question about whether your actions will do what you expect them to do, despite your best efforts, contingent expectations are about what *could* happen as a result of your actions, rather than about what necessarily will happen. "Could" recognizes the rules' fallibility and the possible ineffectiveness of your efforts in carrying out the action. Contingent expectations evoke words like "success" and "failure," but they also allow leeway, as in "close enough" and "best I can do under the circumstances." Your appraisal of the likelihood that your contingent rule will correctly predict the results of your action reflects the *strength* of your contingent expectation.¹²

Noncontingent rules derive from semantic memory and are about your understanding of both social and physical processes. They are based on co-occurrences of events you've observed in the past, from instruction (e.g., classes, from reading, from watching people), or from your ability to imagine how one thing might lead to another.

Noncontingent rules are about things over which you have no control (when there's a downturn in the housing market, you expect the stock market to fall); about natural processes over which you have no

¹² Contingent rules are what is learned in operant conditioning, but they usually are more complex than a rat's simple rule about pushing a lever to receive food.

control (when you see lightening, you expect to hear thunder), or about things that you cannot control simply because you don't understand what causes them and therefore must simply accept that they may or may not happen (every Spring I expect to catch a cold).

Like contingent rules, noncontingent rules aren't foolproof—sometimes a housing downturn isn't followed by a fall in the stock market, sometimes you see lightening without hearing thunder, some years you don't catch a cold. But, their lack of infallibility doesn't make the rules useless; it just means you have to be careful. Because you know so much about the context in which you observed the events upon which you based the rule (or about the credibility of your instructors or the plausibility of your imaginings), and because you know so much about the circumstances in which you plan to apply the rule, you often can make an informed guess about whether it will correctly predict the outcome of interest (event Y). That is, you can estimate how likely it is that the rule will correctly predict Y in this particular situation at this particular time.

Noncontingent rules, like contingent rules, are of the form, "If _ Then." That is, if event X occurs, then you *expect* event Y to occur: Downturn in housing market _ stock market falls, lightening _ thunder, Spring _ colds. These rules set expectations about what *ordinarily* happens in the normal course of events. "Ordinarily" recognizes that noncontingent rules aren't inevitably right but that there are regularities that can be used for making forecasts. Noncontingent expectations evoke words like "probability" and "likely," and "a good chance that..." As with contingent expectations, your appraisal of the likelihood that relying on your noncontingent rule will allow you to correctly predict the results of

others' and nature's actions reflects the *strength* of your noncontingent expectation.¹³

Determinants of Strength of Expectation

The strength of a contingent or a noncontingent expectation will be greater if the rule has seldom or never been wrong in the past, if you believe that whoever taught you the rule knew what he or she was talking about, if you can easily imagine the rule working, and, for contingent expectations, your determination to make the rule work. The strength will be weaker if you have seldom seen examples of the rule working, if you doubt your instructor, if the rule is counterintuitive, or, for contingent expectations, if you lack determination or lack the skill to make the rule work.

Creating Forecasts and Plans

Creation of an extrapolated or action forecast begins with immediate/working memory, which contains your current narrative about what is happening. Extrapolated forecasts are based solely on noncontingent rules; action forecasts are based on both noncontingent and contingent rules.

- Extrapolated forecasting is done by applying your noncontingent rules to the current narrative to predict how it will be transformed by the anticipated actions of other people and nature, with no interference from you.
- Action forecasting is done the same way except that, in addition, contingent rules are added to the process to predict the transformative effects of your own contemplated actions. That is, an action forecast is simply an extrapolated forecast to

¹³ Noncontingent expectations are similar to, but not the same as, the probabilities in the Expected Value equation in Box 1a in Chapter 1. In both cases the focus is on events over which you have no control. Noncontingent rule formation is what happens in classical conditioning; the dog learns to expect food when it hears a bell; it doesn't do anything to get fed, it merely salivates in anticipation.

which are added the predicted effects of actions you plan to take.

We'll discuss plan creation in Chapter 4. Suffice it to say that once you have a plan, you use the anticipated results of the tactics to generate an action forecast and appraise the desirability of the future offered by the forecast. If the offered future is desirable, you implement the plan. If the offered future is undesirable, you must try again.

Normative Rules

Normative rules dictate what is and what is not desirable. They are acquired through instruction and from your conclusions about what is ethical, right, proper, principled, reasonable, appropriate, preferred, and so on; all of which are collectively called your *values*.

Normative rules are of a different form than noncontingent and contingent rules, both of which are of the "If _Then" or X _ Y form. This is because with normative rules you *expect* Y to happen simply because it is moral, ethical, right, proper, polite, and because you want it to happen, not because of any prior event, X, or because of anything you do. In short, there isn't an X and there isn't an _, there's only Y, which isn't contingent on anything, it just is or ought to be. As with noncontingent and contingent expectations arise from their respective rules, normative expectations arise from normative rules and are the degree to which you expect the rule to hold.

Contingent rules and expectations are specific; a rule that works for predicting the behavior of your spouse may not work in predicting the behavior of your boss, a rule about what happens if you work hard at the gym may not work in predicting what happens if you work hard at your job. Noncontingent rules are general, but even they are only applicable to particular classes of persons, objects, or events. A rule that works for predicting the behavior of adults may not work for children, a rule about

how employees react to motivational messages may not work for volunteers, a rule about how to produce a particular kind of chemical reaction may not work for another kind of chemical reaction. In contrast to both of these, normative rules and expectations apply broadly across many classes of persons, objects, and events. Because you regard these rules as universal, you also regard them as imperative and peremptory; they set standards for how things *should* be, and therefore for what is and is not *desirable*. That is, normative expectations are about standards for your own and others' behavior and about what constitutes a desirable state of the world.¹⁴

“Should” reflects the absolute nature of normative rules. Rules and expectations based on your values don't grant much leeway. Normative expectations almost always have a judgmental tone, evoking words like “good” and “bad,” and phrases like “almost isn't good enough” or “close but no cigar.” As with noncontingent and contingent rules, your appraisal of the likelihood that your normative rule will predict the outcome, (in this case, that Y will conform to your rule and therefore to your values) reflects the *strength* of your normative expectation.

What is normatively imperative to you isn't necessarily normatively imperative to everyone else, so not everyone is going to desire what you desire or behave as you think they should behave. As a result, however much it pains, disillusion, or outrages you, you ordinarily recognize that there is less than complete certainty about whether a desirable outcome or behavior that should happen will in fact do so. Phrases like “I hope they will do it,” or “They shouldn't, but I wouldn't be surprised if they did,” and so on, reflect this lack of confidence—usually accompanied by a trace of disapproval. (Of course, you also know that you'll fall short of your own

¹⁴ The Ten Commandments is a set of normative rules for your own and others' behavior. They are imperatives, which is why they aren't called The Ten Suggestions.

standards from time to time, which one might hope would prompt a bit of compassion for other transgressors, but it usually doesn't.)

Strength of Normative Expectations

The strength of a normative expectation will be greatest for a rule that derives from your fundamental, bedrock, primary values, especially if you believe the values are broadly held by the people around you. The strength will be less for rules deriving from more peripheral, idiosyncratic, secondary values, including preferences, because you are less sure of their universal rightness and you know it's unlikely that everyone shares them, so they are less likely to be predictive of what others will do or even what nature will do.

Testing Desirability

Our interest in normative rules stems from their use in testing the desirability of forecasted futures. Testing is done by comparing each of the defining features of the forecasted future with standards about what those features *should* be, based on your normative rules. The forecast is rejected as undesirable if there is a significant difference between the forecasted future's features and your normative standards because a large difference implies that the future, when you get there, will not please you. The mechanics of how this test is performed is the subject of Chapter 5.

Values and Emotions

Because they determine desirability, and therefore hold sway over nearly everything we think and do, we must examine normative expectations more closely. As we have seen, normative rules are different from noncontingent and contingent rules in both form and substance. Different in form because normative rules are not of the "If _ Then" form because there isn't any "If." They are different in substance because normative expectations are not about how results are produced (the

results of others' actions, of natural processes, or of your efforts); they are standards against which to judge the desirability of other's actions, natural processes, or your efforts themselves as well as the desirability of the anticipated results of those actions.

Primary and Secondary Values

Your standards derive from your *values*, which consist, collectively, of your ethics, ideals, and your ideas of equity, justice, solidarity, stewardship, truth, beauty, and goodness, together with your moral, civic, and religious precepts and the responsibilities you assume in the course of performing your daily duties and engaging in social interactions. In addition, they consist of your wants and preferences, which often are no less compelling than the previously listed values even though they tend to be transitory. Collectively, your values dictate what is desirable and what is not, and how much you desire or despise them.

Some of your values are supremely important, called *primary values*, and some are less so, called *secondary values*. Primary values are broadly applicable, enduring, and interconnected, both with each other and with secondary values. They are central to your conception of the world and to your existential beliefs. You take them to be "basic truths" about social reality, about ideology, and about yourself. They are never trivial or simply matters of taste, and they are resistant to change.

Secondary values are in large part your preferences and tastes. They wax and wane in intensity (you want to go to Paris, but having done so you want to do something else) and, if challenged by experience or conflicting information, they are less resistant to change than primary values.

Even though primary values are more imperative than secondary values, they both dictate what you think that you and others should or should not do, or even should want to do, and they dictate what

constitutes a desirable state of your physical and social environment. That is, they set the standards for your own behavior, the behavior of others, and the behavior of nature itself—a tall order, indeed. Small wonder then, that our normative expectations are seldom fully met.

Origins of Values

Values are largely cultural, even though each individual holds a private version. Sociologists have long argued that shared cultural values, called social norms, are what hold societies together. Indeed, within a culture there is a good deal of agreement about values, and there is surprisingly large agreement among different cultures. It is argued that to function successfully within a culture, be it only two people, as in a marriage, or society in general, it is necessary for the participants to be able to assume that others share their values—that there is common ground. But seldom is each individual's private version of these communal values precisely the same as anyone else's. It is this similarity without identity that makes it possible for us to function in a diverse society while allowing us to be individuals instead of cultural clones.

To say that an individual has private values is not to say that he or she is capable of clearly identifying them (see Boxes 3a). Some values are fuzzy, although they remain powerful—most of us somehow believe that freedom is important enough to die for even though many of us aren't quite sure what freedom is. Other values are explicit, if only because they are broadly discussed; most religions have some version of the Ten Commandments and most children are taught about them—most people in the world strongly believe that we shouldn't kill each other or steal or lie.

Primary or secondary, explicit or fuzzy, communal or unique, values all function in the same way; they are imperatives that constrain every thing we do because they define what is and what is not desirable.

Box 3a: Cues for Inferring Values

Most people have difficulty describing their values, in large part because they acquired most of them through observation and example rather than through explicit instruction. But, there are four cues for inferring them from how you or others describe value laden events.

- Expression of obligation (from responsibility to duty) through use of words like “ought,” “should,” “guilt,” or “shame” when describing things you, another person, or an institution did or are considering doing, including the ends you or they are pursuing or considering pursuing.
- Expression of approval (from acceptance to praise) of something you, another person, or an institution did or are considering doing and the ends you or they are pursuing or considering pursuing.
- Expression of fear (from anxiety to terror), anger (from distress to rage), or anticipated regret (from mild to intense) about something you did or are considering doing, about something someone else or an institution did or is considering doing, or about the ends that you or they pursued or are considering pursuing.

Values μ Virtues

Make no mistake, not all values are virtues. We each value things that others would regard as evil, irrational, misguided, or all three. For some people, greed is a primary value and underhandedness is relished for its own sake. Some people disapprove of abortion on moral grounds, and others approve of it on other moral grounds, and both sides are convinced that they are right. When primary values are involved, the fact that other people don't agree with you is of little consequence; you are always right (although you may express it more diplomatically) and anyone who holds conflicting values is flat out wrong. After all, if they depended on other people's endorsement, your primary values wouldn't be primary at all; they'd be secondary and more on a par with tastes and preferences.

Just as values aren't always virtues, they also don't always prevail. You like to think that your values are rock solid, but you know full well that you have failed to live up to them in the past and are likely to do so again. The big problem isn't that you don't have values, or that you hold

them in so little esteem that they don't influence your behavior, it is that you sometimes fail to realize (or perhaps acknowledge) that they apply in a particular situation. In business, for example, there often is a fine line between shrewdness and villainy. It is distressingly common to hear people justify the most outrageous actions with "That's business"—actions that they would never even consider in any other sector of their lives and that probably would infuriate them if someone else behaved that way toward them. A second reason that you sometimes fail to apply your values is because the opportunities offered by the situation are so attractive that they simply swamp your resolve; opportunities and social pressure can overwhelm even primary values in the short run and can corrode them in the long run.

However value-blind you can sometimes be, your primary values are more likely than not to prevail. Few of us would kill someone, or steal anything of significance, or commit perjury. And, if we ever encountered a situation that was so extreme that we were driven to any of these actions, we would feel profoundly guilty, whatever the mitigating factors. Even little violations cause a twinge—you know when you're lying, even if it is to keep from hurting someone's feelings. You know that walking out of a store with some minor item you didn't pay for is theft, even if it was the cashier's fault for overlooking it when ringing up your order. Little things like this may arouse only a little regret and perhaps a little guilt, but that little is often enough to keep you from getting too far out of line—the twinge is a warning that you're pushing the boundaries, that this isn't a path you want to go down. Having said this, it remains regrettably true that we all tend to hold others to a higher standard than we hold ourselves. We are scandalized by things that others do, but when we do them ourselves, they often seem the most natural thing in the world.

Values and Meaning

Whether honored in the observance or in the breach, values give meaning to everything in your life. It is common to talk of action as the pursuit of outcomes that leave you better off than you were, but little is said about what “better off” actually means. In fact, it means that the pursued outcomes hold promise for promoting or complying with your values—which is why you pursue some and ignore or avoid others. Getting things, doing things, making things happen gives intrinsic pleasure when it complies with and promotes your values. For example, you work hard to make money, but money is little more than an electronic blip in your bank’s computer or a piece of paper in your wallet. The real reasons you work are to satisfy your work ethic (a value) and for the intrinsic pleasure your job provides (a value), as well as to make it possible to live the kind of life you want (value) and to do the kinds of things that contribute to your happiness and the happiness of those around you (values). Values, particularly primary values, are the bedrock upon which you build your life.

Because values are so central and because virtually everything you do is aimed at promoting them, it is important that your efforts be successful. You therefore try to do a good job when designing your forecasts and plans, especially when primary values are at stake. As a result of having tried to do a good job, you generally expect your forecasts to be reasonably accurate and you expect your plans to work reasonably well—if you didn’t, you’d immediately revise them. In short, you *expect* your efforts to successfully satisfy your values, although you fully understand that even the best intentions sometimes fall short of your expectations.

Value and Emotion

When you make forecasts, the future you forecast can meet, exceed, or threaten to violate your value-based normative expectations. All three possibilities generate emotions that tell you about the desirability of the forecasted future and that motivate action on your part to avoid or change that future if it is undesirable.

Met Expectations. When the forecasted future meets your normative expectations, all is well—because what happened was what was supposed to happen. Even if they aren't precisely met, close is usually good enough. You must tolerate a little error because the world is in flux and therefore isn't wholly predictable. Just how much leeway you're willing to allow depends on the strength of the expectations—on whether the values are primary or secondary.

Exceeded Expectations. When the forecasted future exceeds your normative expectations, you may be pleased or you may not. Outcomes that exceed your expectations by too much can be nearly as disruptive and unpleasant as outcomes that fail to meet them. This is because you may not be able to absorb them into your narrative and maintain your equilibrium. You probably value money and it may seem like it would be nice to win the lottery, but if it actually happens you may in fact find your life changing in very undesirable ways. You might think you'd simply adjust and enjoy your good fortune, but the example of big lottery winners suggests that this isn't as easy as it sounds.

Threatened Expectations. When the forecasted future violates your normative expectations it really means that your values are only threatened—because the future hasn't happened yet. However, the threat alone engenders strong emotions that motivate you to do something. This can consist of blunting or lowering your expectations, leaving the scene so you escape the threat, or taking action to thwart the threat.

Threats to normative expectations that derive from your secondary values may be so unimportant you can ignore them or simply lower the expectations to accommodate the threats—you seldom feel good about making such compromises, but you often do it because it isn't worth the effort to make everything as perfect as you'd like. Threats to expectations based on primary values, can't be so easily accommodated; so, if you feel powerless to change things, you have to physically leave the scene so the threat is mooted or simply stick your head in the sand and deny that there's a threat—neither of which is particularly practical. But, when you can see a way of thwarting a threat by heading off the forecasted future before it happens, you usually spring into action, formulating a plan (or re-formulating a faulty plan) aimed at producing a future that conforms to your values—a desirable future. Once you have an acceptable plan, you set about implementing it in pursuit of that desirable future.

Expectations and Emotions

Expectations give rise to emotions. Met expectations give rise to feelings of security and contentment. Exceeded expectations usually give rise to feelings of elation and happiness. But the emotions engendered by threats to expectations are of most interest to us because they are the ones that warn us that a forecasted future threatens to be undesirable.

Emotional reactions to threats to normative expectations can vary from being mildly uneasy when the potential violations are limited to secondary values or when they are few in number and magnitude, to being extremely upset about violations of primary values or large numbers or large magnitudes of violations. Beyond mild uneasiness, the emotions most commonly associated with threatened normative expectations are:

- *Fear*. Low level fear is called anxiety, and high level fear is called terror; when the level of fear is somewhere between it these extremes it is called everything from concerned, to worried, to frightened. The more fearful you feel, the more motivated you will be to do something to ensure that the future doesn't turn out as forecasted.
- *Anger*. Just as fear can range from anxiety to terror, anger can range from annoyance to rage. As with fear, the stronger the emotion the greater your motivation to make sure the undesirable future doesn't occur.
- *Anticipated Regret*. Regret is a feeling of loss, of disappointment about what might have been (some researchers regard it as a form of anger). Your ability to anticipate the regret you would feel if you were to let a undesirable forecasted future happen motivates you to take steps to prevent it.

When they occur, the magnitude of your feelings conveys, almost instantly, that the future isn't going to turn out well unless you intervene. As such, feelings are the initial motivators for interventions—and interventions require a plan of action, to which we turn in the next chapter, Chapter 4.

Progress Report

In this chapter we have examined the third of the four propositions listed at the end of Chapter 1:

1. Cognitive narratives provide meaning and continuity by linking the past with the present and allowing us to make educated guesses about the future (forecasts). We addressed this in Chapter 2.
2. *Based on our past experience and our values, we create rules, some of which are used to make forecasts and some of which are used to evaluate the desirability of the future offered by those forecasts.* We have addressed this in this chapter, Chapter 3.
3. When the forecasted future is undesirable, we construct plans that guide actions designed to promote a more desirable future. We will address this in Chapter 4.

4. All decisions are about discrepancies and what to do about them.
We will address this in Chapter 5.

In the next chapter we will examine the third proposition.

Sources and Further Reading

Beach, L. R. (Ed.) (1998). *Image theory: Theoretical and empirical foundations*. Mahwah, NJ: Lawrence Erlbaum Associates

Byrne, R. M. J. (2007). *The rational imagination: How people create alternatives to reality*. Cambridge, MA: MIT Press.

Myers, D. G. (2002). *Intuition: Its powers and perils*. New Haven: Yale University Press.

Myers, D. G. (2007). The powers and perils of intuition: Understanding the nature of our gut instincts. *Scientific American Mind*, 18 (3), 24-29.

Kahneman, D (2003). A perspective on judgment and choice: Mapping bounded rationality. *American Psychologist*, 58, 697-720.

Plessner, H., Betsch, C., & Betsch, T. (Eds.) (2007). *Intuition in judgment and decision making*. Mahwah, NJ: Lawrence Erlbaum Associates.

Chapter 4: Plans

Your current narrative plus your extrapolated forecast constitutes a story about how you expect the past and present to evolve into the future. If the forecasted future is desirable, because it satisfies your values, you can simply keep doing whatever you're doing and let it happen.

If, however, the forecasted future is not desirable, you must do something. You can start by re-evaluating your current narrative and the extrapolated forecast to see if the undesirable future really is as undesirable as you at first thought. If it is:

- You can try to escape it.
- You can simply hunker down and endure it.
- You can actively intervene to change it.

Escape

There are two ways to escape the future—physically and emotionally. Physical escape involves actually getting away from the unacceptable future, going somewhere that precludes it from occurring. Of course, this merely means that some other future will occur and the question is whether it is better or worse than the one you're trying to escape. Examples of physical escape are things like quitting a dead-end job, leaving an abusive relationship, or immigrating to another country in search of a better life—all attempts to avoid what appears to be an undesirable future that you perceive yourself powerless to change if you stay in the situation.

Emotional escape from an unacceptable future involves ignoring, denying, or suppressing your misgivings. Emotional escape occurs when the cost of physical escape is higher than you're willing to pay. Examples are things like staying in that dead-end job, while kidding yourself that it is bound to get better. Or, staying with an abusive spouse because you convince yourself that he or she really is sorry and it won't happen again.

In short, emotional escape is a way of dealing with a bleak future by, in effect, blinding yourself to how bad it is and how bad it will continue to be.

Endure

Enduring a grim future often looks courageous, and sometimes it is. Staying home to take care of an aging, incapacitated parent may have its rewards, but it also can be difficult and trying; someone who does it in full knowledge of how hard it is going to be is to be admired. Similarly, having to fire half your workforce and totally restructure your organization is never easy, if only because you'll feel awful about the firings and uncertain about whether restructuring will save the day, but somebody has to buckle down and do it. On the other hand, endurance often gets more credit than it deserves because it simply is a way of avoiding the uncertainty of taking action to change things. It isn't heroic to endure a rotten job or relationship simply because you dread confrontation or you fear change.

Intervene

If you can't escape the extrapolated future or don't think escape will solve anything, and if you don't feel able or obliged to endure it, you must intervene to take control of it. You can do this in either of two ways, automatically or deliberately.

Automatic intervention occurs when you recognize a situation as being very familiar, either because it was very recently encountered and still is in episodic memory or because it is commonly encountered and you have a ready response to it (a habit) in procedural memory. In the first case, you can do what you did last time if it worked, or modify what you did last time if it didn't work. In the second case, you simply use your habitual response.

Deliberate intervention is more complex. It requires you to think about what makes the forecasted future undesirable; the features that are discrepant from your values. Then you must think about what you must do to change those features so that when the real future arrives, its features will not be discrepant from your values, making it desirable. Your action agenda for making the necessary changes is called a plan and execution of that plan is called *implementation*. The anticipated result of implementing the plan is called the *action forecast*. An acceptable plan is one that offers a desirable future.

The Structure of Plans

A plan is based on the expectation that without intervention the future will unfold pretty much as extrapolated but that by appropriately intervening you can make it more desirable. It consists of a sequence of potential actions designed to influence crucial junctures in the unfolding course of events in order to transform what you otherwise forecast to be an undesirable future into a desirable future.

Even as it is being implemented, a plan is constantly being updated and changed as its effects take hold, as well as to correct for forecasting errors and to work around unexpected changes in the situation caused by the actions of other people and natural phenomena. This makes the plan quite flexible, which allows you to take advantage of unforeseen opportunities and proceed in new and unanticipated directions, sometimes achieving an even more desirable future than the one you began to work toward. (And then, it sometimes works the other way and you have to settle for compromise and second best.)

Tactics

The action components of plans are called *tactics*, which are the contingent rules discussed in Chapter 3. Recall that these rules derive from instruction and from your past experience and are stored in your

procedural memory. The plan is a provisional list of these rules; provisional because when you are constructing it you have yet to create an action forecast with which to determine if the future it promises is desirable. If, when you make this determination, the action forecast's future turns out to be undesirable, you have to make changes in the plan until its action forecast becomes desirable. That or start all over again on another plan.

Tactics are where the rubber meets the road; rules that tie the concrete things you propose to do (or have someone else do) to tangible results that will shape the course of your unfolding future. When you undertake implementation of the plan, each tactic will be executed at the cutting edge of the advancing present; that moment at which the future becomes now. And, even as the tactic is being executed, its results are being anticipated so the next step will dovetail smoothly; the next step being the next tactic in the plan or a substitute that will keep things moving toward your desired future (see Box 4a).

Box 4a: Tools


Tactics usually require tools for their execution. A tool is anything that extends or expands your physical or mental effectiveness in support of a tactic. Examples are drills and saws that allow you to do things to wood that you couldn't do with your bare hands; telephones that allow you to communicate beyond shouting distance; computers that allow you to access information that isn't in your head or on your bookshelf; formal paradigms (e.g., arithmetic, algebra, geometry) that you can apply to a wide variety of problems that differ in content but are similar in structure; and recipes and instructions that guide you as you perform specific tasks, such as baking a cake or replacing an electric socket. Each of these tools supports the tactic in its role as a component of a plan.

Tool use usually requires paradigmatic thought; there are prescribed ways of using them and rules about when they are to be used, and these have to be learned from someone who knows about them. And, like all paradigmatic thought, the goal is to provide needed input to narrative thought. Thus, I use a recipe to bake a cake that is part of a narrative about someone's birthday; I use a saw to cut lumber to remodel my front porch as part of a narrative about making my house more attractive; I use a computer to access the internet to obtain information I need to write this book, which is itself a narrative.

Sequence

At the time a plan is being constructed, it usually is rather vague. You may know the key tactics that will be required, but the others will have to be filled in as circumstances dictate. That is, tactics must be adapted to changing local conditions at the time they are put in motion, and thus they can't all be selected ahead of time. The result is that, in its early stages, a plan looks like the diagram in Figure 4.1. Some tactics are fairly clearly needed and are well defined, designated T_i in the diagram. Some are less well defined but will become so as the time approaches for their execution ("I'll cross that bridge when I come to it"), designated $T_?$. Some tactics, designated T_{cont} , can only be executed after others are because they are contingent on the results produced by the earlier tactic. Some must be executed simultaneously, designated $T_{simult1}$ and $T_{simult2}$. Others, designated T_{alt1} and T_{alt2} , are alternative actions, the one that will be used depends on conditions at the time but both are held in readiness. E_{yes} designates the occurrence of an event over which you have no control, and E_{no} designates the event's nonoccurrence. Your knowledge about these events, E , over which you have no control, derives from your noncontingent rules.

Figure 4.1: A diagram of a plan showing some different kinds of tactics.



The plan in Figure 4.1 says: You will Start implementation of the plan by executing tactic T_i , then depending on what happens, you will execute the as-yet-unclearly defined tactic $T_?$. Next, you will execute both $T_{simult1}$ and $T_{simult2}$ at the same time. Depending on the results of what you've accomplished thus far, you will execute either T_{alt1} or T_{alt2} . Then you will see whether some externally controlled event has or hasn't occurred, E_{yes} or

E_{no} , and, contingent on whether it did or didn't, you will execute either T_{cont1} or T_{cont2} . Finally, you will do what is needed to tie up loose ends, T_2 , and then execute the final tactic, T_l , that will take you to your Goal—which is your desired future.

The point of the diagram in Figure 4.1 is that even fairly comprehensive plans are not carved in stone—they're very flexible. What the diagram can't depict is the additional flexibility that derives from your ability to select tactics in midcourse to replace tactics that are no longer appropriate or that didn't work. Neither can it depict the fact that the specific actions involved in a tactic's execution always are molded to fit existing conditions—both the conditions produced by the plan implementation itself and the conditions produced by the ever-changing world in which implementation takes place.

Sources of Tactics

We've been talking as though you have a box of tactics tucked away in your procedural memory and that you trot out individual tactics to use in your plan. This isn't quite accurate. Rather, when your extrapolated forecast is undesirable, you access memory to retrieve episodic narratives for situations that are conceptually, structurally, metaphorically, or emotionally similar to the current situation. One of the retrieved episodic narratives is likely to be what happened the last time you were in this same situation, or one very like it, and what you did to make its future desirable. If what you did last time worked, you can "recycle" it for use this time. If it didn't work, you can modify it or create a wholly new plan using tactics from your procedural memory.

Box 4b: Recycled Plans

When you find yourself headed toward an undesirable future and want to avoid it, you use the parts of the current narrative that have passed into memory to retrieve conceptually and structurally related episodic narratives from throughout your memory. The retrieved episodic narratives tell you what you did in the past and you can, if it looks like it will work, simply recycle a plan from among them that will be appropriate for the present situation.

Recall that the plan you recycle seldom is the same as it was when you previously set out to use it; it was modified as you implemented it. This means that you usually are recycling plans that worked—eventually. They may not have produced the precise results you wanted when you began their implementation, but what they ended up producing was good enough for you to regard them as successful.

Of course, implementing the recycled plan probably won't produce results identical to those produced before because the circumstances aren't identical, but the point is to reach an acceptable end, not always the one you started out to attain. Moreover, because you will amend and modify it as you go along, the recycled plan won't look much like its old self by the time you get through.

The downside of recycling is that you may unwittingly recycle an inappropriate plan. Simply having a plan at all seems to influence your perception of how appropriate it is; hence the old saying, "When you have a hammer, all problems look like nails." You've seen this often enough, the manager who has a one-size-fits-all plan consisting of a torrent of angry raving followed by mindless slogans and inspirational blather. Common sense suggests that application of a one-size-fits-all plan isn't likely to work (angry raving and mindless slogans don't do much anyway, but that's another story). One reason a broad background of experience is required for most managerial jobs is because, in contrast to novices, experienced managers have access to a larger network of relevant episodic memories on which they can draw.

When our network consists of only a few episodic narratives, we sometimes believe we know what to do when we actually don't. If we proceed with what we think we know, we can inadvertently make a problem worse. Here's an example in the form of a cautionary tale:

Once upon a time there was a bright, well organized, and personable woman, whom we'll call Edith, who decided to go to work after the last of her three children went off to college. She started as a clerk in the shipping department of a small manufacturing company and was immediately recognized as someone who was exceptionally able and reliable. As a result, when her supervisor was promoted to a new job, she was offered the vacant position in preference to others who had worked there longer. She accepted, but with considerable trepidation. First of all, she had only been working a short time and she wasn't all that comfortable yet. Second, the job required her to supervise 13 other people, most of whom had been working since they were old enough to hold a job and many of whom had held their present job for years. Third, some of the people she had to supervise had themselves hoped to be offered the supervisory position.

She had barely been in her new job a week when the first crisis came. Two of her supervisees started arguing about who was supposed to do what and soon were shouting at each other about respecting boundaries and so on. As supervisor, it was up to Edith to settle the argument. But, her first impression

was that the combatants were testing her authority as much as they were actually fighting, and they were doing it for all the others to see.

Although she didn't know much about dealing with supervisees, Edith had raised three children. To her it was immediately clear that she had experienced this situation before—her children frequently had tested her authority with similar fights. As a result she didn't even have to think about what to do, she knew.

Within minutes of hearing the yelling and arguing, Edith marched over to the combatants and ordered them to shut up. She began by coldly berating them for acting like children instead of the responsible adults they were supposed to be. Then she sneeringly asked why they hadn't ask her about task assignments instead of resorting to yelling and arguing, why they couldn't think of a more reasonable way of solving their disagreement, and so on; rhetorical questions, not real requests for information.

She purposely did all of this in front of the entire crew—if these troublemakers wanted an audience, she'd give them an audience. She finished by assigning the feuding parties to tasks at opposite ends of the shop, and telling them in no uncertain terms that she didn't want to hear a word out of them. Then, while everyone looked on, with heads hung low, the culprits shuffled off like martyrs to their new assignments; but, though Edith didn't see it, there was a glint of victory in their eyes.

Newspaper columnists often talk about the Law of Unintended Consequences: Even actions that accomplish what they're supposed to accomplish can have repercussions that are far worse than the original problem. Edith was about to learn about this law. The martyrs' co-workers immediately took up their cause, filing protests with Edith's boss, starting ugly rumors about why she got preference for the supervisor job, ignoring her instructions or only grudgingly complying. She ate lunch by herself because nobody wanted to look like they were on her side. Sensing the hostility, she began to spend less time supervising and more time on paperwork.

If she had had more experience, if she had had the chance to observe a proficient supervisor at work, if she had even given it a little thought, Edith would have known that treating adults like children humiliates them, and doing it in front of their co-workers greatly compounds the humiliation. Even if they brought it on themselves, the humiliation invites sympathy and coworkers begin to take the victim's side. And none of them soon forgets what was done to the poor victims, nor do they soon forgive it.

Alas, poor Edith, her first mistake was to identify the workplace conflict as something she'd seen before in her previous occupation—mother. However, given that she had nothing else in her library of episodic narratives that might fit, she must be forgiven. Her second mistake is less forgivable; had she thought about it, she would have realized that, humiliation, the solution she had found useful for her children's fights, was inappropriate for dealing with supervisees (or with children, for that matter.) Because it came so quickly to mind and because everything seemed to fit so well, Edith never even paused to consider; she just went with the familiar, the tried-and-true. And, it worked—she stopped the fight and demonstrated that she was strong. But, the unintended consequences made her victory hollow. She soon gave up and began looking for a job elsewhere.

The tactics in procedural memory are based upon numerous previous encounters with forecasts that had these flaws or from instruction about how to deal with forecasts that have these flaws. That is, they are contingent rules about how to change the flawed features of forecasted future so that when the real future arrives it will conform to your values and, thus, be desirable. The rules are retrieved from procedural memory by using the flawed features of the extrapolated forecast as probes, so that you don't have to consider every rule in your memory each time you attempt to retrieve one.

Plan Plausibility

A plan is, in fact, a narrative about how you intend to go from the present to a more desirable future—how you intend to direct the development of your current narrative in a desirable direction. Like other narratives, plans are not arbitrary; they must be plausible and coherent. Plausibility means that it must be easy to imagine their successful implementation, even if it will require work, and they must be natural and lifelike, in that they don't rely on implausible or exotic events or unrealistic amounts of effort and resources. Coherence means that they must be reasonably complete even though you know you probably won't end up doing things exactly as planned. A plan that lacks plausibility and/or coherence doesn't inspire confidence, which prompts its revision before it reaches the point of being used to generate an action forecast and determining its desirability.

An Example of Planning

Ed Potter is the (fictional) owner of a (fictional) appliance store that his late father founded in 1952 in the second largest town in the county. Neither he nor his father ever advertised because theirs was the only store in town. This worked fine for a long time, but recently things have begun

to change. Retirees from out of state have moved in, built new houses in developments to the south of town as well as throughout the surrounding countryside. Many of these newcomers have furnished their houses, including appliances, from catalogues and on-line suppliers rather than shopping at local stores. Ed has heard through the grapevine that quite a few of these purchasers have been unpleasantly surprised when what they got wasn't what they thought they had ordered. Disappointment, coupled with the hassle and time lost returning the unsatisfactory goods has made many of them unwilling to use catalogues and on-line sales again—resulting in trips to stores in the county seat, 70 miles away.

Local merchants have discussed ways of convincing people that they'd do better to shop locally, not just because it is more convenient but because they could see the products before buying them. But they can't seem to agree on how to get their message out; so, when all is said and done, more has been said than done.

This lack of progress has convinced Ed that he's on his own. After thinking about it, he is considering advertising for the first time in his business's history. But, he doesn't want to commit until he sees if it is worthwhile. So, he has decided to do a trial run, an experiment—he'll give advertising his best shot for one year. If after a year it has increased sales enough to justify the expense, he'll adopt it as part of his "business strategy" (he's never thought of himself as having a business strategy, he just runs his store as best he can). If his experiment doesn't produce the desired results, he'll drop it and start thinking about selling the store.

Late one night, Ed sat down with a cold beer, tipped back his chair, tipped back his beer, and began to work out a plan. He knew that his plan had to be flexible, so he tried to build flexibility into it. He also knew that the plan had to lead to his desired future—a big enough increase in

sales to justify the expense of advertising. This is his internal monologue (his narrative) as he created his plan:

“My target audience is anyone for about 30 miles around who wants to buy a refrigerator, washing machine, dryer, stove or any other appliance, but my specific target is people who have recently moved here and bought or built a home that needs appliances.

“I’ll begin by advertising in the daily newspaper because they have people who’ll help me put together attractive ads and because it isn’t too expensive. If this works, I’ll also advertise in the county-wide weekly paper. If it doesn’t work, I’ll drop print advertising and switch to TV ads.

“Either way, I’ll supplement my advertisements with offers of both free delivery of all purchases and a 10% discount on purchases totaling \$1,000 or more. If this bumps sales, I’ll simply keep doing it as long as it keeps working. If it doesn’t bump sales, I’ll drop both newspapers and TV and do direct mail advertising by sending out an illustrated monthly flyer.

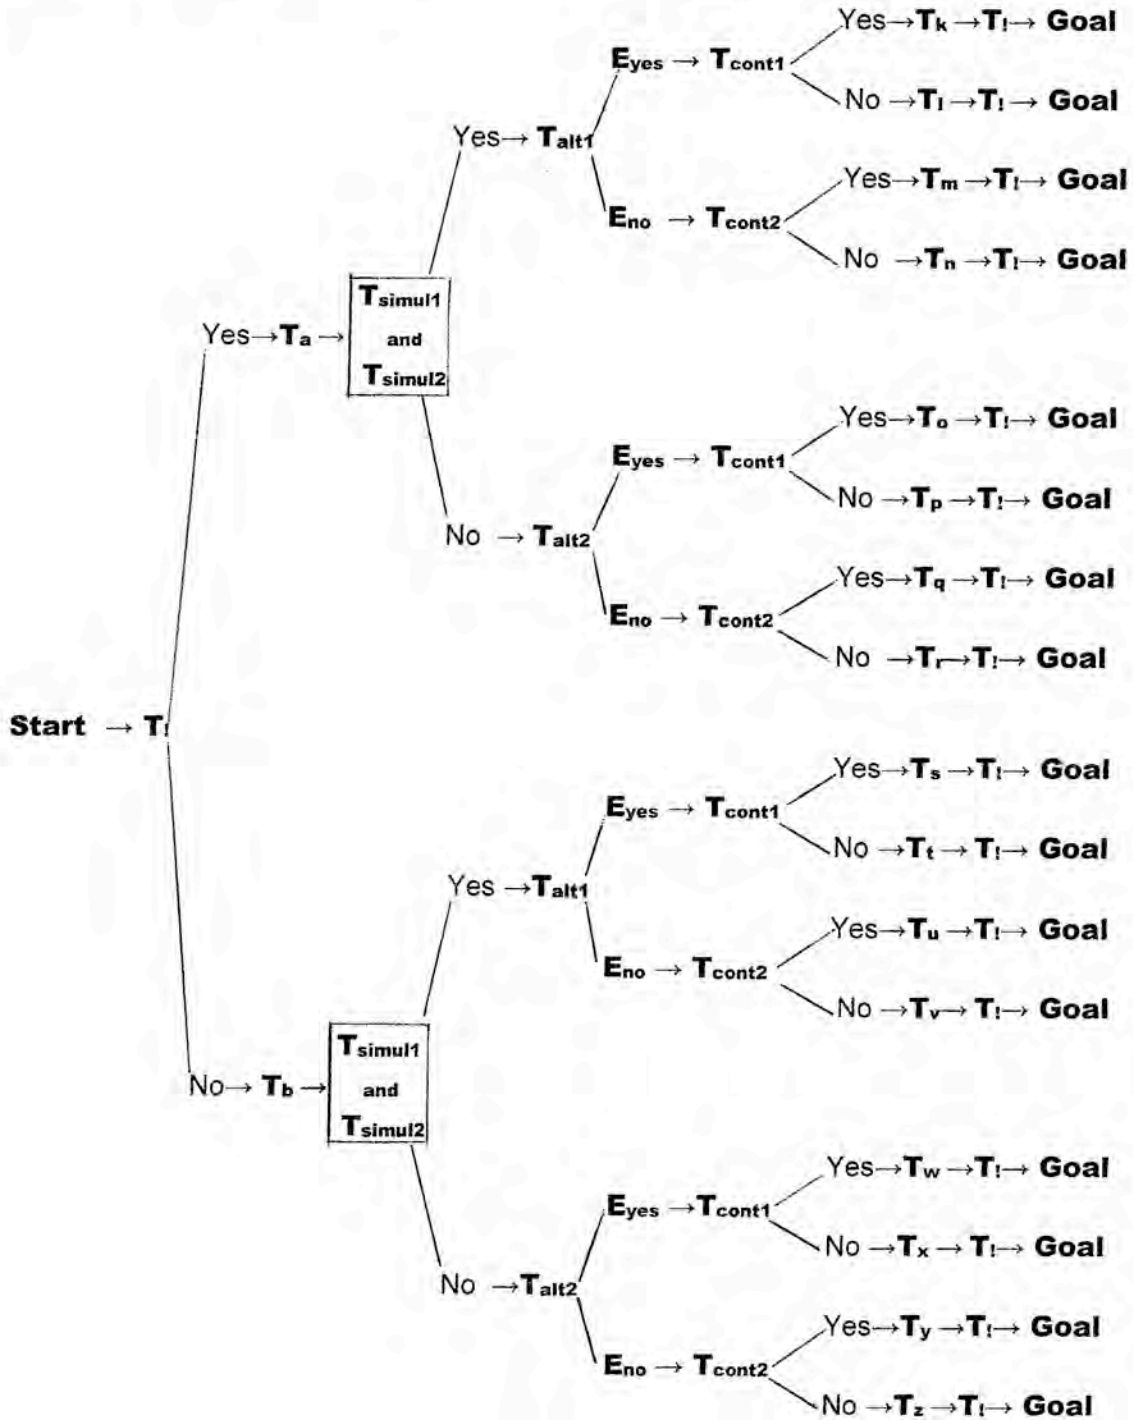
“The fly in the ointment is that Big Box, a nationwide appliance retailer is rumored to be considering a new store in the county seat, 70 miles from here. If this happens, I’ve got to do something—which probably will consist of emphasizing the convenience of shopping locally and offering a kitchen design consultant (Lisa, my niece, who has wonderful color sense and has successfully designed very nice kitchens for herself and several friends) who will work closely with the customer and the customer’s contractor. If Big Box decides not to open a store in the county seat, I can continue what I’m already doing, saving the consulting service for when I need it somewhere down the road. Finally, toward the end of my one-year advertising experiment, I’ll have a big sale and sponsor a circus or throw a town picnic or do something else splashy—after which I will look back over the year to see if I have found a mix of methods of advertising that works and that justifies the expense.”

Translation. As luck would have it, the plan the diagram in Figure 4.1 is the same as Ed's plan. Translating the diagram into his plan: "I'll **Start** by advertising in the local newspaper, T_l . If this works, I'll expand to the weekly paper, T_a ; if not, I'll switch to TV, T_b . Either way, I'll do two things simultaneously: give free delivery, $T_{simult1}$, and offer a 10% discount on purchases totaling \$1000 or more, $T_{simult2}$. If this works, I'll continue doing it, T_{alt1} ; if it doesn't, I'll try direct mail of flyers, T_{alt2} .

"Then there is the problem of Big Box opening a competing store in the county seat. If it does, E_{yes} , I'll start emphasizing convenience of shopping locally and I'll add a design consultant, T_{cont1} . If it doesn't, E_{no} , I'll continue doing what I'm doing, T_{cont2} , and save the local shopping and consultant angles until sometime when I need them. At the end of the year, I'll do my big splashy thing, whatever it is, $T_?$. After that, I'll sit down and go over sales and expenses and see if the former justifies the latter, T_l . Either way, I will have reached my **Goal** of completing my advertising experiment. Of course, the interesting step is T_l , the comparison of sales and expenses. This is the point at which I'll decide whether advertising is moving me toward my desired future of increased sales or if I should start thinking about selling the store."

A Picture of Ed's Plan. Figure 4.2., below, contains a picture of Ed's plan in the form of a tree diagram. Each arrow represents something Ed will do. Each branching point represents the possible results of what Ed will do or what some external force will do. Therefore, the plan's flexibility is at the branching points—what Ed will do next depends on the results of what has been accomplished up to that point. Each path through the tree tells a

(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)
Start	→ T _i	T _?	T _{simul1} and T _{simul2}	T _{alt1}	E _{yes}	T _{cont1}	T _?	T _i	Goal
				T _{alt2}	E _{no}	T _{cont2}			



different story. The top path says that Ed's newspaper advertisements appeared to increase sales so he also advertised in the weekly newspaper. Advertisements in both places offered free delivery and a 10% discount on big purchases, which increased sales. Then Big Box opened a store in the county seat, so Ed started emphasizing how convenient it was to shop locally instead of taking a 140 mile round trip to the competitor's new store, and he added a design consultant. This worked, so he finished the year by sponsoring a small traveling circus in conjunction with a big sale. Then he took stock to see if advertising was worth it, which finished off his year-long experiment.

Other paths would tell other stories, but I'll leave it up to you to trace through and figure out what they are. It is pretty safe to say that the paths along the lowest branches of the tree yield pretty sad stories; if things unfold this way, Ed probably will give up advertising and sell his store.

Box 4c: Incremental and Master Plans

A plan is a sequence of tactics designed to move you toward a more desirable future, where each tactic is a rule based on past experience or instruction. Because you have varying degrees of confidence (strength of expectation) in each of these tactics, you also have more or less confidence in the plan of which they are constituents. In general, your confidence in the overall plan will be greater the fewer its number of constituent tactics, if only because there are fewer opportunities for things to go wrong than there are for complex, multi-step plans that take a long time to implement.

As a result, most day-to-day plans are short and simple, but there are a lot of them. That is, as soon as a short, simple plan has been implemented, it is followed by a new short, simple plan, then another, and another, and so on. In this way, you "feel your way along," toward your desired future. Each new mini-plan, which usually consists of only one or two tactics, takes into account feedback about what happened as the result of implementing the previous mini-plans. This allows you adjust what you're doing to allow for unforeseen obstacles and curves in the road, and to compensate for the failure of one or more of the mini-plans, while keeping things headed in the right direction. Thus, by using a sequence of mini-plans that each move things forward incrementally, you get where you wanted to go. This is far better than having to follow a complex master plan that was wholly worked out before you began

and therefore was constructed in ignorance of the unforeseen obstacles and curves that lay ahead.

For example, suppose that you have a very difficult boss and your forecasted future offers only more abuse. You want relief, but you don't know exactly what that means other than his being less difficult. Your incremental plan starts with a general tactic, being pleasant and enthusiastic, during the execution of which you keep an eye out for improvement in his behavior. If this tactic appears to be producing good results, you can move to more specific sequence of tactics—perhaps offering to take on some task that needs doing and otherwise showing that you are a willing worker. Eventually, your working relationship with your boss may mellow enough that your distress goes away or is reduced to an acceptable level.

On the other hand, if your first tactic, pleasantness and enthusiasm, doesn't do much, you have to think of some other way to reduce your distress. Let's say that the first thing you think of is moving to a new job in the same company, and then it occurs to you that, while you're at it, you might as well search more broadly. At first your search plan is rather casual— newspaper want-ads, postings within your company, and the like. Before long you discover that your skills are in demand in a variety of industries, and you begin to imagine various alternative futures. Among these, making a big change, perhaps even moving to a different city, begins to emerge as an attractive idea. Whereupon you begin to consider where you might like to live and what kind of industry might provide the best opportunities, and you launch a more structured search. Eventually, you end up working in Amsterdam for a multinational company—all because your boss was difficult.

This illustrates how plans change as a function of both feedback and opportunity, where opportunity often is nothing more than new information about possibilities and a new way of thinking about yourself and what you want. Note too that comprehensive and incremental plans are not mutually exclusive; most plans are necessarily a mixture. There are parts where it is very clear what you must do and parts where you'll just have to wing it and see what happens.

Implementation

After a plan is formulated, it is joined to the current narrative to create an action forecast. The future offered by this forecast is evaluated in the same way an extrapolated forecast is evaluated. The features of the forecasted future are compared to your normative expectations—your primary and secondary values—which define your desired future. If they meet your expectations, fine. If they don't—if there is a discrepancy between them and your expectations—it counts against the desirability of the forecasted future and, as a result, against the adequacy of the plan. If the overall discrepancy is large, you reject the

plan and either revise it or start over. If the overall discrepancy is zero or within acceptable limits (which we'll discuss in the next chapter), you accept the plan and set about implementing it.

Implementation has two components, executing the planned tactics (action) and monitoring the results. Execution of tactics is rather straightforward; you do what your contingent rules lead you to expect will create a desirable future. Execution of each tactic produces results, some of which are desirable and some of which aren't (i.e., some of which move you closer to a desirable future and some of which don't). Desirable results motivate you to press on because you are making progress. Undesirable results motivate you to revise or abandon the plan.

Monitoring

When you are implementing a plan, the implementation becomes a central part of your current narrative because it is what is going on at the moment. As a result, when you make a new extrapolated forecast, it takes your plan into account (which makes it, in fact, an action forecast—but let's not complicate things). And, new extrapolated forecasts can be generated repeatedly throughout implementation. If there are fewer discrepancies between the latest forecast and your normative expectations—your values—than for previous forecasts, you can assume that progress is being made toward achievement of your desired future. If the discrepancies stay the same or increase, you must reconsider your plan.¹⁵

¹⁵ Sometimes you deliberately do something that moves you away from the desired future in order to set things up for future progress—sort of like falling back to make a pass. Because you know you're doing it, you can take it into account and regard it as progress.

Progress Report

In this chapter we have examined the second of the four propositions listed at the end of Chapter 1:

1. Cognitive narratives provide meaning and continuity by linking the past with the present and allowing us to make educated guesses about the future (forecasts). This was addressed in Chapter 2.
2. Based on our past experience and our values, we create rules, some of which are used to make forecasts and some of which are used to evaluate the desirability of the future offered by those forecasts. This was addressed in Chapter 3.
3. *When the forecasted future is undesirable, we construct plans that guide actions designed to promote a more desirable future.* This has been addressed in this chapter, Chapter 4
4. All decisions are about discrepancies and what to do about them. This will be addressed in Chapter 5.

In the next chapter we will examine the fourth proposition.

Sources and Further Reading

Alexander, L. D. (1985). Successfully implementing strategic decisions.

Long Range Planning, 18, 91-97.

Beach, L. R. (1990). *Image theory: Decision making in personal and organizational contexts*. Chichester, UK: Wiley.

Beach, L. R. (2006). *Leadership and the art of change*. Thousand Oaks, CA: Sage.

Jungermann, H., & Thüring, J. (1987). The use of causal knowledge for inferential reasoning. In J. L. Mumpower, L. D. Phillips, O. Renn, & V. R. R. Uppuluria (Eds.), *Expert judgment and expert systems* (pp.131-146). New York: Springer.

Mintzberg, H. (1978). Patterns in strategy formation. *Management Science*, 24, 934-948.

Mintzberg, H., & Quinn, J. B. (1991). *The strategy process*. Englewood Cliffs, NJ: Prentice Hall..

Thüring, M., & Jungermann, H. (1986). Constructing and running mental models for inferences about the future. In B. Brehmer, H. Jungermann, P. Lourens & G. S. Sevon (Eds.), *New directions in research in decision making* (163-174). Amsterdam: North Holland.

Chapter 5: Decisions

As we have seen, narratives are the stories we tell ourselves about the past, present, and future. But, because the future hasn't happened yet, it is still full of potential. Decisions are the way you shape that potential so when the future gets here, it is desirable because it conforms to your values. The decisions begin when you decide about the desirability of the future offered by an extrapolated forecast. If you decide that it is undesirable, you formulate a plan and generate an action forecast. If you decide the action forecast offers a desirable future, you implement the plan that generated it. Then, as implementation progresses, you make repeated decisions about whether it is bringing you closer to a desirable future. To understand how all of this happens, we must understand how we make decisions.

Rudimentary Decisions

Let's begin at the beginning. The most rudimentary decision any creature makes is categorization and the most rudimentary categorization is recognition or identification of other creatures, objects and events—which, for ease of exposition, we'll call *entities*. Recognition means that the entity can be categorized as identical to a single entity that you have encountered before—"Here's X again." Identification means that an entity can be categorized as similar to other entities that you have encountered before—"Here's another X." In both cases, once it has been assigned to a category, you can generalize what you know about the category or about other entities in the category to the entity at hand. This gives you more information about it than what is immediately observable, which enriches your understanding of it.

A category is defined by a specific set of criteria. In recognition, one of the criteria is that the category can have but a single, “resident,” member; although it may have been encountered numerous times before. Assigning the observed entity to the category implies that it *is* the resident member, encountered once again. In identification, one of the criteria is that the category has multiple resident members; each seen once or numerous times before. Assigning the observed entity to the category implies that it is but another example of that collection of entities—this particular entity may or may not have been seen before. Of course, any entity is both unique and similar to other entities, so it is both the sole member of its own category and one of the multiple members of one or more other categories. For example, my friend Terry is himself as well as being a man, a professor, a fisherman, and so on.

The category’s criteria strictly define it, but the category’s name is rather arbitrary, although not completely arbitrary. That is, a single-member category may have many names—my spouse, my wife, Barbara, my love, my best friend, but not spouses, wives, Barbaras, lovers, friends. Similarly, a multiple-member category can have many names—professors, teachers, mentors, advisors, but not my professor, my teacher, my mentor, my advisor, and most especially not an individual’s name such as Dr. Mitchell or Terry.

Categorization Decisions

Determining whether a candidate entity should be assigned to a category requires (1) a *process* that compares the properties of the candidate to the category’s criteria and (2) a *decision* about whether this candidate meets the criteria. The decision itself is merely “Yes” or “No,” but the comparison process leading up to that decision is somewhat more elaborate. Together, the process and the decision deriving from that process are called *decision making*.

Because recognition, rather than identification, is the kind of categorization we'll be interested in from here on, let's examine a simple recognition task. Suppose I gave you a stack of photographs and asked you if the person who mugged you yesterday is in one of the photos.¹⁶ Presumably, you would compare the features of the image in each photograph with the features of your episodic memory of the mugger, the features of which would constitute the criteria for admission of a photo to the single member category, "My Mugger." Then you would decide, "Yes" or "No."

Good enough, but let's further assume that in the interest of speed and efficiency, that before you even look at a photo your default assumption is that it is a picture of the mugger unless there is sufficient evidence to the contrary. In the very simplest case, sufficiency of the evidence would depend on how many non-trivial discrepancies you detect between the features in the photo and the features in your memory of the mugger. Thus, for each photo:

- If the number of discrepancies is unreasonably large, you probably would not assign the person in the photo to your My Mugger category, which is to say that you wouldn't recognize the person in the photo as the person who mugged you. After all, a photo and your memory should be more alike than not if they are representations of the same person.
- If the number of discrepancies were reasonably small, you probably would assign the person in the photo to your My Mugger category, which is to say that you would recognize the person in the photo to be the person who mugged you. You probably would tolerate a few discrepancies because, after all,

¹⁶ On TV cop shows this task is always referred to as "identifying the 'perp'", but it really is a recognition task. In common parlance, recognition and identification are used interchangeably, but inaccurately.

a photo and a memory are unlikely to be exactly the same, even if they are representations of the same person, but there shouldn't be many discrepancies.

This example illustrates the underlying logic of what is called the *discrepancy test*, a conceptual model of the categorization process for recognition decisions. The input to the test is nontrivial *discrepancies* between the features (the admission criteria) of the category's resident member and the corresponding features of the candidate for admission to the category. The test's operative assumption, called the *identity assumption*, is that the two sets of features, and thus the category resident and the candidate, are the same unless proven otherwise (see Box 5a). The test's *decision rule* is to reject the identity assumption if the number of discrepancies exceeds a criterion, called the *rejection threshold*. Rejection of the identity assumption means that the decision is "No, the candidate entity does not meet the criteria for admission to the category." Failure to reject the assumption means that, "Yes, it does meet the criteria." In other words, "No" means that you don't recognize the candidate entity and "Yes" means that you do.

Box 5a: The Primacy of Discrepancies

To understand why the discrepancy test relies on discrepancies to the exclusion of nondiscrepancies, consider this little story (for which I thank Prof. Guido Peeters, University of Leuven, Belgium):

Assume that there is a creature that lives solely on fungi. Mushrooms are abundant and edible. Toadstools may or may not be abundant, but they are poisonous. The creature holds as its working hypothesis that every fungus it finds is a mushroom. However, that hypothesis is quickly rejected if a fungus has one or more features of a toadstool, even if its remaining features are those of a mushroom. The reverse rationale doesn't apply. That is, evidence that disconfirms the working hypothesis that a fungus is a mushroom determines the decision about its edibility, confirming evidence is irrelevant. This strategy is safest because it favors false negative decisions [rejection of edible mushrooms that have features of poisonous toadstools] over false positive decisions [acceptance of poisonous toadstools that have features of edible mushrooms].

Interpreting this story in terms of the discrepancy test: the creature works on the (identity) assumption that every fungus it encounters (a candidate entity) is an edible mushroom (is a member of the Mushroom category) unless there is evidence to the contrary; in which case it is presumed to be a poisonous toadstool. However, it requires very little evidence to reject this assumption—even if almost all of the fungus’s features match the defining features of a mushroom, those few that don’t suggest that the fungus is unsafe to eat—that it might in fact be a toadstool. In short, it is the mismatches (discrepancies), not the matches (nondiscrepancies), between the fungus’s features and the defining features of mushrooms that determine whether it is accepted as a mushroom.

Not only is this disconfirmatory strategy safe, it is efficient. Strictly speaking, even one item of disconfirming evidence is sufficient to disprove the identity assumption, but no amount of confirming evidence is sufficient to prove it because there is always the possibility that there is disconfirming evidence yet to be discovered. That is, you can prove the identity assumption is wrong by a *sufficiency* of disconfirming evidence, but proving it is right requires an *exhaustive* search of all possible evidence and every bit of it must confirm the assumption.

Stopping Search

A problem arises when the identity assumption actually is true because there won’t be any disconfirming evidence. But you can’t know that unless you observe every possible bit of evidence. To avoid this, you need some sort of “stopping rule” to keep the search for discrepancies from going on and on. The fact that our actual searches usually aren’t very long suggests that we have such a rule, even if it isn’t hard and fast. In some cases search stops because there is only a small amount of pertinent information, so an exhaustive search is feasible. (In discrepancy analysis we assume that the number of defining features of the resident category member is reasonably small, so a candidate entity that is in fact identical to it will be recognized quickly and admitted to the category.) In other cases, when there is a great deal of pertinent information, as the number of observations mount without any disconfirming evidence being found, we tend to become increasingly confident that there isn’t any disconfirming evidence to find. At some point it seems unlikely that the identity assumption isn’t going to be rejected, so we stop searching for discrepancies, assume the entities are identical, and act accordingly. One might think that this stopping point might be determined by something like the ratio between discrepancies and nondiscrepancies, or the difference between them, but empirical evidence indicates otherwise. It appears to simply boil down to the absolute number of nondiscrepancies, although the factors that determine the critical number isn’t yet well understood. It is likely that the stopping point is different for different individuals and that it differs for an individual in different circumstances.

Incidentally, research on decision makers’ use of confirmatory and disconfirmatory information produces rather paradoxical results. It is found that their searches focus on confirmatory information but that their decisions are based on the disconfirmatory information that their search uncovers. Moreover, decisions usually are made after observing only a few disconfirming observations, three or four at the most.

The Null Hypothesis

If you are familiar with statistics, disconfirmation will be familiar to you as the logic underlying null hypothesis testing. In the classic example, it is assumed that in spite of having been subjected to the experimental treatment,

the experimental group remains identical to the control group (the word “null” refers to an absence of a difference between them). Instead of assuming the null hypothesis is false until it is proven true, which would require an exhaustive search of all pertinent information, it is more efficient to assume it is true until it is proven false; to assume that the experimental treatment had no effect until empirical data prove otherwise. In principle, the search for disconfirming data should end when you encountered a single counterexample, which would prove the null hypothesis to be false. In practice, because data contain various forms of error, you tolerate some degree of disconfirmation before you declare the null hypothesis to be false. The question is, of course, at what point do you stop being tolerant? The convention is that tolerance ends when the degree of disconfirmation produces a statistic (t, Z, F, etc.) that could be obtained by chance only 5% of the time or less. (But this is merely a convention; its critics argue that it is too mechanical and that experimenters should use their good sense and professional judgment about the weight of disconfirming evidence rather than slavishly following the 5% rule.)

A Model of the Basic Discrepancy Test

As is customary in science, decision researchers formalize their ideas about decision making in the form of theories and models. The discrepancy test is a theory about recognition decisions that we now are going to cast in the form of a model. The purpose of a model is to make the assumptions and logic of the theory clear.

We'll start with a very simple model of the discrepancy test, which we'll call the basic model. It is simple because the speed with which most recognition decisions are made suggests that the underlying process is likely to be quite simple. Moreover, the scientific principle known as the Law of Parsimony (also called Occam's Razor), “the simplest explanation is best,” dictates that we start with the most parsimonious model we can conceive of, complicating it only when doing so is warranted.

Elements

- *Entities*: Persons, other creatures, objects, or events.
- *Features*: Entities are defined by their characteristics, called features.
- *Defining Features*: The features of the category's resident entity that uniquely define it.

- *Candidate*: The entity being considered for admission to a single-member category.
- *Comparable Features*: Those features of the resident entity for which the candidate entity has corresponding features.
- *Identity Assumption*: Unless or until there is sufficient evidence to the contrary, it is assumed that there are only trivial discrepancies between the resident entity and the candidate entity on the comparable features; that is, that the entities are identical.
- *Evidence*: Nontrivial discrepancies between the defining features of the resident entity and the comparable features of the candidate entity constitute evidence against the identity assumption.
- *Rejection Threshold*. The number of nontrivial discrepancies between the defining features of the resident entity and the comparable relevant features of the candidate entity that constitutes sufficient evidence to reject the identity assumption.

Process

- Compare the candidate entity's comparable features with the resident entity's corresponding defining features.
- Note nontrivial discrepancies between the entities on each comparable feature.
- Determine the weight of evidence against the identity assumption, which is the number of nontrivial discrepancies, **D**.
- The number of nontrivial discrepancies, **D**, that is sufficient to reject the identity assumption is the *rejection threshold*, **T**.

Decision Rules:

- Reject the identity assumption if $\mathbf{D} \geq \mathbf{T}$. That is, reject if the weight of evidence, **D**, equals or exceeds the rejection threshold, **T**.
- Do not reject the identity assumption if, after an exhaustive search of the resident entities defining features, $\mathbf{D} < \mathbf{T}$. That is, do not reject if,

after the all of the comparable features have been compared for discrepancy, the weight of evidence, **D**, is less than the rejection threshold, **T**.

Decision:

- If the identity assumption is rejected, the decision is “No, the candidate entity is not recognized as the resident member of the category.”
- If the identity assumption is not rejected, the decision is “Yes, the candidate entity is recognized as the resident member of the category.”

Note that nondiscrepancies contribute nothing to the model; only discrepant features contribute to the overall discrepancy between the two entities, **D**. As explained in Box 5a, ignoring the nondiscrepant features makes sense because in decisions about feature similarity, which underlie recognition, nondiscrepancies cannot cancel out or balance out discrepancies. Strictly speaking, even one discrepancy is sufficient for proving the identity assumption is wrong, but no amount of non-discrepancies can prove that it is right—the next observation may turn out to be a discrepancy. But, having said that, both prudence and common sense recommend allowing for observational error by setting the rejection threshold at greater than one discrepancy.

Note also that the search for discrepant features is limited to those features of the candidate entity that are comparable to the defining features of the resident entity, not to all of their comparable features. Thus, while it may have many features, the resident entity's defining features—the features that make it unique—comprise the set that will be compared for discrepancies. Of course, information may not be available about the candidate entity's status on all of those features, so actual comparison may occur for an even smaller set. These limits on the set of

features that will be compared have two implications: First, this set is almost always smaller than the full range of either entity's features. Second, the search for discrepancies stops in either of two ways: (1) when the rejection threshold is reached or (2) when the set of comparable features has been exhausted (see Box 5a).

Finally, note that the rejection threshold, **T**, is not rigidly defined. This and other thresholds in the model are discussed in Box 5b.

Box 5b: Thresholds

In the study of perception, there are two kinds of thresholds, absolute and difference. An absolute threshold is the minimum perceivable amount of sensory stimulation; the minimum level of light, pressure on the skin, odor, sound, and so on. A differential threshold is the minimum perceivable difference between two things; differences between the weights of two objects, between the lengths of two lines in a visual display, between the pitches or the loudness' of two sounds, and so on. In neither case is the perceivable minimum actually the physical minimum—very low levels of stimulation or very small differences go undetected.

In spite of the name, absolute thresholds aren't absolute. That is, if detection is important enough, observers become increasingly observant and perceive lower than usual levels of stimulation—although they tend to make false positive errors (they say they perceive the stimulus when it isn't there). Similarly, difference thresholds decrease if detection is important enough, again with false positive errors.

Just as thresholds decrease when it is important to detect stimuli or differences, thresholds may increase when detection isn't important. In this case the errors are false negatives (failure to perceive the stimulus or difference when it is there). For example, if you work in a room with a ticking clock, after a time you cease to hear it. This doesn't mean you become deaf, it merely means that you adapt to the sound. If the ticking suddenly becomes louder or is brought to your attention, you hear it again. This variability in perceptual thresholds also is found in judgmental thresholds: judgments of absolute length, breadth, distance, time as well as judgments of differences among them.

There are two thresholds in the model discussed in the text; the discrepancy threshold and the rejection threshold. The discrepancy threshold is a judgmental difference threshold; the degree of difference between the features of two entities below which a difference is judged to be trivial and ignored and above which a difference is judged to be nontrivial and counted as evidence against the identity assumption. That is, below the discrepancy threshold, observers may very well know that there is a difference between the entities but they don't regard it as big enough to count as a discrepancy.

Like all thresholds, discrepancy thresholds are different for different observers and are different for the same observer at different times and under different circumstances. Research shows, for example, that observers'

discrepancy thresholds decrease as a function of how many prior discrepancies they have observed. That is, a difference that is judged to be too trivial to count as a discrepancy when it is the first difference you discover may be judged to be nontrivial, and counted as a discrepancy if it is discovered after you've discovered two or three others. (This is why judges seldom allow prosecutors to inform juries about a defendant's prior convictions; knowledge about previous discrepancies between the defendant's behavior and the law could induce jurors to convict on the basis of weaker evidence than they might otherwise have required.)

The rejection threshold is a judgmental absolute threshold; it is the number of discrepancies, **T**, above which the identity assumption is judged to be rejected and below which the identity assumption is judged to be retained. This threshold appears depend on how important it is to discriminate between a forecast and the desired future.

Adapting the Basic Model

Of course, our interest isn't in categorization decisions, as such. Our interest is in decisions about whether the features of a forecasted future are similar to those of your desired future. While it may not look like it, categorization and desirability decisions are logically parallel, so our model for categorization decisions can help us understand decisions about the desirability of forecasted futures.

To make the parallel clear, think of the desired future as the resident entity of a category named The Desired Future and the future offered by an extrapolated forecast or an action forecast as a candidate for admission to that category. This is the same thing as asking, "Is this (extrapolated or action) forecasted future identical to my desired future?" Generalizing the discrepancy test described above to answer this question requires you to compare the features of the forecasted future with the defining features of your desired future and count the nontrivial discrepancies. If the sum, **D**, is acceptably small, your decision must be, "Yes, the forecasted future is the same as my desired future," in which case you would sit back and allow the forecasted future to become the actual future because it is desirable. If the sum, **D**, is unacceptably large, your decision must be, "No, the forecasted future is not the same as my

desired future," in which case you would search for a way to make the actual future, when it arrives, acceptably similar to your desired future.

To complete the parallel, we must restate the basic discrepancy model for decisions about categorization, given above, as a model for decisions about forecasted futures' desirability:

Elements:

- Entities: The desired and forecasted futures.
 - *Desired Future*: The future as it would be if it conformed to the decision maker's normative expectations; it is the resident entity of its category.
 - *Forecasted Future*: The future as it is predicted to be; it is a candidate entity.
 - ✓ The extrapolated forecast is a projection of the current narrative assuming that no action is taken to change the natural flow of events (noncontingent rules).
 - ✓ The action forecast is a projection of the current narrative assuming that a plan of action is implemented to change the flow of events (contingent rules).
- *Features*: Characteristics of the entities.
- *Defining Features*: The features of the desired future for which normative rules (values) create strong normative expectations.
- *Comparable Features*: The features of the forecasted future that are comparable to the defining features of the desired future.
- *Identity Assumption*: A forecasted future is assumed to be identical to the desired future unless there is sufficient evidence to reject the assumption.

- *Evidence*: Nontrivial discrepancies between the defining features of the desired future and the comparable features of the forecasted future constitute evidence against the identity assumption.
- *Rejection threshold*. When the weight of evidence—the number of discrepancies—becomes too great to sustain the identity assumption, it is rejected and the forecasted future is assumed to be undesirable.

Process:

- Compare the forecasted future's comparable features with the desired future's defining features.
- Note the nontrivial differences between features of the desired future and the forecasted future; each difference constitutes a *discrepancy*, labeled \mathbf{d}_c , where \mathbf{c} indicates one of the desired future's \mathbf{n} desirable features, $\mathbf{c} = 1$ to \mathbf{n} .
- Determine the weight of evidence against the identity assumption, which is the number of discrepancies, $\mathbf{D} = \sum_{c=1}^{\mathbf{n}} \mathbf{d}_c$.
- The number of discrepancies, \mathbf{D} , that is sufficient to reject the identity assumption is the *rejection threshold*, \mathbf{T} .

Decision Rule:

- Reject the identity assumption if $\mathbf{D} \geq \mathbf{T}$. That is, reject if the weight of evidence, \mathbf{D} , equals or exceeds the rejection threshold, \mathbf{T} .
- Do not reject the identity assumption if, after an exhaustive search $\mathbf{D} < \mathbf{T}$. That is, do not reject if the weight of evidence, \mathbf{D} , is less than the rejection threshold, \mathbf{T} .

Decision:

- If the identity assumption is rejected, the decision must be “No, the forecasted future is not desirable”.
- If the identity assumption is not rejected, the decision must be “Yes, the forecasted future is desirable”.

Implications:

- If the decision is “No,” the forecasted future must be judged to be undesirable. If the forecast was an extrapolated forecast, a plan must be created to make the actual future, when it arrives, desirable. If it was an action forecast, the plan must be repaired or replaced and the result re-submitted to the discrepancy test until the future it offers is desirable.
- If the decision is “Yes,” the forecasted future can be accepted as desirable. If it was the result of an extrapolated forecast, you can let it happen. If it was the result of an action forecast, the plan upon which it was based can be implemented.

Enriching the Discrepancy Test

The foregoing model treats all discrepancies as of equal importance and magnitude, which makes it simple enough to account for rapid, rudimentary recognition decisions, but which clearly makes it too simple to account for deliberative decisions. Deliberation allows more to be taken into consideration than merely the number of discrepancies, namely the *differential importance* of discrepant features as well as the *differential magnitudes* of the discrepancies themselves. This requires an enriched version of our basic discrepancy test:

Elements:

The elements of this enriched model are the same as for the basic model for forecasts except that the following is substituted for its counterpart in the basic model:

- *Evidence*: Nontrivial discrepancies between the defining features of the desirable future and the comparable features of the forecasted future constitute evidence against the identity assumption. The contribution of each discrepancy to the overall weight of evidence incorporates:
 - *Magnitude*: The size of a nontrivial discrepancy.
 - *Importance*: The strength of the decision maker's normative expectation for each comparable feature.

Process:

- Compare the forecasted future's comparable features with the desired future's defining features.
- Assign a score of 1 (one) to a discrepancy, d_c , and a score of 0 (zero) to a nondiscrepancy.
- The contribution of an individual discrepancy to the weight of evidence against the identity assumption is the discrepancy's score, d_c , multiplied (weighted) by its magnitude, m_c and by its importance, i_c , which is labeled $d_c i_c m_c$.
- The weight of evidence, D' , against the identity assumption is a function of the sum of the $d_c i_c m_c$ across the desired future's n defining features, $D' = \sum_{c=1 \text{ to } n} d_c i_c m_c / 100$, where 100 is a scaling unit that makes D' lie between zero and n (See Appendix 2 for an explanation of scaling units).

Decision Rule:

- Reject the identity assumption if $D' \geq T$. That is, reject if the overall weight of evidence, D' , equals or exceeds the rejection threshold, T .
- Do not reject the identity assumption if, after exhaustive search, $D' < T$. That is, do not reject if, after all of the desired future's

defining features have been compared for discrepancy with the corresponding features of the forecast, the weight of evidence, D' , is less than the rejection threshold, T .

Decision:

- If the identity assumption is rejected, the decision is “No, the forecasted future is not the same as the desired future.”
- If the identity assumption is not rejected, the decision is “Yes, the forecasted future is the same as the desired future.”

Implications:

- If the decision is “No,” the forecast must be rejected and a remedy sought.
- If the decision is “Yes,” the forecast can be accepted.

In the following discussion, we will see how the basic and elaborated versions of the discrepancy model relate to decisions of differing complexity.

Levels of Decision Making

We have developed a model of decision making based on categorization and then we have generalized it to the decisions about the desirability of forecasted future. But, of course, some futures, both forecasted and desired, are more complex than others. As a result, some decisions require more time and effort than others. Decisions about simple futures can be made very quickly and easily—virtually automatically. Decisions about complex futures require more work. We will refer to this difference in complexity as a difference in *decision level*.

Level 1 Decisions

Level 1 decisions are about desired futures that have a small number of defining features and, usually, short time horizons. The small number of defining features means that the discrepancy test will be simple; the short time horizon means that the test must be quick. Indeed,

most level 1 decisions are so simple and quick that they typically are described as “intuitive” or “unconscious” because you aren't even aware of having made them.

Level 1 decisions are about routine things; necessary but not requiring much thought. Afterward you may not even remember them—putting the cat out at night, turning the stove after cooking dinner, brushing your teeth before bed. In fact, you don't regard some of these decisions as decisions at all; a car swerves into your lane and you put on your brakes—a decision? Your cell phone rings and you quickly reach to quiet it—a decision? You see your son's shoes in the middle of the living room floor and you pick them up and put them in his room—did you decide to do it or merely do it?

Prior to each of the decisions listed in the last paragraph, there was a forecast about what the future would be if you didn't do anything: the cat would be in the house, the stove would remain on, your teeth would remain foul, the other car would hit yours, your cell phone would continue to bleat, your son's shoes would remain on the living room floor. And there was a desired future: the cat should be out, the stove off, your teeth brushed, the accident averted, your cell phone quieted, the shoes put away. In each case the discrepancy between the forecasted and the desired futures prompted action. But, because so little was involved in each decision, the discrepancy and the corrective action were virtually automatic; so automatic that they made little impression on your consciousness, thus making it difficult to remember them.

The basic discrepancy test is an adequate model for level 1 decisions because it boils down to simply counting discrepancies for a small number of defining features that you can easily keep in mind—usually fewer than 9, which is pretty much the maximum number of items that most of us can hold in immediate memory. As a result,

everything happens so fast that when the decision is “No, the future isn’t desirable,” you react with a habitual response; you don’t need a plan because you already know what to do.¹⁷

Level 2 Decisions

Level 2 decisions are about desired futures that have a moderate number of defining features and, frequently, moderate time horizons, which means the decision is moderately complex but you have time to think about it. These are the day-to-day decisions that contribute the run-of-the-mill material to your narratives without changing the storyline very much. They are, however, recognizable as decisions and you would refer to them as such in a conversation. They require conscious deliberation; perhaps not a lot, but some. And, because consciousness is involved, they can take more into account than level 1 decisions can—specifically, the importance and magnitude of discrepancies. Actions that follow level 2 decisions usually are recycled plans from your episodic memory, but new plans are not infrequent.

The enriched version of the discrepancy test models level 2 decisions.

Level 3 Decisions

Level 3 decisions are about desired futures that have a large number of defining features and, usually, a long time horizons. These are the momentous decisions that mark major turning points in the development of your narrative and your life, and to which you devote a great deal of cognitive effort—if only to worry about them. Actions following a level 3 decision are largely unique because you have seldom, or never, encountered the situation before, so you have no past

¹⁷ And, as we shall see, if you don’t know what to do, the level 1 decision gets kicked up to level 2 so you can consciously think about what to do.

experience upon which to draw—although you may have helpful general knowledge such as how to dial 911 or how to hire a lawyer.

The enriched version of the discrepancy test also models level 3 because, like level 2 decisions, they involve consciousness and take into account the importance and magnitude of discrepancies.

Cognitive Overload

These three levels of decisions actually represent regions on a continuum—simple, moderate, and complex, with all shades in between—but it is easier to talk in terms of distinct levels.

Complex desired futures have more defining features than simple ones. As a result, some Level 2 decisions and all level 3 decisions involve more defining features than you can keep in mind at once (on average between about 5 features for some people and about 9 for others, but never very many for any of us). The result is cognitive overload that makes the decision seem elusive and difficult to keep track of—you focus first here and then there, never being able to get everything clearly in mind at once. As a result, the more complex the desired future is, the more difficult it is to be confident in your decision about it and, therefore, about what you should do next. These are the decisions that keep you awake at night.

More About Level 1 Decisions

In order to save time and cognitive effort, it is prudent to deal with the decisions at the lowest level possible. As a result, every forecasted future is first approached as a level 1 decision. If that doesn't work, it is treated as a level 2 or level 3 decision. The point is, level 1 gets the first cut at every decision. If this first cut test detects enough discrepancies to justify deciding the forecasted future is undesirable, you don't have to go any further. After all, you needn't look at any more features once the rejection threshold has been reached—and knowing about importance and magnitude is irrelevant.

These first cut decisions at level 1 serve as an early warning system—if the forecasted future is badly flawed, the basic discrepancy test is likely to detect it and set off alarm bells—you almost instantaneously know that the forecasted future is undesirable. Because this happens so quickly and because all you experience consciously is the emotions evoked by the discrepancies' violations of your normative expectancies, it sometimes seems a little mysterious, thus meriting the label, “intuitive.” Forecasted futures that are only slightly flawed (and few are flawless) are likely to survive the level 1 first cut decision and be passed on to level 2.

All of this is important because you'll never make a level 2 or level 3 decision about the future offered by an extrapolated or action forecast that hasn't already survived a level 1 first cut decision. Moreover, no matter at which level you decide that an extrapolated forecast's future is undesirable, level 1 will get the first cut at the action forecast for the plan you devise to remedy things. The result is that level 1 is an early warning system for severely flawed plans as well as for severely flawed extrapolated forecasts. For example, if your financial future looks particularly bleak it will set off alarms telling you its undesirable and urging you to do something. However, alarms also will go off if your plan is to fix things by forging checks, swindling old ladies, or holding up convenience stores—all, presumably, violations of your values and therefore undesirable. Nor are you likely to let yourself get away with staying drunk for a month or hiding under your bed, because they won't produce a desirable future. All of these plans would be rejected at Level 1, where silly, immoral, futile, or useless plans are routinely screened out because their action forecasts are so flawed that even level 1's basic discrepancy test can detect it and prompt their rejection. The happy result is that you do far fewer stupid things than you otherwise might.

If this isn't enough, here are two more reasons why it's important to know about level 1 decisions:

- Level 1 decisions have veto rights over level 2 and level 3 decisions.
- In the interest of reducing cognitive overload, you tend to simplify level 2 and level 3 decisions so you can move them downward toward being simple, intuitive, level 1 decisions.

Veto Rights

To illustrate veto rights, consider the post-choice oversight of plans exercised by a intuitive level 1 decision after a deliberative level 2 or 3 decision: If intuition tells you that the house you plan to buy is just not right, you'll back out of the deal: "Something just doesn't feel right about it," which means that a subset of the features of the future offered by your action forecast—buying the house and living in it—are discrepant from your desired future, making the plan undesirable, and, therefore, making the purchase of the house undesirable. You may not be able to explain why you changed your mind; you just know that you felt uneasy—your intuition told you not to buy.

Similarly, every consultant has had the experience of working closely with a client on deliberations for a level 3 decision, painstakingly reaching a conclusion about a plan to solve a complex business problem, only to watch with chagrin as the client sits back his or her chair and utters these fatal words: "I'm just not comfortable with it." Hours of rigorous thought and days of intense collaboration are torpedoed by a little intuitive feeling, a niggling hunch, a nagging discomfort. And once a plan has been vetoed, clients seldom change their minds, no matter how compelling the analysis may seem to others. In a conflict between reason and intuition, reason seldom wins.

Level 1 oversight and veto serves an important function. We often get so caught up in the effort of dealing with complex decisions, and so frustrated with our inability to keep everything in mind, that we lose track of how we really feel about things or we just blunder through just to get things over with. Level 1 oversight is the reality check that saves us from mistaking hard work for wisdom and recklessness for decisiveness.

Decision Simplification

Level 1 decisions are everybody's secret vice because they are virtually effortless and we attribute a good deal of credibility to the emotions that accompany them. Because most level 2 and 3 decisions are too complex to get our minds around, even the most thoughtful among us simplify them to make them as much as possible like level 1 decisions. That is, we tend to focus on only a few of the features of the decision to the exclusion of others, or we lump features together to reduce the number we must consider, or we ignore possible ramifications and remote consequences in an effort to pare the problem down to a manageable size—where manageable is defined as level 1 or something very close to it.

We seldom can simplify them all the way down to level 1, but we try to get them as simple as we can. We usually know, somewhere in the back of our minds, that simplification is dangerous, but we have little choice because we lack the skills, attention span, and the cognitive capacity to manage the details of complex decisions. Whatever its dangers, simplification sets the decision up to be faster and easier than it otherwise would be (perhaps faster and easier than it should be).

Decisions During Plan Creation and Implementation

We've been discussing the role of the discrepancy test in decisions about the acceptability of extrapolated and action forecasts. However, its role in decisions about action forecasts goes beyond what has been

discussed thus far. The test also plays a role in both the creation of effective plans and the monitoring of a plan's progress during implementation.

Plan Creation

In some cases, of course, a plan already exists when the occasion for its use arises—you used it before and it is stored in your episodic memory. In this case, it can be recycled, although it usually must be modified to fit the unique circumstances of the current situation. After all, as the old saying has it, “You can't step in the same stream twice;” the progress of events almost always means a plan from the past requires adjustment to fit the present and future.

To adjust a recycled plan, you link the plan to your current narrative and use them both to make an action forecast. Then you conduct a discrepancy test to see if the forecasted future is discrepant from your desired future. If the recycled plan's forecasted future isn't significantly discrepant from your desired future, you can implement it without adjusting it. If it is significantly discrepant, you must revise or replace some of the plan's component tactics, make a new action forecast, and do another discrepancy test. If your adjustments eliminated the discrepancies, the action forecast and the desired future will be sufficiently similar to permit you to accept the adjusted plan. If not, you must adjust further, revising or replacing tactics, and go through the process again, continuing in this way until you find a plan that yields an action forecast that is not significantly different from your desired future. (Of course, adjustment of a tactic is logically the same as replacing it with another tactic that is only slightly different, so adjustment and replacement are fundamentally the same thing—so we'll refer to both of them as replacement.)

Recall that a plan is a sequence of tactics, which are contingent what-to-do rules with their corresponding contingent expectations. What recommends one tactic over another in the first place, or recommends it to be replaced or to be a replacement, is your contingent expectation that it can help eliminate one or another of the discrepancies between the plan's action forecast and your desired future. And, remember, contingent expectations take into account your efforts to make the underlying contingent rule work; so it isn't just the rule's track record that counts, it's also your appraisal of your skill and determination to make it work.

Implementation Monitoring

During plan implementation, tactics are transformed into concrete action and progress is or is not made toward achievement of your desired future. Monitoring is simply an attempt to make sure that these actions and their results don't stray from the plan and that the plan remains a viable path to your desired future. It consists of periodic updates of the plan's action forecast given what has happened up to that point. That is, starting from the conditions that exist right now (your current narrative up to now), does the as-yet-unimplemented portion of the plan produce a forecasted future that is significantly different from your desired future? If the decision is "No," then continue with the implementation. If the decision is "Yes," then you must adjust the as-yet-unimplemented portion of your plan in the manner described above.

Level 1's basic discrepancy test permits adequate monitoring for most implementations, yielding intuitive feelings that things are progressing nicely or that something isn't right. However, if something isn't right, plan adjustment usually requires at least a little conscious deliberation, pushing it to level 2 or level 3, depending on the extensiveness of the adjustments.

Choices

Most of us equate “decision” with a choice about which plan of action to adopt from among two or more plans of action. This implies comparison between the competing plans, which in most cases means that choices are level 2 or 3 decisions because comparisons involve magnitudes of discrepancies and therefore require consciousness.

To choose a plan, you must perform a discrepancy analysis for each of the competing plans’ action forecasts; the one that is least discrepant from your desired future is the one to choose. Which is to say, you consciously compare the results of the discrepancy analysis for each plan and choose the plan for which the magnitude of the overall discrepancy, **D'**, is smallest. Note that the comparison is between discrepancy analysis results, not between the plans themselves. It might be instructive, in some way, to compare the plans, but such a comparison would be irrelevant to the choice because the goal is to select a plan that will best promote your desired future. This may fly in the face of general opinion and admonitions to “compare the alternatives,” but unless the comparisons among the various plans in some way informs the comparison of each plan’s action forecast with the desired future (which comparing the plans with each other doesn’t), it is wasted effort.

We’ll discuss the mechanics of choice in Chapter 6.

Box 5c: Recognition and Identification

In the text, classification for recognition provides the foundation for our discussion of decision making and the discrepancy test. However, both recognition and identification are both important aspects of cognition and decision making.

Whenever a new entity (the candidate) presents itself to your perception, you use its salient features to probe memory to see if you recognize it. This means that you search for memories that allow you to decide if this entity is identical to one that you have previously encountered. This may require two steps, first to identify what the entity is—a person, a screwdriver, a lamp—and a second to recognize the individual person, screwdriver, or lamp. However,

the speed with which recognition occurs suggests that the process may be a single step, or at least a very rapid search. So, for example, when you approach the checkout counter at the grocery store, you see the clerk, identify her as a person and a female, and then narrow the search so you can recognize that she is the clerk you dealt with yesterday. Or, it may be that recognizing her as the clerk you dealt with yesterday tells you that she is a person and a female—after all, recognition implies identification but the reverse isn't necessarily true. It's the job of perceptual psychologists to settle such questions like this, but, either way, you recognize the clerk based on her features and the context (current narrative) in which you encounter those features.

The classes for either recognition or identification are constructed from linked episodic memories; the context constrains the search so not everything in your memory has to be considered. A single-member class that serves recognition is compounded from the various episodes in your memory that are relevant to a single entity having the features that the candidate entity has. Similarly, a multiple-member class that serves identification calls upon episodic memories of all context-pertinent entities having the features that the candidate entity has. In both cases, the discrepancy test allows you to decide whether the retrieved single-member or multiple-member class fits the features of the candidate entity well enough to permit its recognition and/or identification—which is to say, well enough to permit the candidate entity to be assigned to a single-entity class and/or to a multiple-entity class. And, in both cases, classifying the candidate entity allows you to go beyond the perceptually available information by accessing knowledge about it based on either your previous experience with it or your previous experience with other examples of it.

The knowledge you access after recognition is specific to the candidate entity but the knowledge you access after identification is general to entities like it. Of course, you sometimes recognize an entity (single-member category), only to realize that you know very little about it and have to revert to what you know about entities like it (multiple-member category). For example, if you've only seen the grocery clerk once before, you probably don't know much about her and recognizing her isn't very useful. Far better to fall back on what you know about grocery clerks in general, and how to interact with them, so you can efficiently purchase your groceries and be on your way.

Unaided and Aided Decision Making

Decision researchers distinguish between unaided and aided decisions. Unaided decisions are made entirely in your head. The limitations of having only your head to depend upon are well documented; if nothing else, we all lack the attention span and cognitive capacity we need to keep numerous considerations in mind at once. Aided decisions, on the other hand, are made using pencils, paper, calculators, computers, and the like, in conjunction with step by step

instructions, called *decision aids*. These aids are designed to overcome the attentional and cognitive limitations that make unaided decisions so difficult. They don't completely mitigate the limitations, but they can help.

Decision aids are paradigms; step by step rules for taming decision complexity. They tell you how to identify the crucial elements of the decision, how to consider each element separately from the others and reach a conclusion about it, how to record the conclusion so it isn't forgotten when you move on to consider the next element, and, after all the elements have been considered and conclusions reached, how to summarize everything in order to reach a final decision.

Extraordinarily elegant decision aids have been developed for the Economic Man view of decision making (Chapter 1); procedures that help decision makers do what he would do were making the decision. Government, corporations, and other institutions pay large fees to decision analysts to help them through the sometimes staggering complexity of these aids; few people can do it on their own, especially for level 3 decisions. One must admire the ingenuity that has gone into elaborating the simple gamble analogy into such an impressive technology. On the other hand, lovely as they are, these decision aids aren't much use to most of us.

As you might suspect, decision aids growing out of the narrative-based decision viewpoint we've been exploring in these pages aren't going to be the same as those growing out of the Economic Man viewpoint. In fact, as you'll see in the next chapter, the paradigms that have thus far grown from the narrative viewpoint are not particularly elegant, if only because they aren't very rigorous. The measurements they require allow for a great deal of vagueness, and the decisions they prescribe are more in the way of suggestions than prescriptions. The only real claims that can be made for them is that anybody can use them

without having to hire an analyst and that most people find them intuitively compelling and seem satisfied with the decisions they make using them. (No doubt, part of this satisfaction results from not having anything better to help them; the aids are at least something to hold onto while making complex decisions.)

But maybe rigor is over-rated. Most of us are less interested in rigorous decision procedures than in successful decision outcomes—more desirable futures. We understand that a decision is only a beginning, not an end; the decisions we make are followed by real actions in a real world that changes rapidly, requiring more decisions to take the changes into account, and then more after that. Therefore, decision precision is a luxury, not a necessity. What you really need is something that gets you started in the right direction. Feedback will quickly swamp the original decision anyway, making you revise the plan and then revise the revision many times over before you get to your desired future. Moreover, the precise nature of desired future probably will change even as you pursue it, necessitating even more revisions in what you're doing to attain whatever it is becoming. Therefore, the function of our decision aids is not to tell you what to do, it is simply to help you figure out which way to push the first domino—how to get moving in the right direction.

Ambiguity and Uncertainty

We'll end this chapter with a few words about ambiguity and uncertainty in decision making. There are three major sources of ambiguity that, together, determine what we usually think of as the uncertainty in a decision. These sources are the extrapolated forecast, the action forecast, and the desired future.

An extrapolated forecast is ambiguous because it is about the ever-encroaching future, which is a moving target, so there seldom is sufficient time for a detailed forecast—and even if there were, everything is

changing so fast the added detail would be of little use. At best, the forecast is a simple extension of the current narrative's plotline, often nothing more than a montage of images and emotions. Perhaps it is best to think of a forecast as a sketch of the future, not a portrait.

An action forecast is ambiguous for the same reasons that an extrapolated forecast is. But, it also is ambiguous because the plan that generated it is probably incomplete. All but the simplest plans contain contingent tactics that will be selected, as needed, during implementation; which makes the plan ambiguous, but in a good way. More ambiguity is added when feedback is ambiguous; is an event's occurrence evidence of good plan successfully implemented or is it due to something or someone else? In much of what we do, feedback is more a matter of oblique inference than concrete fact, and obliqueness adds to the ambiguity.

Desirability is ambiguous because it derives from your primary and secondary values, which aren't always as clear as you'd like to believe. If asked, most of us can list what we think are our values, but our behavior frequently suggests the list is wrong—or at least incomplete. Our ignorance about something that is so very vital results from the fact that values aren't things. They're not even ideas. They're emotions—powerful but somewhat diffuse. As a result, when you are evaluating the desirability of an extrapolated or action forecast, other than generally positive or negative feelings, you often aren't altogether sure what values apply or why. This ambiguity about desirability further compounds the inherent ambiguity of extrapolated and action forecasts because the criteria against which their suitability is tested are ambiguous.

These three sources of ambiguity combine to produce the uncertainty you feel when you are making decisions; uncertainty about

your predictions of the future,¹⁸ uncertainty about the suitability of plans, and uncertainty about whether you actually will find the results of plan implementation as desirable as you think you will. Uncertainty is what leads you to regard very small discrepancies between the features of forecasts and desires as trivial; they may simply be due to ambiguity—rather like measurement error. It also is what leads you to tolerate a few nontrivial discrepancies without rejecting the identity assumption; you know that when two ambiguous entities are compared, apparent discrepancies aren't unlikely, even if the entities are identical.

Theory versus Reality

You may be asking yourself if what is being described here is what actually goes on in your head. I doubt it. In fact, it probably isn't possible to really know what is going on in your head—it's all neurons and electrical impulses, and while studying them is instructive, it usually doesn't tell us as much as we'd like. If you think about it, the study of mental functions is a pretty iffy proposition because the thing doing the studying (scientists' minds) is the same as the thing being studied (minds in general). It has been less than 100 years since scientists began to figure out how to do this with any degree of objectivity and rigor—not a long time in science.

As is true in any science, cognitive scientists, including decision researchers, construct interesting and useful stories, called theories. These theories are composed of hypothetical concepts that are thought to resemble, however roughly, the phenomena to which the theory pertains. The scientists have no illusions that their theories are in any sense “true” because they know that no matter how ingenious they make them, real cognition is bound to be far more subtle and complicated than anything their theories could ever capture; only a fool underestimates the richness

¹⁸ This is the kind of uncertainty that prompted the development of probability theory.

of the human mind. However, if a theory, with all its necessary faults, provides useful insights—if it helps us think more productively about cognition and suggests ways of helping people understand or enhance their cognitive skills—then it puts us further ahead than we were without it, so it is valuable. By and large, whether scientific theories of any kind live or die is based on their usefulness, not on whether they are “true” in some transcendent sense—because we never can know what’s true, only what works.¹⁹

For example, consider a familiar cognitive theory, IQ. Everyone thinks they know what IQ is, but they don’t, because in some sense it doesn’t exist. In fact, IQ began as simply a score on a test that was devised to predict success in school. Over time IQ became a more complicated concept, with scientists arguing about whether “it” was a single entity or made of many components (the latter won), about the degree to which “it” was a function of hereditary or environment (a bit of both), and so on. In fact, different kinds of IQ have been proposed, many of which focus on cognitive skills other than those required to succeed in the classroom.

The point is, IQ is just an oversimplified, but very useful, way of thinking about people’s ability to cope with various cognitive tasks. In reality, that ability is bound to involve bewilderingly complicated cognitive processes that differ for different kinds of tasks, that vary with experience, that mature at different rates, that differ in quality and efficiency from one person to another, and so forth. Nobody really understands what those processes are—not yet. But, until we do (if we ever do), IQ is useful because it allows us to think about cognitive abilities

¹⁹ In fact, theories generally are regarded as true in proportion to their usefulness (in the broadest sense of the word “usefulness”) on the assumption that a theory wouldn’t work well enough to be useful if it weren’t reasonably close to true.

in ways we never had before and it has helped educators tailor ways of enhancing those abilities.

Similarly, the concepts being presented in this book probably aren't true in any transcendent sense. It is unlikely that your nervous system differentiates between level 1 and level 2 decisions and it's even more unlikely that you make decisions using anything as simple as the discrepancy test—real decision making is bound to be cognitively more complicated. But, thinking this way about decisions is useful because it allows us to get a handle on something that we otherwise might think we can't do anything about. In fact, as we shall see, thinking this way helps us design tools that can enhance decision making, which, after all, is a good thing. Therefore, it behooves you to “believe” what you're reading insofar as it can be of benefit to you, but to keep in mind that it is, like any narrative, simply a useful story designed to help you make sense of, and take control of, a sector of your experience.

Progress Report

In this chapter we examined the fourth of the four propositions listed at the end of Chapter 1:

1. Cognitive narratives provide meaning and continuity by linking the past with the present and allowing us to make educated guesses about the future (forecasts). This was addressed in Chapter 2.
2. Based on our past experience and our values, we create rules, some of which are used to make forecasts and plans and some of which are used to evaluate the desirability of the future offered by those forecasts and plans. This was addressed in Chapter 3.

3. When the forecasted future is undesirable, we construct plans that guide actions designed to promote a more desirable future. This was addressed in Chapter 4.
4. *All decisions are about discrepancies and what to do about them.* This has been addressed in this chapter, Chapter 5.

In the next chapter (Chapter 6) we will examine decision aids for level 2 and 3 decisions.

Sources and Further Reading

Beach, L. R. (1964). Recognition, assimilation, and identification of objects. *Psychological Monographs*, 78, (Whole No. 583).

Beach, L. R. (1990). *Image theory: Decision making in personal and organizational contexts*. Chichester, UK: Wiley.

Beach, L. R., and Strom, E. (1989). A toadstool among the mushrooms: Screening decisions and image theory's compatibility test. *Acta Psychologica*, 72, 1-12.

Connolly, T. (1988). Hedge-clipping, tree-felling and the management of ambiguity: The need for new images of decision-making. In L. R. Pondy, R. J. Boland, Jr., & H. Thomas (Eds.), *Managing ambiguity and change* (pp. 37-50). NY: Wiley.

Kahneman, D. (2003). A perspective on judgment and choice: Mapping bounded rationality. *American Psychologist*, 58, p. 697-720.

Miller, G. A. (1956). The magical number seven, plus or minus two: Some limits on our capacity for processing information. *Psychological Review*, 63, 81-97.

Myers, D. G. (2007). The Powers and Perils of Intuition: Understanding the Nature of Our Gut Instincts. *Scientific American Mind*, 18, No. 3, 24-31.

- Peeters, G. (1971). The positive-negative asymmetry: On cognitive consistency and the positivity bias. *European Journal of Social Psychology, 1*, 445-474.
- Peeters, G. (2002). From good and bad to can and must: Subjective necessity of acts associated with positively and negatively valued stimuli. *European Journal of Social Psychology, 32*, 125-136.
- Peeters, G. (2003). Positive-negative asymmetry in the human information search and decision-making: Five basic and applied studies on voting behavior. S. P. Shohov (Ed.), *Advances in psychology research* (pp.61-92). NY: Nova Science Publishers.
- Plessner, H., Betsch, C., & Betsch, T. (Eds.) (2007). *Intuition in judgment and decision making*. Hillsdale, NJ: Lawrence Erlbaum Associates.

Chapter 6: Decision Aids

In Chapter 1 we discussed the differentiation between narrative thought and paradigmatic thought. Then, in Chapter 2, we discussed how narratives give continuity and meaning to experience and allow you to forecast the future. For the most part, narratives are logical, but not always, not rigidly, and they tolerate a good deal of ambiguity. Over the course of human history, narrative thought has been supplemented and extended by the invention of paradigms for thinking about a wide range of topics and problems for which narrative thought doesn't do a good job or can't handle at all. Foremost among these paradigms are mathematics and the scientific method. Universities, which began as bastions of narrative thought, have become even bigger bastions of paradigmatic thought as philosophy, theology and literature have taken a back seat to the sciences, engineering, and "how-to" professional training.

Although paradigms are designed to supplement narratives, they require a substantially different way of thinking. Narrative thinking is integrative; it absorbs details into an ongoing story that ties everything together and links it to your subjective experience. Paradigmatic thought is analytic; it breaks things down into details that are to be treated in prescribed ways and links it to abstract theories and empirical facts. Links to your private experience, if they exist at all, are through narratives. Atomic theory, for example, is only linked to most atomic physicists' private experience through narratives about atoms, not through direct experience. Other than the threat of a bomb or our reliance on electronic devices, atomic theory has little or no obvious linkage to the private experience of the rest of us.

Judging by how difficult it is for people to learn to think analytically, it is clear that, for most of us, paradigmatic thought is less congenial than

narrative thought. This may be true in part because we all learned narrative thinking through immersion in a sea of stories when we were young; much as we learned to speak our native language by being immersed in it. In contrast, paradigmatic thought was, for the most part, learned in school—after we already had 5 or more years of immersion in narratives.²⁰

Aids for Narrative Decisions

Decision aids are decision researchers' contribution to the catalogue of paradigmatic inventions for extending narrative thought. They consist of step-by-step procedures for making a decision. In this chapter, we are going to examine two decision aids, both of which use the enriched version of the discrepancy test from the previous chapter (Chapter 5).

The paradigmatic analysis of a decision consists of procedures for decomposing the decision into its component parts, procedures for dealing in a consistent and reasoned manner with those parts—usually computationally—and procedures for recomposing the results into a recommended decision. Of these, decomposition is the most difficult. Computational procedures are cut and dried, as are the procedures for recomposition of the results, but the procedures for decomposition, while

²⁰ If you don't use a paradigm frequently, you tend to forget the details. For example, I've forgotten how to do the calculations required by most statistical tests because I haven't had to do them since computers came along. On the other hand, I retain a detailed narrative about what tests are appropriate, when, and how to interpret their results. Similarly, I took trigonometry, never used it, and remember nothing about it—but almost 60 years later I retain a narrative about where the classroom was and what it looked like, and that the course was taught by Miss Gibson, who was far more enthused about trigonometry than I was. And, I once memorized a paradigm for memorizing things, used it once, decided it was too much work, and have completely forgotten it—but 58 years later I retain a narrative about learning it in an evening class that met on the downtown campus of the University of Denver; a class that was taught by a Mr. Dolan, and that required us to memorize the contents of each page of an issue of a tabloid-style newspaper, which probably was the Rocky Mountain News. These narratives have a strong visual component; I can see them in my mind's eye, as well as a verbal component that tells the story.

crucial, are more arbitrary. Therefore, any decision aid must focus heavily on helping users think objectively and analytically in order to achieve the required decomposition. The more complex the decision is, the more crucial it is for the aid to provide help with decomposition.

Let's be clear about what decision aids are and aren't—at least what the ones we'll be talking about are and aren't. They aren't a substitute for thinking. Quite the opposite; they are designed to help you think more purposefully and more clearly, which means that they require work. Second, they don't make the decision for you; they help you think through the decision and make it for yourself. Indeed, if you devote the effort that you should to using the aid, it is very likely that you'll know what you want to do far before the computations are done and the aid's recommendation is known. Decision aids, like almost all paradigms, are aimed at educating your intuition—once you've thought about the decision, examined its various aspects and weighed their implications, you'll pretty much know what will best satisfy your values. When all is said and done, decisions boil down to figuring out how to pursue a desirable future—it doesn't have to be a super future, just something with which you can be comfortable. The aid is merely to get you started in the right direction.

Because decision aids are about helping you think clearly, it is important to make them transparent—you don't want to get so bogged down in the details of using the aid that you're distracted from the decision itself; that just befuddles intuition rather than educating it. Therefore, the decision aids we'll be discussing are as straightforward as possible, the only obscurity being a little mathematical magic at the end. (And, you really don't have to understand the magic; just follow the directions in Appendix 2 and all will be well.) Aside from this, the measurements and computations are so simple you can do them in your

head; although, if you're as bad at math as I am, you might want to use a calculator.

Formally, a decision aid is a set of procedures that derive from a model, as well as instructions for application of the procedures. The procedures are about decomposition, computation, and recombination of the decision and about how to interpret the results. In addition, because one of the goals of using an aid is to transform a complex decision into a simple decision and an abstract task into a concrete task, additional tools may be required; Figure 6.1 contains one such tool, a *decision sheet*, which requires additional tools; a pencil and, just to be safe, a calculator.

Fig. 6.1: Example of a decision sheet.

Option _____

Features of	Desired Future	Importance	Magnitude	Impt X Mag
	i_c	m_c	$i_c m_c$	
1	_____	_____	_____	
2	_____	_____	_____	
·	·	·	·	·
·	·	·	·	·
n	_____	_____	_____	
D = $\sum i_c m_c / x$ =				_____

A Decision about the Extrapolated Future

Before we begin, recall from Box 2c, in Chapter 2, that one source of the desired future is the “repaired forecast;” what the forecasted future would be if its faults were repaired so it would be desirable. That will be the source for our first example; subsequent examples will illustrate desired futures from other sources.

Let’s suppose that you have been thinking about your present job, wondering if it is leading anywhere and if you’re getting bored with it. As you ponder, you daydream about the future; that is, you formulate an extrapolated forecast of what things are likely to look like in a couple of years. Just then the phone rings and your attention is diverted to a more pressing task.

Later that day, you return to your daydreams and try to be honest with yourself about how you feel about your future in your present job. You toss things around in your head and realize that you can’t quite get a grip on them. So, you drop by my office and I give you a blank decision sheet like the one in Figure 6.1. After I show you how to use the sheet, you go home, sit down with a cup of coffee and begin work:

- Begin by reviewing your extrapolated forecast and writing a short description of what your job will be like in two years; describing it as you would to a friend. Be as exhaustive as possible. Then go back and underline each descriptor in your description. These are the features of your forecasted future.
- List the forecast's features in the column on the left side of the decision sheet.
- Consider each feature on your list in terms of how important that particular feature of your future is to you. Indicate importance by assigning the feature a number from 1 to 3 (1 = low importance, 2 = medium importance, 3 = high

importance), and writing the number in the column labeled **i_c**.

- Consider each feature on your list in terms of how discrepant it is from what you feel it should be. Assign a 0 (zero) to any feature that is already close enough, and a number from 1 to 3 (1 = small discrepancy, 2 = medium discrepancy, 3 = large discrepancy) to indicate how far each of the others falls short of being what you feel it should be. Write your answers in the appropriate spaces in the column marked **m_c**.
- For each feature, **d_c**, multiply the numbers in the two columns and write the product in the corresponding blank in the column on the right of the sheet, **i_cm_c**.
- Finally, add the products in the right hand column, divide the sum by 9, which is the scaling unit, and write the summary number, **D'**, at the bottom of the sheet. **D'** indicates the degree to which you find the future offered by your job to be *undesirable*. (The magic scaling unit 9 and how to interpret **D'** are both explained in Appendix 2).

Of course, even if your analysis demonstrates that your job is pretty undesirable, it doesn't mean you should simply quit. Your first option is to figure out a plan for repairing its flawed features. Even if it can't be repaired, you shouldn't necessarily quit. It may be that the good features are good enough to make you endure it. Therefore, it is prudent to look over the good features on your list and balance them against the bad features; if the good overbalance the bad, you may think it reasonable to stick with the job. However, this doesn't make the job any more desirable, it just means that you are willing to compromise, to endure an undesirable job because the good things are good enough—but it is unlikely that you'll feel good about it. (Of course, we all have endured jobs we didn't

like; which is why our employers paid us. Even a bad job is better than starving.)

A Decision about an Opportunity

Suppose that about the time you were pondering your future in your present job, you are offered a job at a firm on the other side of the city. Now you have a choice; stay with the old job or take the new job.²¹ So, you come back to my office and I give you another decision sheet and tell you to analyze the new job just the way you did your old job, starting by imagining what the future would be if you accepted the new job.

After you've done the analysis for the new job and written its **D'** at the bottom of the sheet, you are ready to make the choice. The decision rule is simple, if there is a meaningful difference between the **D'** for your present job and the **D'** for the new job, the job with the smallest **D'** is the most desirable.

Again, the good features of the new job may weigh in favor of taking it even though it is less desirable than your present job—but that doesn't change the fact that the new job is less desirable than your present job. In fact, if you have already decided your present job isn't very desirable and the new job turns out to be even less desirable, then taking the new job because of something good (colossal pay, for example) is never going to make you like the new job any better—it still is undesirable, you're just willing to endure it.

Notice that we didn't compare the features of the two jobs directly, we evaluated the desirability of each independently and then compared the results. This is because it isn't a question of choosing the more desirable of the two; they both may be undesirable. It is a question of appraising the desirability of each and then making the choice based on whether either or both are desirable and the mitigating circumstances

²¹ In this example the future offered by the new job is a suggested desired future (Box 2c).

(the good features of each). In short, the decision aid doesn't tell you what to do, it merely helps you sort out your feelings about the alternatives; the choice flows naturally from that knowledge combined with knowledge about mitigating circumstances. The cold fact is, the least desirable job could end up being the best choice because of mitigating circumstances—but you're not likely to be happy about having chosen it.

A Choice among Plans

Let us assume that you found the new job more desirable than the old job, whose good features couldn't over-balance its undesirability. But before you accept the new job you have to make another decision. Your old job was within walking distance of your home, but the new job is across town. Because you are unwilling to move, which would necessitate moving your children to inferior schools, you have to commute. So, you must make a decision about how best to do it; each alternative mode of transportation is, in fact, a plan for getting to work. So, your decision is a choice among alternative plans for getting to work.

You have driven to work for jobs you've held in the past, and most people drive to work, so driving is the first plan that comes to mind. But, you realize that, if nothing else, driving is expensive, so you think about it a little more and end up with three competing plans: drive, transit, or carpool. The question is, which of the three methods (plans) of commuting, if any, is the best choice for achieving the future offered by your new job.

Beginning with driving, apply the aid to each of your plans individually. After all three have been analyzed, compare the results and make your decision.

Drive

- Begin by imagining what your experience would be if you drove to work. Then write a description of the experience (your action

forecast). Go through your description and underline or highlight each descriptor; these are the defining features of your forecasted future for driving.

- Create a decision sheet like the one in Figure 6.2.
- Write the name of your first plan—drive—at the top and list the features of your forecasted future down the left side of the decision table.
- Consider each feature on your list in terms of how important it is to you. Indicate importance by assigning the feature a number between 1 and 3. Write the numbers in the column labeled i_c .
- Consider each feature on your list in terms of how discrepant it is from what you feel it should be. Assign a 0 (zero) to any feature that is already close enough, and numbers between 1 and 3 to indicate how far each of the others falls short of being close enough. Write the numbers in the column marked m_c .
- For each feature, multiply importance, i_c , and discrepancy magnitude, m_c , and write the product in the column on the far right of the decision sheet labeled $i_c m_c$.
- Compute the sum of $i_c m_c$, divide it by the scaling unit, 9, and write the result at the bottom of the column. This is the degree, D' , to which driving is an undesirable plan for commuting to the new job.

In the example, $D' = 2.66$, which is equivalent to having 2 highly important large-discrepancy features in *any* analysis you might perform and a smattering of lesser ones that contribute the .66 part of the total of 2.66.

Transit and Carpool

Create a decision sheet for Transit and another for Carpool and perform the analysis for each by following the steps listed above for Drive.

Fig. 6.2: The decision sheet for a plan for commuting to a new job.

Option: Drive

Features of

Desired Future Importance Magnitude Impt x Mag

	i_c	m_c	$i_c m_c$
Money	<u>3</u>	<u>3</u>	<u>9</u>
Travel Time	<u>3</u>	<u>1</u>	<u>2</u>
Reliability	<u>3</u>	<u>0</u>	<u>0</u>
Comfort	<u>3</u>	<u>0</u>	<u>0</u>
Privacy	<u>2</u>	<u>0</u>	<u>0</u>
Traffic	<u>3</u>	<u>3</u>	<u>9</u>
Convenience	<u>1</u>	<u>1</u>	<u>1</u>
Safety	<u>1</u>	<u>1</u>	<u>1</u>
Environment	<u>1</u>	<u>2</u>	<u>2</u>

$$D = \sum i_c m_c / x = 24/9 = \underline{2.66}$$

Choice

After completing all three analyses, you will have three D' , one for driving, one for transit, and one for carpool.

- Eliminate from further consideration any of the three plans for which the D' is so large that you feel uncomfortable about its desirability as a way of getting to the new job

- If no plan survives, reject the new job because you have no way of getting to work—unless you want to go back and consider bicycling.
- If only one plan survives, choose it and accept the new job as your desired future.
- If two or more plans survive, choose the one for which **D'** is least and accept the new job.

What About Costs?

It may seem as though you shouldn't choose a plan until you've balanced its implementation costs against the future it promises—perhaps it is so expensive you would rather settle for another, less costly, plan or a less desirable future. In fact, the features on the left of the decision table (the action forecast) should include those costs because when you arrive at the future you will already have expended the money, energy, time, etc. to implement the plan, thereby impoverishing that future by the amount of the expenditure. In a sense, you always are borrowing implementation costs from the future, so the future is a little less desirable for costly plans than for less costly plans.

In our example about commuting, implementation costs are included in:

- Driving your car will involve *money* for upkeep, parking, gas, and so forth.
- Using transit will involve *money* for fares, convenience in terms of scheduling, *privacy* in terms of sharing with other transit users, and, perhaps, *comfort* in terms of weather and sometimes having to stand rather than sit.
- Carpooling will involve *money* for sharing expenses, *convenience* in terms of having to pick up others and/or being picked up, *privacy* in terms of sharing a car with others.

All of these are costs. If you think carefully about the features for your list, you shouldn't have to balance costs against **D'**.

Aiding Complex Decisions

Choosing between jobs and choosing between plans for commuting to work were both level 2 decisions because neither of them involved many defining features. We turn now to level 3 decisions, which have more defining features and therefore require an expansion of the aid in Figures 6.1 and 6.2.

Expansion consists of decomposing the level 3 decision into a number of level 2 decisions, subjecting each level 2 decision to the kind of analysis illustrated above, combining the results of these analyses into an overall result, and using that overall result to reach a decision. In short, the strategy for level 3 decisions is the same as for level 2 decisions—it merely is applied to parts of the level 3 decision rather than to the whole decision.

As we did for level 2, we'll illustrate the decision aid for level 3 decisions with an example.²²

Herb is a (fictitious) 38 year old man who has been married for seven years to Fran, who is 33. They have one child, a 3 year old girl named Angela. Their marriage is sound, although Herb thinks they may have been drifting apart a little lately. They both have demanding jobs that they like. Herb has been working hard to get promoted to regional manager, which will demand even more of his time because he will have to do a lot of traveling. Fran has a very responsible job at a cancer research laboratory, where her particular skills are central to the laboratory's ability to obtain research grants. She loves her job, although

²² In this example the future offered by Fran's mother's pressure for Herb and Fran to have another child is a suggested future (Box 2c).

things become very stressful whenever existing grants come up for renewal or new ones must be submitted for funding.

Because Fran's mother clearly wants more grandchildren, Herb and Fran have discussed having another child, but they never seem to get very far. They agree that if they are going to do it, they should do so soon because they think four years is a good space between children. They both derive great satisfaction from raising Angela, although neither of them is home as much as they feel they should be. Angela has been in daycare since she was very small. As a result, she has great social skills and seems to have far more friends than either of her parents.

Last month Fran's sister gave birth to her third child, setting Fran's mother off again. As a consequence, Fran has begun to wonder aloud if they should stop avoiding the decision about another child; either decide to do it or decide not to do it and make it clear to her mother what they decide.

The question Herb and Fran must answer in order to know what to do is: Which course of action presents the more desirable future, not having a second child (the extrapolated forecast) or having a second child (an action forecast)? Herb and Fran freely admit that they like their lives and that if it weren't for a few little worries, they would be perfectly willing to continue just the way things are. However, those few little worries are indeed worrying, and Fran's mother's pressure for them to have another child is getting more urgent. They know that if they decided not to have another child that Fran's mother would accept it, but they aren't totally sure they don't want a second child, so they haven't said anything to her. The fact is they need to decide and then deal with Mom afterward.

Here is what I'd tell Herb and Fran if they asked me how to make their level 3 decision about having another child.²³

The Extrapolated Future

- Think about how future will be if you don't have a second child (your extrapolated future without another child). Describe this future in writing and then highlight all of the descriptors; these are the forecasted future's features.
 - Transcribe each descriptor onto an index card.
 - Sort the features (cards) into sets of similar or related features and give each set of features an identifying label.
 - Sort the labeled groups of sets into larger groups.
 - Write the labels and features in the column on the left of a decision table (Figure 6.3).
- Evaluate the importance of the groups, sets, and features of your desired future.
 - Assign each group (I and II in Figure 6.3) a number from 1 to 3 to indicate how important the contents of the group are to you (1 = low importance and 3 = high importance).
 - Assign each set (A,B,C in Figure 6.3) in each group a number from 1 to 3 to indicate how important contents of the set are to you.
 - Assign each feature in each set (1 through 19 in Figure 6.3) a number from 1 to 3 to indicate how important it is to you.
- For each individual feature, indicate the magnitude of any discrepancy between your forecast of what you think it *will* be and what you feel it *should* be by writing a number from 0 to 3 in the column labeled **m_c**. Assign a 0 (zero) to any feature that is already

²³ Remember, the desired future is the touchstone. We don't compare the plans directly; we evaluate the desirability of each plan's action forecast relative to the desired future and then compare the results of the evaluations in order to choose a plan.

close enough, and numbers from 1 to 3 (1 = small discrepancy and 3 = large discrepancy) to indicate how far each of the others falls short of being close enough. Write the numbers in the column, m_c .

- Calculate the importance of each feature of the desired future relative to the other features:
 - Multiply the number you gave to group I by the ratings you gave each set in the group (A, B, C).
 - Multiply the number you gave to each feature by the product you calculated in the previous step for the set (A, B, or C) of which the feature is a member. Write the product in the column labeled i_c .
 - Repeat for all features.
- For each feature, multiply the number in the Importance column, i_c , and the number in the magnitude column, m_c , and write the product in the column on the far right of the decision table, $i_c m_c$.
- Sum the numbers in the column at the far right, divide by the scaling unit, 81 (see Appendix 2), and write the result at the bottom of the column. This is D' , the undesirability of a future without another child.

The Action Forecast

- Create another decision sheet and perform precisely the same analysis for a future that includes a second child.
- Compare the undesirability, D' , of a future without a second child with the undesirability, D' , of a future that includes a second child.

Decision Rule

- Because the two options, not having a second child and having a second child, are mutually exclusive and exhaustive (you have to do one or the other, although there is a middle ground, waiting):

- If not having the child is undesirable and having it is desirable, choose to have the child. If not having the child is desirable and having it is undesirable, choose not to have the child.
- If both options are desirable, select the one with the smallest **D'**.
- If both plans are undesirable, select the one that leaves you with the greatest freedom to revisit the decision in the future; in this case, not having a second child right away but leaving the option open.

Sensitivity Analysis

The numbers you write on the decision sheets are unlikely to be very precise or reliable—who among us can claim to clearly know the features of the future, exactly how important they are, or precisely how discrepant they are from our values and preferences? Fortunately, the analysis tends to be fairly insensitive to imprecision, although not always. You can see how sensitive/insensitive it is by arbitrarily changing the numbers you are least sure about and doing the computations again. If arbitrarily increasing or decreasing the evaluations by ± 1 changes the decision prescribed by the analysis, then you should give more careful thought to each number so you can make the analysis more precise.

Back to Herb and Fran

So, Herb and Fran sat down together and jointly filled out the decision sheets. They could have each filled them out separately and then compared the results, but they decided to do it together and to talk about each step until they agreed on what should be written down. They disagreed on some points, so some compromising was required. (Also, when they did the sensitivity analysis, afterward, they first put in one person's numbers and then the other's to see if they disagreed enough to affect the overall results—they didn't.)

Figure 6.3: The decision sheet for Herb and Fran's decision about having another child.

Option: <u>Don't Have Another Child</u>			
Features of Desired Future	Importance	Magnitude	Impt x Mag
	i_c	m_c	$i_c m_c$
I. Us (3)			
<u>A. Herb (3) [9]</u>			
1. Good Health (3)	1 <u>27</u>	1 <u>0</u>	1 <u>0</u>
2. Grow & Mature (2)	2 <u>18</u>	2 <u>1</u>	2 <u>18</u>
3. Fulfilling Job (3)	3 <u>27</u>	3 <u>0</u>	3 <u>0</u>
<u>B. Fran (3) [9]</u>			
4. Good Health (3)	4 <u>27</u>	4 <u>1</u>	4 <u>27</u>
5. Grow & Mature (2)	5 <u>18</u>	5 <u>2</u>	5 <u>36</u>
6. Fulfilling Job (3)	6 <u>27</u>	6 <u>0</u>	6 <u>0</u>
<u>C. Our Marriage (3) [9]</u>			
7. Material Well-being (3)	7 <u>27</u>	7 <u>0</u>	7 <u>0</u>
8. Becoming Closer (3)	8 <u>27</u>	8 <u>2</u>	8 <u>54</u>
9. Time together (2)	9 <u>18</u>	9 <u>1</u>	9 <u>18</u>
II. Angela, Relatives & Friends (3)			
<u>D. Angela (3) [9]</u>			
10. Good Health (3)	10 <u>27</u>	10 <u>1</u>	10 <u>54</u>
11. Grow & Mature (3)	11 <u>27</u>	11 <u>2</u>	11 <u>54</u>
12. Good Childhood (3)	12 <u>27</u>	12 <u>1</u>	12 <u>27</u>
<u>E. Relatives & Friends (2) [6]</u>			
13. Becoming Closer (2)	13 <u>12</u>	13 <u>1</u>	13 <u>12</u>
14. Time Together (1)	14 <u>6</u>	14 <u>2</u>	14 <u>12</u>
D' = $\sum i_c m_c / x = 285/81 = \underline{3.52}$			

They began by thinking about what their future would be if they didn't have a second child, discussing many things while slowly constructing a description of the future. When the descriptors and themes were translated to cards and sorted, two themes emerged—one was **Us** and the other was **Angela, Relatives & Friends**. At first they thought they ought to include Angela in **Us**, but they realized that their marriage was separate from her in that it existed before she came along and would endure after she left home.

Next they concentrated on **Us**, including Herb's future, Fran's, and their joint future and their respective components. Then they concentrated on **Angela**, health and growth & maturity, and then on **Relatives and Friends**, closeness and time spent with them.

Having identified the features of the extrapolated future, Herb and Fran turned their attention to evaluating the importance of each feature. They started by assigning numbers from 1 to 3 to the two big groups—**Us** and **Angela, Relatives & Friends** to indicate the relative importance of each, writing their answer in the adjacent parentheses. Then, within each group they assigned importance to each set (Herb, Fran, Our Marriage) to indicate their relative importance, again writing their answers in the adjacent parentheses. Finally, they assigned importance to each of the features within each set, writing their answers in the adjacent parentheses.

The reason for doing these different levels separately is to break the task down into manageable bites—it is difficult to know how important one's job is relative to, say time with friends, so it is best to compare related things with each other. In other words, they decomposed their desired future into meaningful chunks (groups) so they could evaluate the chunks and the contents of the chunks, simplifying evaluation of each chunk so it resembles a level 2 decision.

Then they began the computations to derive the relative importance of each feature. To do this, they multiplied the number in the parentheses adjacent to the name of the first group (**Us**) by the number in the parentheses adjacent to the name of the sets and wrote the answer in the brackets adjacent to the name of the set. They repeated this for the next set, and the next, and then repeated the whole thing for the second group (**Angela, Relatives & Friends**).

Finally, they multiplied the number in each bracket by the number in the parentheses adjacent to each feature, writing the answer in the appropriate blank in the column for importance, **i_c** .

Then they evaluated discrepancy magnitude, **m_c** , performed the multiplications, summed the products, and divided by the appropriate scaling unit, in this case 81, to arrive at **$D' = 3.52$** for not having a second child. They laid out a scale from 0 to 19 on a piece of paper²⁴ and marked the place on the scale corresponding to their **D'** for not having a second child. Recall from Appendix 2 that a **D'** of 3.52 is the equivalent of having 3 highly important large discrepancy features in any analysis one were to perform, including this one, with a smattering of lesser ones that account for the .52.

Next they went through the same procedure for the second option, having a second child—the joys and tribulations of raising a baby and of the impacts having another child would have on their lives—both good and bad. Then they constructed a new decision sheet, with a list of features from their forecast of the future were they to have another child. As they considered what numbers to write in the blanks, Herb and Fran noticed that importance and magnitude reflected their fears about their jobs as a result of having another child—Herb didn't feel he could accept the promotion if they had a baby because he would have to be away

²⁴ Because there were 19 individual features in the column on their decision sheet.

from home too much and Fran would require his help if they had two children. Fran felt she would have to cut back her hours at the lab in order to devote more time to both her children, perhaps showing her employers at the lab that she wasn't as crucial as they thought. As a result, they both expected their respective contributions to their joint income to decline just when more money would be needed for the new baby.

On the other hand, they thought that Angela might be better off with a brother or sister, although she had it pretty good as an only child and would have to share more than she might want to. Herb and Fran agreed that having a new baby would bring her closer to her mother, who would be delighted—which, to a large degree, was why they were even considering a second child. It also would bring their family closer to their respective brother's and sister's families, both of whom had three children, as well as to their friends, almost all of whom had children.

After writing numbers in the blanks for the second decision sheet they computed D' ; the undesirability of the forecasted future if they were to have a second child. Then they placed a second mark on their 0 to 19 scale to indicate the D' for having a second child; allowing comparison of the undesirability of a future without the child to that of a future with one.

Even before they compared the two D' —indeed, even before they finished the computations for having a second child—Herb and Fran knew what they were going to decide, but they finished the job anyway. Then they poured themselves fresh cups of coffee and talked about the results. Did they agree with the analysis? How would they feel if they did what the analysis prescribed? How would they feel if they did the opposite? In short, having thought through all of the issues, having discussed them and reached agreement about them, what did their intuition tell them to do?

Of course, this is just a fictional example, so I don't have to tell you what they did, you'll just have to figure out what you'd do if it were you.

The whole point of any analysis is to think through the decision—what it means to stay with the status quo and what it means to change. As I said before, the goal of using a decision aid is to educate your intuition so that when you sit down with that cup of coffee afterward, you can say to yourself, “Yes, that really is what I want to do; it will be for the best.”

Some Rules of Thumb

The most difficult part of using a decision aid is making sure you include all the defining features. Fortunately, merely attempting to do so tends to make you think more deeply, which reveals the features. As you ponder, simply list the features on a piece of paper. Then, when you find yourself thinking about things you’ve already thought about once or twice before, stop and arrange the features like we did in Figure 6.3.

In the course of arranging the features for the decision sheet, and subsequently evaluating them, you’ll find that some of them are redundant, some are too much like others to warrant listing both, some have been left out. Use your good sense to straighten things out, but always pay attention when your intuition tells you something doesn’t fit right. Keep revising until you feel comfortable with the list and its arrangement. (Actually, this seldom takes as long as you might imagine.)

Then start entering your evaluations on the decision sheet. As before, listen to your intuition. Discomfort about an evaluation means it probably isn’t right. Discomfort sometimes is a cue that the list of features isn’t structured correctly and you therefore can’t make a clean evaluation. If you think this is the problem, go back and repair the list of features. Discomfort also is a cue that you really don’t understand a feature or that it in some way doesn’t belong on the list; if so, think about it more and make the appropriate changes.

Finally, when you’ve done all the calculations, made a 0 to n scale, and compared the options’ **D**, and reached a decision, there is another

little test you can do to help you know if the decision is right. Ask yourself how you would have felt if the decision had come out the other way. If your decision analysis adequately reflects your thoughts about the various plans, you shouldn't have more than minor qualms about your choice. But, if you've fudged somewhere, or the analysis has left something out, you may not feel so comfortable—so it's back to the drawing board.

Finally, our use of 3 and 4 point scales may seem to oversimplify measurement. But most of us aren't able to reliably make subtle judgments about importance and magnitude, so using scales with numerous gradations doesn't add precision, it just adds the illusion of precision. Moreover, when scales with more gradations are used, say scales from 1 to 10, most of us tend to shy away from using the higher values. The result is that our judgments imply that none of the features are all that important and none of the discrepancies are particularly large—which probably isn't true or we wouldn't be making a decision at all. After all, if the future isn't important or the forecast is pretty much like the desired future, why bother? And we certainly wouldn't bother using a decision aid. The mere fact that the decision is worth bothering about suggests that some of the features are important and that discrepancies are to be taken seriously, so it is a mistake to use scales that obscure this fact. That said, however, if you feel that you can make more nuanced judgments, you are not limited to the scales we've used in these examples. You can use any scales you wish, but remember that you have to adjust the scaling unit to fit the selected scales, as described in Appendix 2.

What Has Been Accomplished

In this chapter we have seen how the discrepancy test leads to paradigms for aiding decision makers in their efforts to clarify their thinking about their forecasts and their desires for the future, and reaching a

decision. No doubt these aids look forbidding—especially the magic scaling unit—but it really isn't all that complicated. Their major value is in encouraging you to lay your decision out on paper so you don't lose track of one issue while thinking about others. The calculations can be important if you really are having trouble deciding, because they make you think through how you feel about the various features of the forecasted future in light of your desires (values and preferences). But the fact remains, by the time they've listed its features, most people have a pretty good idea about how they feel about the forecasted future, and they know which features must be repaired if they are going to make ensure that their future turns out well. This is true no matter what the source of the desired future (Box 2c). Remember, the point of a decision aid is to help you sort out your thoughts and feelings, not to tell you what to do. Even at that, it only gets you started in the right direction. After that its up to you to keep on track toward the future you desire.

Progress Report

In Chapters 2 through 5, we examined all four of the propositions listed at the end of Chapter 1. In this chapter (Chapter 6), we have shown how the four propositions, together with the Simplifying Strategy, translate into procedures for making conscious level 2 and 3 decisions. In Chapter 7, we will see how everything said thus far applies to organizational decisions.

Sources and Further Reading

Klein, G., Ross, K. G., Moon, B. M., Klein, D. E., Hoffman, R. R., & Hollnagel, E. (2003). Level 2 cognition. *IEEE Intelligent Systems (May/June)*, 81-85.

Chapter 7: Organizational Decisions

Most of the decisions we've discussed in previous chapters were private. That is, they were decisions made by individuals in pursuit of a desirable future for themselves. But, even private decisions aren't made in a vacuum. You're a member of numerous groups—family, friends, co-workers, the community at large, and so on. This means that in most decisions you make—most level 2 decisions and probably all level 3 decisions—you try to take into account these other people's *interests* (which are their narratives, forecasts, plans, expectations, values, and preferences). Failure to do so leads to trouble when the actions you take in pursuit of what you think is the desired future interfere with other people's pursuit of what they think is the desired future. The resulting conflict hinders everyone's plan implementation and, if bad feelings result, may tarnish the desirability of everyone's future.

Of course, you can't always figure out others' narratives, nor anticipate what they think are the organization's forecasted and desired futures, but in the interest of social tranquility, you usually try not to be too obviously selfish and insensitive. Because you make this effort, very few decisions that you make can be thought of as purely private; they almost always reflect what you perceive to be the interests of those around you.

Organizations

For our purposes, any group of two or more persons who are affiliated in some way constitutes an organization. Examples are a marriage, a family, a weekly poker game, a sports team, the sales staff of a furniture store, the office staff at local charity, a global corporation, a governmental unit, a political party. Affiliation doesn't mean the members of the group all like each other or even that they work well together to attain some common future—dysfunctional groups are organizations too. Moreover, a subset of persons within a larger affiliative group also counts

as an organization—a social clique on a shipping dock or a cabal of employees conspiring to get their supervisor fired could both be organizations within a department, and the department is an organization within the company, which is an organization within its industry, and so on. In short, all affiliative groups, big or small, functional or dysfunctional, unique or encapsulated within a larger organization, are organizations.

Decision Agents

It is common to talk as though organizations make decisions, but they don't; people do. When we say, for example, that a big automobile manufacturer decided to close one of its assembly plants, we actually mean that a single person or group of people, acting as decision making agents of the organization, evaluated the extrapolated future, found it flawed, and chose to close the plant as part of a plan to correct those flaws. In short, a member of an organization may act as a single decision making agent of the organization or as a member of a decision making group that acts collectively as an agent of the organization.

Single Agent

Most positions of responsibility within an organization require the jobholder to make decisions for a specific sphere of the organization's activities. In a business, the marketing manager is in charge of marketing and makes decisions related to it. In a social club, the various officers and committee heads have their specific areas of responsibility and, within their area, they are allowed to make routine decisions on behalf of the club. In both a business and a social club, as in most organizations, the person responsible for a specific sphere is given latitude, but it is assumed that he or she will factor in the interests of others, usually by seeking input before decisions are made.

Factoring in others' interests has two goals. One is to make decisions reflect the interests of the organization as a whole, insofar as that is

possible. The other is to reduce conflict and facilitate cooperation during implementation of the decision.

You want your decisions to reflect the interests of the organization because you are acting as a decision making agent for the organization, not as an independent agent or an agent for a coalition. True, most decisions are colored by your own interests as well as the interests of people with whom you are closely affiliated. But, ideally, you try to be as evenhanded as possible, making sure that the broader organization's interests are served.

Because the organization is made up of many people, it may seem almost impossible to factor the diversity of interests into a decision. However, this seldom is as difficult as it appears. Every organization, even a marriage or a coffee klatch, has a culture that reflects the core values and preferences shared by its members. This culture defines what is good and bad for the organization and its members, what behaviors are and are not acceptable, how things are generally done, what qualifies as success or failure, how success will be rewarded and failure penalized, and what members can expect from each other and from the organization as a whole.

An organization's culture evolves over time, but at each step in its evolution it provides the organization's members with a common core of understanding that binds them together as a group. Therefore, when a decision maker is entrusted with the responsibility to make decisions for the organization, knowledge of the culture allows him or her to factor in the narrative, plans, expectations, values, and preferences that make up the culture. This doesn't mean that everyone will be pleased by the decision or support it; after all the culture is the general viewpoint and individuals differ in the degree to which it represents what they personally think and feel. But, it represents the views of most members of the organization and

even dissidents will subscribe to some aspects of it. Therefore, decisions that incorporate it will be most easily accepted by most members, eliciting the support that you need for implementation.

The problem is, of course, that no matter how hard you try, some people will argue about whether your decision actually reflects the organization's interests—as they understand them. If these critics are more powerful than you, you usually must accommodate them, even if it is against your better judgment. Power may be in terms of holding a higher-level position in the organization than you, of representing an important constituency, of being able to influence the opinions of others through charm or persuasiveness, of possessing superior technical knowledge, of having control over resources needed for implementation, or of possessing moral authority. In virtually every organization, relative power among the members is reflected in their relative influence on decisions. When the power is possessed by people who actually have something to contribute to the decision, this can be a good thing. When the power is possessed by people who merely enjoy exercising it, it usually isn't so good. Inappropriate intrusion of power into decision making is what most people are referring to when they bemoan "playing politics."

Power's biggest influence on decisions is the ability (in some cases, the right) to impose what the desired future should be and/or to constrain plans. This isn't necessarily as sinister as it may sound; a boss' job is to set goals and direct his or her subordinates toward particular ends—both of which define the desired future

Collective Agent

Human history can be read as a search for a satisfactory mode of organizational decision making; either a single individual, which runs the danger of that individual becoming a despot, or a collective, which runs the danger of the collective becoming so compromised by conflicting

internal interests that it is ineffective. (Of course, collectives can become despotic too—military juntas are collectives.)

When decisions are made by collectives, diversity of the members' interests and the interests of those they represent usually becomes very apparent. Indeed, membership in decision collectives often is designed to reflect this diversity—members are selected to represent competing viewpoints or competing groups. Because of this, the time spent in collective decision making is seldom spent on the decision itself; it's spent on the negotiations that lead up to the decision. That is, in order for the collective to begin decision deliberation, it must work through the differences among its members' narratives, plans, expectations, values, and preferences. We'll discuss these negotiations in a moment.

Decision Simplification

Both single agents and collectives resort to the same simplification strategy that individuals use when making private decisions. And, the result is the same as it is for private decisions; complicated level 2 and level 3 decisions are simplified so decision makers can get a grip on them.

Even though it is similar, simplification is more dangerous for organizations than it is for individuals because organizational decisions impact so many people, both internal and external to the organization. What is more, in contrast to individuals' flexibility of plan implementation and revision, organizations have difficulty making timely changes because so many different functions and so many different people are involved. Once set in motion, organization's plan implementations often seem to take on a life of their own. Even when everyone is aware that a plan isn't working, it is very difficult to make rapid corrections. The ungainliness of organizations has been likened to trying to turn a large ocean liner in mid course—it takes time and a lot of patience, and heaven knows what will happen before it's done.

Whatever the dangers, however, there seldom is an alternative to simplification. Non-routine organizational decisions tend to be complex level 3 decisions involving a large number of interdependent features and long time horizons. Most decision agents simply lack the cognitive ability to deal with such decisions, so they must simplify if they are to make any progress at all. Therefore, the case for using decision aids for organizational decisions is precisely the same as it is for private decisions, but more so.

In both private and organizational decisions, aids help decision makers avoid oversimplification by laying the decision out in an orderly manner so all the complexities don't have to be kept in mind at once. Nobody claims that the decision aids we discussed in the previous chapter ensure optimal decision making.²⁵ Rather, the claim is merely that they are helpful within the context of narrative-based decision making because they educate intuition.

Qualms

Of course, whether for organizational or private decisions, decision aids require the user(s) to acquire skill in their use. But it isn't a high level of skill; almost anyone can use a decision aid if they put their mind to it, or hire someone to help them. One might think, therefore, that decision makers would clamor for instruction in how to use aids and would use them for every decision of any size. But, one would be wrong.

Experience shows that even those decision agents who know about decision aids are not always open to using them. Aside from the work involved, qualms about their use results from the transparency they bring to the decision process. That is, decision aids require identification of the

²⁵ Proof of optimality would require a far more formal, rigorous theory than what is being presented here. On the other hand, the problem with formal theory is that proof of optimality usually requires imposition of highly constraining assumptions. If those assumptions aren't met by a real-world decision, there is no assurance of optimality, and the real world seldom is as well-behaved as formal theories require.

defining features of the forecast and the desired future and explicit evaluations of importance and magnitude, all of which reveals the narrative, plans, and values underlying the decision. Common sense suggests that this is a good thing, but many would disagree.

Virtually every organizational decision of any importance will meet with opposition. If the underpinnings of the decision are clear, those who oppose it will know precisely where to aim their arguments. After all, merely opposing a decision won't necessarily get its opponents very far, but if they can produce doubts about the legitimacy of its underlying narratives, plans, and values, they may produce enough opposition to scuttle the decision. Largely because they dislike conflict, decision agents prefer obscurity to transparency, which means that many of them resist using decision aids. Think, for example, of a school board making collective decisions. Virtually every decision gets challenged by one powerful interest or another. Reducing transparency at least makes it harder for critics to pinpoint their attacks.

Of course, from the opposition's viewpoint, transparency is desirable because they think the decision agent is biased against their own interests, which is to say they think the agent has an incorrect narrative, unworthy plans, and questionable values. Opposition isn't necessarily obstructionist for the sake of obstruction—opponents of the decision usually are sincere and often are right. Opposition can provide a corrective, but it also can produce stalemate—and it almost always involves conflict, which most of us abhor.

If nothing else, opposition prevents the agent, single or collective, from becoming tyrannical. But, that said, it seems to be the natural order of things that decision agents tend to favor reduced transparency, even if they claim otherwise, if only so they can make the decision and get on with its implementation without a lot of argument and trouble. And,

opponents tend to favor increased transparency so they can be sure that decisions reflect what they regard as the “proper” narratives, plans, and values—which is to say, so they can be sure that the decision is congruent with what they perceive to be the organization’s interests (and perhaps their own).

Sometimes, however, decision makers conclude that transparency is a good thing. For example, a few years ago a decision collective within a governmental agency (which must remain nameless) was faced with a decision that they were sure would arouse the ire of an especially vigilant environmental group. Rather than trying to finesse things by keeping their deliberations secret and then taking on the environmental group in an all out fight, they elected to do just the opposite. They got together with representatives of the environmental group to figure out the defining features of a mutually desirable future. Of course, there was a good deal of back and forth about what should be included, but that was the point of meeting.

Eventually they worked things out so they had a list that contained both sides’ major concerns. In the course of compiling the list, something unexpected happened. They discovered that they were in agreement about many more things than either side had realized. This meant that the two groups only had to negotiate features about which they disagreed. After they had hammered out a mutual agreement about the features and their relative importance, the environmental group was perfectly willing to let the agency get on with the decision—although they reviewed the chosen plan and made suggestions about improving it before it was implemented, and they kept an eye on the implementation.

Sometimes attempts at openness backfire and the two sides discover they have irreconcilable differences. I once worked with a very dysfunctional church that invited a colleague and me to run a meeting at

which the congregation attempted to identify the features of a desirable future. This was supposed to be followed by a meeting to discuss how to attain that future. But the second meeting never happened because the congregation couldn't agree on what they wanted the future to look like. One faction wanted the church to be community oriented, doing outreach to the poor and in other ways helping people who needed help. The other faction wanted a more traditional church, with emphasis on bible studies and spreading the faith. If nothing else, the meeting made it clear that neither faction was interested in finding a middle ground. A few months later the more traditional members of the congregation departed to form their own church.

The Process

Events leading up to organizational decisions parallel the events that lead up to private decisions. Difficulties are detected when one or more persons in the organization detect discrepancies between how things are now, or are likely to be in the future (extrapolated future), and how they should be (desired future). Alarmed by the discrepancies, they will start to articulate it as a problem that must be addressed: "Sales cannot continue the way they are if we are to make our quarterly goals;" "Forecasts of our pension obligations are so much larger than anticipated that we must find a way to deal with them;" "Workforce training is too slow to provide the skilled workers we will need to remain competitive;" "Membership in the Mothers' Guild has plummeted and we must reconsider its role in the church's overall program;" "Nobody remembered to book a room for the Stamp Club's meeting and we have to come up with something fast."

Single Agent

When they're part of your job, many routine decisions are almost reflexive; you've been in the situation before and your retrieved episodic

narratives instantaneously remind you of what you did, so you can do it again. Even allowing for differences in detail between then and now, what you did before provides a general plan for what you might do now. Multiple experiences with situations of this kind provide you with procedural memories; general rules for dealing with this sort of problem. These decisions are level 1 or very elementary level 2 decisions, and they make up the vast majority of day-to-day workplace decisions.

Non-routine, complex level 2 and level 3 decisions are a different matter. Experience may provide guidelines about what to do, but because the decision is non-routine and has far-reaching effects, you are unlikely to have a pat solution. You've got to think carefully about what is wrong and what to do about it. And, just as it would be for a private decision, your strategy will be to simplify the decision as much as you comfortably can so you can get your head around it. I won't belabor the steps you must take to use the expanded version of the discrepancy test, the decision aid illustrated in Figure 6.3. The process is very like what was described in Chapter 6, except that your assessments of feature importance and expectation must be made from the organization's viewpoint, not your own.

Collective Agent

Collective use of the decision aid is generally the same as single agent use or individual use, except that it requires input from the group. It is the group aspect that complicates things because negotiations are taking place the entire time.

Note that when Herb and Fran used the decision sheet in Figure 6.3 to decide whether to have another child, they were acting as a collective agent making an organizational decision. That is, they are an organization, a marriage, and they made the decision by pooling their thoughts while filling in the decision sheets together. When describing the

example, I brushed over the negotiation part of their decision making, but you can be sure that there was a lot of it going on. After all, Herb knew that if they had a child he'd have to travel less and forego a promotion; Fran knew she'd probably end up giving up or cutting down on the job she loved in order to take care of the child. Both knew that the child would have repercussions, both good and bad, for their daughter Angela. They both knew that deciding not to have another child was going to meet with opposition from Fran's mother. Naturally, all of this would have generated discussions over dinner, in the car on the way home from a movie, and at night after the lights were off. A great deal of the discussion would have consisted of gently testing the waters as expectations were explored; neither person wanting to impose their interests on the other but both wanting to make sure their interests were taken into consideration.

A common understanding must have emerged from all the talking, if only because Herb and Fran were able to fill out the decision sheets without fighting. No doubt compromises were made, some dreams modified, some narratives revised. If not, the discussion would still be going on and the decision aid would be stowed in the back of a drawer somewhere. Of course, in Herb and Fran's case, failure to reach a decision is, in a sense, a decision not to have another child—to stick with the status quo while apparently leaving the question open about another child. This might be a perfectly good solution, but it seldom is satisfying because it leaves things hanging. Moreover, what might work in a marriage, won't work in a corporation—decision by indecision isn't generally regarded as desirable in the business world, although it isn't uncommon.

Pre-decisional negotiations consist of each member of the collective trying to get the others to understand (perhaps adopt) their

current narrative for the task at hand, as well as getting the others to respect (perhaps adopt) their plans and values.

If you seriously want these negotiations to attain a common ground so the decision can go forward, you have to discern which of your own and your colleagues' interests are negotiable and which aren't. It has been my observation that most collectives' final negotiated narratives, plans, expectations, values, and preferences—the ones that guide the collectives' decisions—are the non-negotiables brought to the table by the various members of the collective, particularly those of the more powerful members. That is, if George is powerful and strongly insists on X, then X will survive, and if Sally is powerful and feels strongly about Y, then Y will survive. (Deadlock occurs when powerful people have strong, but opposing, views and neither will compromise—often because the constituencies they represent won't let them). The result of all of this wielding of power is that the underpinnings of the collective's decision often are a strange, frequently inconsistent, muddle; more reflective of the ability of the powerful and eloquent to exert their influence than of any real unity among the collective's members. Decisions based on this muddle seldom make anybody particularly happy.

On the other hand, if the members of the collective can forego power plays and seek legitimate agreement about the organization's interests, the resulting decision usually is reasonable and it is well-received by the other members of the organization. It may not be the very best decision, the optimal decision, but it ought to get things started in the right direction and elicit cooperation throughout the organization. Getting started in the right direction is half the battle; the plan can be revised to fit the evolving circumstances. Even factoring in organizational inertia, the plan will be easier to change in mid-course if it has been widely accepted

because the organization's members will tend to give you the benefit of the doubt.

An Example

Steller Art Frames, a fictitious company founded in the 1920's, is located in a Southwestern American city near the Mexican border. It began when George Steller, a cabinet maker, began to produce custom picture frames using turn-of-the-century motifs from the arts and crafts movement in domestic architecture and interior design. The frames were particularly effective for western art in conjunction with mission-style furniture. As a result of the quality and design of its frames, Steller prospered from the moment it was founded. After George passed away, other Stellers stepped forward to take over, but management skills didn't seem to be inherited so outsiders were eventually brought in to run things. Some of these people were less able than others and the board of directors, all descendents of George Steller, has run through a distressingly large number of CEOs over the years.

Skip forward to the present: Steller has fallen on hard times and the board of directors has hired yet another new CEO; a woman who goes by the single name, Carson. The board has made it clear that it values the participation of both itself and the rank and file employees in big decisions, which is fine with Carson because this plays to her strengths and she has made up her mind she is going to succeed where other CEOs have failed.

As she has worked with the board and with employees to find a way out of Steller's difficulties, it has become abundantly clear that continuing to do what they've been doing isn't an option—the extrapolated forecast is unacceptable and all efforts to repair it have failed. It is time to come up with a new, inspiring, vision for the future and Carson is the person to do it.

Carson listened as the Chief Financial Officer described the problems facing the company. She had known most of what he was telling her before she agreed to take the job; everyone had been very upfront about what was happening. But now that she was getting the details, she realized that things were more serious than she had thought.

That night, as she ate Chinese take-out in her new apartment, she thought about what she should do. In fact, she was thinking of a way of coming up with a plan for saving the company. But she didn't just want to save the company, she wanted to transform it so it wouldn't get in this same fix again.

She had no doubt that the extrapolated future the CFO had described was unacceptable. But, other than survival, she didn't know exactly what the desirable alternative to it might be; she simply hadn't been part of the company long enough to fully understand what was possible and what wasn't. Therefore, her first job was to tap the knowledge and intuitions of those who had been around longer in order to figure out what the desirable vision of future might be like. Then, her second job was to work with those same people to design a plan for attaining that desirable future.

Getting Organized.

Carson began by holding a meeting of all 109 employees. She and the CFO described the future that could be expected if things continued as they were (the extrapolated forecast). Then, Carson asked the group to elect three of their colleagues to work with a member of the Board of Directors (a grandson of the company's founder), the CFO, and Carson herself to devise a vision for the company—a description of the future that would be both desirable and reasonable.²⁶ She was somewhat dismayed when two of the people who were elected turned out to be chronic

²⁶ In this example, the vision is an envisioned future (Box 2c).

critics of management and apparently suspicious that the “crisis” was merely a pretext to justify laying off employees

The six members of the newly named Vision Group met the next day. Carson knew that she had to overcome the suspicions of the two dissenters before anything could be accomplished. She began by granting them, and anyone else who wanted it, access to any financial documents (other than specific salaries) they needed to appraise for themselves the problems facing Steller. They were given three entire days to probe as they wished, and every manager was instructed to be available to them and to answer their questions candidly.

Her strategy worked; three days later the two worrisome employees attended the second Vision Group meeting with decidedly different attitudes. They seemed a bit intimidated by the size of the problems facing the company, but they soon began to suggest ways of solving them. Carson listened for a while and then interrupted. She thanked them for their suggestions and asked them to remember their ideas so they could be brought up later, when plans were being formulated. For the moment, however, the group’s task was to figure out what Steller should look like when whatever they decided to do had been done. In short, they needed to formulate a vision for the company.

Idea Generation

Carson set aside two days for idea generation, with everyone dreaming, imagining, wishing, and sometimes just fantasizing about what Steller might become. Carson was careful to prevent premature closure—working to keep the discussion rolling rather than letting one set of attractive ideas dominate to the point of becoming a premature consensus. She wanted lots of ideas before they began to narrow things down.

To facilitate idea generation, Carson suggested that one of the group members, a carpenter named Mark, be the Recorder. The walls of the conference room were covered with huge sheets of white paper and the Mark's job was to use a black felt-tipped pen to write down ideas (or shorthand names for ideas) as they came up in the discussion. By noon of the first day, he had nearly covered the walls; there had been lots of ideas at first but they began to peter out by the afternoon. Sensing that not much more was to be gained by continuing, the group adjourned for the night, agreeing to meet again the next morning.

All in all, things had gone far smoother than Carson had anticipated. There had only been one flare-up and a little adroit negotiation had solved it before it became disruptive. The third of the three elected members of the group was a 70 year old woman, Edith Sorenson, who had worked at Steller since she was 18 and who had immense moral authority among her co-workers—everybody liked her and nobody crossed her. She clung to the idea that a major feature of Steller's future should be a strong "presence" in the community. She argued that Steller had always sponsored local sports teams and contributed money to worthy projects and it should continue doing so. The Board's representative to the group, Ed Steller, disagreed. He thought that in these difficult times, the company should focus on its own survival; in fact, he was in favor of moving all production to Mexico, leaving only a small business office in the United States. Edith and Ed had both become a bit red in the face by the time one of the other members stepped in. This third person pointed out that it probably was possible to reduce Steller's overhead without withdrawing from community responsibilities. In fact, she said, citing the most recent annual report, the money Steller spent on community relations was relatively small in the total scheme of things. Besides, she continued, if it were decided that Steller should eventually

expand, community good will would pay off when it came time to search for talented people to hire; nobody wants to work for a company that is seen as failing so badly that it must turn its back on the community of which it has always been a part. That settled it; Edith got to keep community “presence.” (In fact, Ed had quickly realized that without pulling rank as George’s grandson, he wasn’t going to win against Edith, who was powerful in her own right and who regarded community “presence” as non-negotiable. But, in the process of letting her have her way, he made sure he got two things; he got to emphasize the importance of reduced expenses as a key feature of the desired future and he got to introduce the idea of moving to Mexico—for which he argued strongly later on.)

Midway into the second morning it became clear that there weren’t going to be any new ideas—people were trying, but the results were increasingly farfetched. Sensing that the well had run dry, Carson made a final call for ideas. When nothing new was offered, she said it was time to consolidate everything that was written on the walls in the hope that some sort of order would emerge.

Consolidation

Mark the Recorder began by using various colors of pens to tie together ideas that were conceptually linked. He explained aloud what he was doing and why, and others spoke up when they disagreed or thought that an idea had been overlooked or improperly linked. After a while, all the ideas had been joined by different colored lines to indicate conceptual linkages, although some ideas had been judged to be so similar that they could be combined. This took all afternoon and by 5:00 everyone was tired. Carson asked them to return at noon the next day, at which time she would provide them with a tentative list based on the work they had done thus far.

Early the next morning, Carson sat in the conference room drinking a cup of coffee and looking over everything that was written on the walls. After a while she began to see how the ideas that were linked with a particular color of ink formed a general concept. Soon, all of the constellations of ideas began to make sense to her—although there some ideas that didn't fit well with any of the others.

She spread out another big sheet of paper on the conference table and began to tinker with the concepts, arranging them into lists like the ones in the decision sheets we've reviewed in the previous chapter. Things began to fall into place when she took out all the ill-fitting ideas and divided them into two new concepts. Then, she drew up a tentative list to present to the group at noon.

Carson began the noon meeting by describing what she had done and asking for discussion. The odd thing about lists is that they have a compellingness of their own; they seem more final than they necessarily are. So Carson's problem was that the group was less willing to revise her list than she thought they should be. She had to let long silences build before anyone would venture a suggestion. As it turned out, most of these reluctant suggestions were valuable. For example, the two new concepts that Carson had constructed from the ideas that didn't fit in the other concepts really weren't very convincing. The group discussed them and ended up putting one of the ideas into another concept and creating a separate new concept for the other—a concept that, now that they had it, was obviously what was needed.

The final list of features of the desired future, Steller's vision, is listed on the left side of the decision sheet in Figure 7.1.

Planning

Carson asked everyone to change mind-set. Instead of thinking about what constituted a desirable future, they now should think about

Figure 7.1: The decision sheet for Steller Art Frames.

Features of Desired Future	Option _____		
	Importance (i_c)	Magnitude (m_c)	Impt x Mag ($i_c m_c$)
<u>I. The Company</u> () []			
<u>A. Survival</u> () []			
1. Increased Income ()	1 _____	1 _____	1 _____
2. Decreased Overhead ()	2 _____	2 _____	2 _____
3. Diverse Product Lines ()	3 _____	3 _____	3 _____
<u>B. "Family"</u> () []			
4. Keep Culture ()	4 _____	4 _____	4 _____
5. Retain Employees ()	5 _____	5 _____	5 _____
6. Hire People Who "Fit" ()	6 _____	6 _____	6 _____
<u>II. The Community</u> () []			
<u>D. Local Community</u> () []			
7. Employment Opportunities ()	7 _____	7 _____	7 _____
8. Economic Development ()	8 _____	8 _____	8 _____
9. "Presence" in Community ()	9 _____	9 _____	9 _____
<u>E. Broader Community</u> () []			
10. International Reputation ()	10 _____	10 _____	10 _____
11. Maintain Ties with Mexico ()	11 _____	11 _____	11 _____
12. Innovate ()	12 _____	12 _____	12 _____
			$D' = \sum i_c m_c / x =$ _____

how to attain Steller's new vision. Accordingly, she changed the group's name to the Planning Group.

Once again the walls were swathed in paper and the felt-tipped pens were laid out. Mark the carpenter had been such a good Recorder that he was pressed to serve again as the group members described ideas about how to attain the desired future. Many of the ideas were fragments that might move things along, but weren't themselves plans for doing the whole job—they were tactics and short sequences of tactics. Everything was written on the walls until new ideas stopped coming and old ones started recycling. Then Carson stopped the discussion and asked Mark to draw colored lines among related ideas as he had before. But, he quickly concluded that what had worked for vision construction didn't work for plan construction.

The problem was that merely linking related ideas didn't create coherent plans. In fact, plans required sequences of different tactics, not similar ones, in order to create a forward thrust. On the other hand, the hodgepodge of ideas on the walls suggested identifiable themes, and these themes could then be fleshed out, using various of the ideas, to create something approximating complete plans. After a little work, the Recorder came up with seven themes. They weren't wholly independent of each other because many of them contained ideas (tactics) that also were part of other themes; which caused some discomfort until the group realized that the themes were the plots of plausible stories (narratives) about how to move Steller from where it was now to where they wanted it to be. Stories often have similar elements without being the same, so plans could too.

The seven plans they came up with were:

- Plan 1: Wait and see--do nothing differently and see if things get better.

- Plan 2: Move the company to Mexico to reduce costs and be closer to suppliers.
- Plan 3: Sell the company to a Midwestern competitor who has made overtures in the past.
- Plan 4: Open a plant in Mexico while moving the parent company to less expensive quarters in the local economic development zone.
- Plan 5: Create new product lines.
- Plan 6: Expand markets, primarily in Europe and Canada.
- Plan 7: A combination of plans 4, 5, and 6.

Evaluation

Carson thought it was time to get back to the employees about the group's progress. She held another meeting at which she asked one of the group's elected members (a shipping clerk) to describe the vision the group had come up with, followed by another elected member (a frame finisher) describing the seven plans. She then explained that the next step was to evaluate the expected effectiveness of each plan and to select the most effective one for implementation. There were numerous questions about the vision and the details of the plans, some of which the group could answer and some of which they couldn't. The ones they couldn't answer were written down with an assurance that the answer (once they figured out what it was) would be sent to everyone via e-mail. The questions were valuable; one led to the realization that there was a poorly thought out portion of one of the plans that had to be investigated—whether the Midwestern competitor would offer the price it had offered in the past, now that Steller's problems were public knowledge.

Carson was worried that the six group members had got so used to working together that they might not be as critical as they should be when they moved on to the evaluation and decision phases of their work.

Therefore, at the meeting she asked employees to elect 4 more members for the group, bringing membership to ten. She also reported that the Board had asked to have an additional representative (the Directors were feeling a little left out of the loop, which Carson noted would require some delicate diplomatic work sometime soon).

The eleven members of the Planning Group met the following Monday morning. The new members were brought up to speed and their questions were answered. The glitch in Plan 3, was addressed (the offer had indeed decreased—a lot).

Then a decision sheet (Figure 7.1) was given to each member of the group and the first task was explained:

- Evaluate the importance of the groups, sets, and features of Steller's desired future.
 - Assign each group (I and II in Figure 7.1) a number from 1 to 3 to indicate how important the group as a whole is to Steller's desired future (1 = low importance, 2 = medium importance, and 3 = high importance).
 - Assign each set (A,B, in Figure 7.1) in each group a number from 1 to 3 to indicate how important the set as a whole is to Steller's desired future.
 - Assign each feature in each set (1 through 12 in Figure 7.1) a number from 1 to 3 to indicate how important each feature is to Steller's desired future.

The eleven group member's decision sheets were collected and their assessments were averaged across the group members to derive a single set of numbers (the arithmetic mean was used, but other measures of central tendency would work too). These averages were written on yet more paper on the wall, along with an indication of the degree of consensus for each of them (the range was used, but other measures of

dispersion would work too). Members challenged some of the averages and when everyone agreed, some were adjusted a little (but not much). Where consensus was low, the evaluation was discussed and reassessed in hopes of increasing it (one of the features never received much consensus but the mean was right in the middle of the distribution so everyone figured that was good enough).

After everyone was reasonably satisfied with the importance averages, it was time to calculate the importance of each feature of the desired future relative to the other features, i.e., their relative priority.

- The Recorder multiplied the average evaluation given to group I times the average ratings for each set in the group (A, B).
- Then he multiplied the average evaluation for each feature by the product calculated in the previous step for the set (A or B) of which the feature was a member and wrote the answer in the blank under Importance on the right of the decision sheet (ic).
- He repeated this procedure for all twelve features.

The Recorder's decision sheet was then duplicated and seven copies were given to each member of the group; one sheet for each of the seven plans.

- For each plan, they were instructed to assign each feature (1 through 12 in Figure 7.1) a number from 0 to 3 (0 = no discrepancy, 1 = small discrepancy, 2 = medium discrepancy, 3 = large discrepancy) to indicate the degree to which they felt that each feature of the future offered by the plan under consideration was discrepant from the feature listed on the decision sheet. Zero indicated that the plan's future was not felt to be discrepant from the desired future on that feature of

the future. They were to write their answer in the column on the plan's sheet labeled magnitude (m_c).

After they had finished evaluating discrepancy magnitudes for each plan, each group member performed some of the preliminary calculations for each of his or her seven decision sheets:

- For each feature, each member of the group multiplied the number in the Importance column (i_c) of his or her decision sheet and the corresponding decimal number in the Magnitude column (m_c), and wrote the answer in the column on the far right of the decision sheet ($i_c m_c$).

Mark the Recorder drew seven 12 (features) by 11 (group members) grids on the wall, one grid for each of the seven plans. Each member of the group came forward and wrote the 12 numbers ($i_c m_c$) that were on his or her decision sheet in one of the 11 columns of each plan's grid. After all the numbers had been written in the appropriate columns of the appropriate grids, the Recorder averaged (the arithmetic mean again) across each row in each grid. He wrote the average for each row on the margin of its grid and then summed the averages, and divided by the scalar for a 3-tier hierarchy ($x = 3^4 = 81$, for the 3-point scales Carson used), which yielded the group's D' for each plan. As it turned out, Plan 7 had a considerably smaller D' than any of the other plans, which by this point surprised nobody.

Decision

It was moved, seconded, and passed that Plan 7 be adopted for pursuit of Steller's vision. The group had reached its decision. Next came the job of convincing others that it was the right decision and enlisting their cooperation in turning it into action.

Follow-up

The next morning, at a special meeting, Carson and the Board's two representatives presented the vision and the plan to the Board of Directors and explained how the plan would achieve the vision. After a barrage of insightful questions, both the vision and the plan were accepted—in part because they looked good and in part because it was clear that Carson liked them. Most Board members had come to the conclusion that she was the best thing that had happened to the company in a long time and that it was their job to support her.

An hour later, Carson held a meeting of the employees and made the same presentation she had made to the Board. Again there were questions, although not as many as the Board had had, followed by a vote to accept both the vision and the plan. Two hours later, Carson and her husband were on a plane for Jamaica and a long weekend in the sun.

The following Tuesday, Carson began working on the details of the plan so she could get everyone started on its implementation.

Progress Report

In this chapter we discussed single agent and collective agent decision making in organizations. We identified the organization's culture as its value system and discussed how it can contribute to resistance to decisions and their implementation. In addition, we talked about the effects of differential power on decisions as well as negotiations that lead up to decisions. We ended with an elaborated example of organizational decision making using the multi-leveled decision aid that was developed in Chapter 6 (Figures 6.3 and 7.1).

Sources and Further Reading

- Beach, L. R. (1993). *Making the right decision: Organizational culture, vision, and planning*. Englewood Cliffs, NJ: Prentice-Hall.
- Beach, L. R. (Ed.) (1996). *Decision making in the workplace: A unified perspective*. Mahwah, NJ: Lawrence Erlbaum Associates.
- Beach, L. R., & Connolly, T. (2005). *The psychology of decision making: People in organizations (2nd ed.)*. Thousand Oaks, CA: Sage.
- Beach, L. R. (2006). *Leadership and the art of change: A practical guide to organizational transformation*. Thousand Oaks, CA: Sage.

Chapter 8: Commentary

The purpose of this book is to address the second and third of the three questions about human subjective experience and thinking (the “fundamental puzzle”) that were posed at the beginning of Chapter 1. The first of the three questions, about sensory functioning and basic perception, is beyond the scope of this book. The second question, about the meaningful flow of subjective experience from the past to the present and into the future, was addressed in our discussions of narratives. The third question, about how expectations about the future give rise to actions, was addressed by our discussions of decision making.

Our endeavor to cast light on these two questions has led us to a theory of thought and decision making that we will call the Theory of Narrative-Based Decision Making (NBDM). The theory is intended to provide an alternative to the commonly accepted paradigmatic prescriptive theory that derives from economics. The foundational argument is that narratives are the “natural language” of thought and that formal paradigms play a supporting role, having been developed to help us think more rigorously about particular areas for which narratives lack the required precision. This is not to say that paradigms aren’t part of the human cognitive repertory; obviously they are. But, they are useful adjuncts to narrative cognition rather than the basic mode of thought.

As we saw in the introductory chapters, narrative thinking consists of the stories that help us make sense of our experience and paradigmatic thinking consists of knowing about and using formal systems of procedures, often computational procedures, for performing classes of tasks. The usual goal of paradigmatic thinking is to inform and add precision to narrative thinking.

Overview of NBDM

Before we proceed, let's briefly review the elements of NBDM so we can keep the big picture in mind. These elements derive from four propositions and two adjunct concepts, which amount to the basic assumptions of the theory:

- *Cognitive narratives provide meaning and continuity by linking the past with the present and allowing us to make educated guesses about the future (forecasts).* In Chapter 2 we saw that decision makers construct and use coherent stories to make sense of what is happening in light of what has preceded it and to extrapolate the story in order to forecast the future.
- *Based on our past experience and our values, we create rules, some of which are used to make forecasts and some of which are used to evaluate the desirability of the future offered by those forecasts.* In Chapter 3 we saw that the desirability of an extrapolated forecast of the future derives from its congruence with the decision maker's primary and secondary values, including his or her preferences—the desired future.
- *When the forecasted future is undesirable, we construct plans that guide actions designed to promote a more desirable future.* In Chapter 4 we saw that the decision maker anticipates, however sketchily, the sequence of actions (tactics) he or she will perform as he or she endeavors to make the actual future more closely approximate the desired future.
- *All decisions are about discrepancies and what to do about them.* In Chapter 5 we saw that discrepancies are mismatches between the features of entities and that they prompt action to correct them.
 - The basic model for discrepancy analysis is:

- ✓ The decision maker assumes that comparable entities are essentially identical in all their relevant features until or unless there is sufficient evidence to the contrary.
- ✓ The decision maker compares the entities' features and notes the discrepancies that are large enough to be considered nontrivial.
- ✓ If the number of discrepancies exceeds the decision maker's level of tolerance (rejection threshold), he or she rejects the identity assumption and regards the entities as different.
- The prompted actions are:
 - ✓ Discrepancies between the features of an extrapolated forecast's future and the desired future prompt the formation of a plan to make the actual future closely resemble the desired future.
 - ✓ Discrepancies between the results offered by a including a tactic in a plan and the results required by the plan prompt replacement of the tactic with one that offers better results.
 - ✓ Discrepancies between the future offered by a plan's action forecast and the desired future prompt revision or replacement of the plan.
 - ✓ Discrepancies between feedback about what an implemented plan is accomplishing and what is needed to attain the desired future prompt revisions of the plan to keep it on track toward the desired future.

In the latter part of Chapters 5 and continuing through Chapters 6 and 7, we examined two adjunct concepts:

- The first concept was that decisions differ in *magnitude*.

- Level 1 decisions are simple, rapid and unconscious.
- Level 2 decisions are the stuff of everyday life and require conscious deliberation.
- Level 3 decisions are complex and frequently exceed the decision maker's cognitive capacity to deal with the details.
- The second concept was the *simplification strategy*.
 - Lacking the ability and cognitive capacity to manage complexity without the aid of an appropriate paradigm, unaided decision makers simplify complicated level 2 and all level 3 decisions and treat them as though they were lower-magnitude decisions.

In Chapters 6 and 7, we examined paradigms, called decision aids, for helping decision makers deal with the complexity of personal and organizational decisions.

The Narrative about NBDM

The NBDM narrative is: You construct narratives to make sense of what is going on around you and how it relates to what has happened in the past. In addition, you extrapolate the narrative to forecast what will happen in the future. If the forecasted future is desirable you continue doing what you're doing and let the future unfold as it will. If the forecasted future is undesirable, you consider ways in which you might prevent the actual future from turning out that way by intervening to make something more desirable occur happen instead. Deciding whether the forecasted future is desirable and deciding on ways to create an alternative future both require you to have some idea about what you would like the future to be. This desired future is based on your primary and secondary values, including your transient preferences.

A plan to achieve a desirable future is a composite of things you've done in the past to deal with situations that structurally or functionally

resemble the current situation. These past actions are components of episodes stored in your memory. Each episodic memory is part of a network of similar memories, so you usually have a fairly large pool from which to draw both strategies and tactics in your efforts to formulate a plan. The strategic and tactical components of your plan are selected to reflect your values as well as their ability to advance the plan toward its goal of achieving your desired future.

Having designed a plan, you begin to implement it while monitoring its progress toward producing your desired future. Monitoring consists of making periodic (action) forecasts of what will happen if you continue to follow your plan. An action forecast that includes your desired future prompts you to continue the plan's implementation and a forecast that doesn't include your desired future prompts you to revise your plan until it permits an improved action forecast.

Because the environment in which plan implementation takes place is dynamic, being influenced by both the implementation itself and the actions of outside forces, you must be flexible. This means that you must revise your plan in light of current conditions until you can forecast that it will achieve your desired future, but you also must be open to changes in the specifics of the desired future. This flexibility allows you to learn from experience as well as to embrace futures that you might never have dreamed of until their possibility was uncovered by the implementation.

Decisions arise when you make comparisons and all decisions are about discrepancies discovered during the comparison process. The major decision points are (1) when you compare the extrapolated forecast with your desired future and (2) when you compare the action forecast offered by a potential plan with your desired future.

The discrepancy test (discrepancy analysis) is a model of decision making that derives from one of the most basic survival skills possessed by every creature, the ability to classify objects and events in order to recognize or identify them. The test is a formal paradigm that is designed to clarify and add rigor to our consideration of how psychological discrepancy analysis might occur. No doubt the latter is far less rigid than the model, but the model can profitably be thought of as the essence of what happens during decision making.

The basic discrepancy test is presumed to model level 1 decisions; simple, rapid decisions that are largely unconscious. It assumes the decision maker unconsciously makes comparisons between the extrapolated state of affairs and the desired state of affairs, encodes the number of discrepancies, and rejects the identity assumption when the number exceeds a threshold level. Rejection leads to use of habit to remedy the discrepancies which, if it succeeds, ends in quiescence. If it doesn't succeed, consciousness is enlisted to search for or create a remedy, raising the decision to level 2.

Once consciousness is enlisted, the decision ceases to be simple or rapid. The difficulty is that higher level decisions often involve comparisons that are too multifaceted for the decision maker to keep everything in mind at once. Lacking tools to help them with these decisions, unaided decision makers usually resort to simplifying them so that something approximating the basic discrepancy test can be applied.

The NBDM decision aids outlined in Chapters 6 and 7 are paradigms designed to allow decision makers deal with more complexity than they can without an aid. They simplify complex decisions in the same way decision makers usually do, but avoid oversimplification through use of a structured decision sheet and pencil and paper. The result is that more of the decision's complexity can be included than would be the case if the

decision aid was not used, but the complexity is laid out in an understandable manner. Once the decision is captured on a decision sheet, the decision maker can use his or her normal, narrative way of thinking about plans to make evaluations of the decision's components. The real reason for doing all of this is less to let the aid reach a conclusion than to allow the decision maker to grasp the issues involved in the decision and make his or her own, intuitive, narrative-based decision. Decision makers tend to know what they are going to decide long before the analysis is complete.

Research Strategies for NBDM

As yet, NBDM lacks empirical support, but its antecedents do not. There is a literature on narratives dating at the very least to the early 1980's. There is a literature on the psychology of planning going back equally far. The literature on expectations dates to the 1960's. Research on the discrepancy test (see Appendix 1), dates to the 1980's. This being said, however, there has not yet been any empirical research on NBDM per se—because this book is its debut.

Rethinking Research Strategies

Before research begins, thought must be given to the strategies to be used. Just as NBDM breaks with the past, the research that investigates it must break with the past. It is important to test the theory on its own ground rather than on ground that is simply convenient because it is familiar and amenable to conventional methods. In this vein, it is time to abandon the familiar paraphernalia of the decision laboratory: choices among bets, colored balls drawn from bags, intricate word problems, complicated experimental games, and all the other contrived little puzzles we give participants in hopes of inferring something meaningful from the errors we have virtually assured that they will make.

Perhaps we should look upon these established methods as simply an early stage of development in the study of decision making. In the 1960's, when decision research was getting under way, it made a degree of sense to think that the products of the human mind (specifically, expected value theory and probability theory) might mirror the "natural" ways in which humans make decisions—simply more sophisticated versions of everyday thinking. This viewpoint led experimenters (and I was one of them) to examine the parallels between the procedures involved in applying the paradigms to a decision problem and what unaided experimental subjects did when confronted with the same problem. As reported in Chapter 1, surprisingly few parallels were found.

From the beginning, there were critics who questioned the appropriateness of comparing the unaided performance of humans with the procedures and output of formal paradigms. Some suspected that, whatever research subjects were trying to do, it wasn't necessarily what experimenters thought they were trying to do—so the paradigm was somewhat beside the point. Others thought it odd to ask unaided subjects to perform tasks for which paradigms exist; after all, the reason for developing paradigms is because there are tasks that people don't naturally do very well. It isn't any wonder that unaided human decision performance is found wanting when the subjects are denied the tools—both the paradigms and pencil and paper--needed to perform the tasks they are asked to perform.²⁷ In light of these criticisms, one has to ask whether the results of these studies have much value.

The answer is both yes and no. "Yes" because anyone who assumes that unaided decision making is the same as paradigmatic decision

²⁷ I sometimes wonder if the vast bulk of decision research won't come to be seen merely as evidence for the rather obvious fact that people need paradigmatic tools, and the skill to use them, for some cognitive tasks (see Box 8a).

making—which many economists do—should be disabused of the idea. Behavioral Economics, where decision research meets economic theory, has quite properly taken up the cudgel for more realistic economic theory.

“No” because many of the results of decision research probably don’t tell us much about unaided decision making and, in fact, they don’t tell us much about how to create decision paradigms that are useful to the average person. Narrative thinking is so different from the formal paradigms of economics and related disciplines, paradigms that have developed far beyond the intuitive understanding of most of us, that it makes little sense to use these extremely abstract systems as criteria for “right thinking.” It is time for those of us involved in psychological decision research to turn our attention to how decisions are made in the context of narrative thinking, because that’s the thinking we all rely upon for most of our decisions (even economists and decision researchers, when we’re not at work).

Three Instructive Examples

Three lines of research provide instructive examples of alternatives to the usual decision research studies. If the examples appear similar to NBDM, it is not an accident. NBDM is heavily influenced by their results and is an attempt to incorporate those results into a theory of narrative based decision making.

First example. Nancy Pennington and Reid Hastie have examined the creation of narratives by studying the way jurors use evidence that is presented to them in the course of a trial.²⁸ Volunteers from the jury pool for the Massachusetts Superior Court, who had served on an average of three juries prior to the study, were shown a videotape of a trial and were asked to act as a jury member. The tape contained the usual information

²⁸ Pennington & Hastie (1986, 1988, 1992).

that is forthcoming in a trial—including the inaccuracies, omissions, and disorderliness inherent in differing testimonies.

Afterwards, each individual juror talked to the experimenters about the case and his or her verdict. This was tape-recorded for later analysis. Results showed that the jurors constructed coherent stories about the crime using information from the testimonies, in combination with their general knowledge—including what they thought they or others would do in similar circumstances. These stories were organized and economical; irrelevant information was ignored and missing information was filled in through inferences based upon the parts of the story that had already been constructed. Finally, the verdict that most closely “fit” the story was the one that was chosen. Even though all of the stories shared a general structure, different jurors constructed different versions and their verdicts varied accordingly.

Second example. Gary Klein and his associates have studied how decision makers use their prior experience, episodic memory, to make decisions.²⁹ To do this, he and his colleagues interview real-life decision makers such as firefighters, military commanders, and police officers about non-routine decisions they have faced. The purpose is to identify what options existed, which of these options were actively considered, what information was important, what the goal (desired future) of the decision was and how that goal shifted in the course of the decision, and the way in which the decision was made. Interviews are recorded and impartial analysts identify the critical decision points and how the decision maker went about making the required decision.

Examination of these decisions showed that they tended to be unanalytical and heavily based on recognizing the situation as one for which this or that aspect of past experience or training provided a

²⁹ Klein (1989, 1993, 1996)

solution. Depending on the expertise of the decision maker, and on the degree of latitude he or she had in approaching the decision, from 40%-80% of the decisions involved recognition in at least the initial stages. The more experience the decision maker had in the area in which the decision arose, the greater the role of recognition—which is to say, in our terms, the greater the role of episodic memory.

The process begins with a decision about the familiarity of the situation—has it been encountered before or has it been covered in training? If it is not familiar, more information must be obtained. If it is familiar, a forecast is made about what will happen. If what is going to happen isn't good, the question becomes one of what to do about it. Potential actions are not usually reflexive (although this can happen under extreme time pressure). Rather, the decision maker performs a mental simulation that permits him or her to imagine what might happen if the actions are performed. If the simulation suggests that the actions will successfully deal with the situation, they are implemented. If the simulation reveals potential problems, the actions are modified and the simulation is run again. This is repeated until a set of actions is derived that produces a successful mental simulation. Of course, if the situation is changing over time, this process must be constantly repeated in order to keep the actions attuned to the changing conditions. Thus, Klein's model, called the Recognition Primed Decision Model, has four main components: recognition of the situation, understanding of the situation, serial evaluation of the potential courses of actions for solving the problem, mental simulation of the possible results of using an action in order to evaluate its potential, and choice and implementation of a course of action. All of this must sound familiar to you by this point.

Third example. As we saw in Chapter 7, an organization's culture embodies its values and the more discrepant an action forecast is from

those values, the less the acceptability of its underlying plan. My colleagues and I tested this using three business firms.³⁰ One firm manufactured athletic shoes and was experiencing large executive turnover at the time the study was done. The second firm, a winery, was growing rapidly but was managing change carefully. The third firm, which manufactured ski wear, was old, stable, and most of its employees had worked there for a long time. The firms were picked because they differed in how much their employees were likely to understand and agree about the firms' culture (values): athletic shoes least, because of its turnover; wine next, because of its rapid change; ski wear most, because of its stability.

To ascertain each firm's values, some of its executives were interviewed about the firm's culture—which turned out to be very different for the different firms. Then, for each firm, various plans were drawn up for the introduction of a hypothetical new product. Each plan was designed to have some features that were discrepant from that firm's values and some that were not; different plans had different numbers of discrepant features. The plans for each firm were presented to all of its executives with the request that they rate each plan for its compatibility with the firm—from extremely incompatible to extremely compatible.

As expected, for all three firms, the greater the number of discrepancies between a plan's features and a firm's values, the lower the rated compatibility. But of equal interest was the finding that the agreement among the executives differed from firm to firm. As predicted, it was lowest for the athletic shoe manufacturer with the high executive turnover, medium for the winery which was growing and going through a lot of controlled change, and highest for the ski wear manufacturer which was an old company with little turnover. In short, the more stable the

³⁰ Beach, Smith, Lundell, & Mitchell (1988).

organization, the clearer its values to its executives and the more agreement among their decisions about potential plans.

The results suggest two conclusions: First, the greater a plan's discrepancy from an organization's values, the less compatible with the organization it is perceived to be. Second, the longer the tenure of the organization's members and the greater the organization's stability, the greater agreement there is about the merits of proposed plans of action. The former speaks to the role of discrepancies in evaluation of potential plans. The latter speaks to the problems encountered by single agent decision makers who base decisions on their understanding of the organization's values only to find that the decision conflicts with others' understanding of its values.

Comparison of NBDM with Similar Theories

In a previous book, *The Psychology of Decision Making*, my co-author, Terry Connolly, and I reviewed a number of what might be called mini-theories that have been advanced as alternatives to various aspects of Economic Decision Theory.³¹ Our argument was that these mini-theories would have to be incorporated into any general psychological theory if it were to be complete. In the construction of NBDM, the aim has been to do just that.

Image Theory

Image Theory, received the most attention in the Beach and Connolly review of mini-theories, and NBDM is a fairly straightforward elaboration of it.³² The basic idea of Image Theory is that a decision maker's store of knowledge can be partitioned into three categories,

³¹ Many of these mini-theories are regarded as part of a movement within psychological decision research that is called Naturalistic Decision Making; the emphasis being on how decisions actually are made rather than how they "should" be made—as defined by Economic Decision Theory. NBDM and its predecessor, Image Theory, are part of the naturalistic movement.

³² Beach & Mitchell (1990), Beach (1990).

called images because they are his or her vision of what constitutes a valuable and properly ordered course of events. The categories are labeled the value image, the trajectory image, and the strategic image. The constituents of these images can be further partitioned into those that are relevant to the decision at hand and those that are not. The relevant constituents define the decision's frame, which gives meaning to the context in which the decision is embedded and which provides standards that constrain how it will be understood.

Image Theory's value image consists of the decision maker's morals, ethics, and preferences; it sets standards for how things should be and how people ought to behave. This concept has been retained in NBDM as normative expectations.

Image Theory's trajectory image is the decision maker's agenda of goals. This concept has been retained in NBDM as the desired future.

Image Theory's strategic image is the decision maker's plan for achieving his or her goals. This concept has been retained in NBDM as the plan and its action forecast.

Image Theory posits two kinds of decisions: adoption decisions and progress decisions. The former is about adding new goals or plans to the trajectory and strategic images; the latter is about evaluating the progress of plan implementation toward goals. NBDM collapses these two kinds of decisions into one general kind about whether there is a significant discrepancy between the forecasted future and the desired future. Then, if there is a significant discrepancy, a second decision must be made about what must be done to correct it. This applies both to instigation of a course of corrective action because the forecasted future is undesirable as well as to changes in that course of action when it shows signs of failing to be corrective.

In Image Theory, new goals for the trajectory image or new plans for the strategic image are screened by the compatibility test in light of the decision maker's values. The compatibility test has been retained in NBDM and renamed the discrepancy test. However, in NBDM, it is used to evaluate the mismatch between forecasts of the future, or forecasts of the effects of a plan of action, and the desired future—which is defined by the decision maker's values.

In Image Theory, when more than one candidate goal, or candidate course of action, survives the compatibility test, the decision maker breaks the tie using paradigmatic choice strategies, such as Economic Man's maximization of expected value. In contrast, NBDM provides choice paradigms based on the discrepancy test rather than on logic that is external to the theory.

Research on Image Theory has focused almost exclusively on the compatibility test (now called the discrepancy test). The results are relevant to NBDM and are summarized in Appendix 1 at the end of this book.

There was never much research on images because they were generally misunderstood. My colleague, Terrance Mitchell, and I knew what they were, but nobody else did. Image Theory's images were always being confused with visual images, although the concept was very much broader. Work on narratives was just getting underway at the time Image Theory was proposed. Had it begun earlier, we probably would have used narratives rather than images as our key concept, if only because everyone knows what stories are and most of us understand how much implicit information is packed into even the simplest of them. The familiarity and appreciation of narratives' richness makes NBDM more approachable than Image Theory, and changing the names (and, it must be admitted, the characterizations too) of the value, trajectory, and

strategy images to normative expectations and to forecasts is even more helpful. Image Theory always seemed to me to be almost embarrassingly simple and intuitive, but I was in the minority. Its awkward terminology was part of what made it difficult for others to see its virtues.

My hope is that by couching the general ideas of Image Theory in the framework of narratives, it will make them more easily understood. The added advantage is that NBDM is simpler than Image Theory; we don't need a concept of framing and decisions aren't about screening on the one hand and choice on the other, although screening and choice both are encompassed by the NBDM. Moreover, making discrepancy from the desired future the foundation of decision making places values in the forefront, which is where they belong. Decision making is about managing your future in light of your values. It is as simple as that.

Recognition Theories

Recognition theories assume that the context in which a decision is to be made provides information that allows the decision maker to access his or her past experience and existing store of knowledge in order to determine what to do.³³ In some cases this elicits an existing policy for dealing with situations like the one at hand. In other cases it provides background that, together with imagination, allows the decision maker to devise new ways of dealing with the situation. The parallels with N-BDM are both obvious and intentional; it was designed to include the insights provided by recognition decision theories.

Scenario Theory

Scenario theory describes how plausible stories can be constructed to forecast the future and guide planning.³⁴ A scenario consists of propositions that describe known and inferred causal relationships in the

³³ Simon (1979), Klein (1993, 1996).

³⁴ Jungermann (1985), Jungermann & Thüring (1987)

form of a causal model—a mental representation of what is going on and the forces that drive it. Forecasts result from “running” the model; using the causal relationships to answer the question, “What will happen if my propositions are accurate?” Other questions are, “What will happen if my propositions are wrong?” and “What would happen if I assumed this-or-that about my propositions instead of assuming that they are accurate?” The latter permits generation of alternative scenarios and what must be changed about the propositions to bring about more desirable alternatives.

The parallels with NBDM are not quite as straightforward as they might at first seem, but they aren't far off. The scenarios discussed above are not quite like the current narrative and its extrapolated forecast; they are more like plans in which a variety of tactics can be considered and an action forecast derived. Different tactics lead to different action forecasts, and those that produce an action forecast that isn't significantly discrepant from the desired future can be adopted and implemented.

Explanation Theory

The Pennington and Hastie research on jury decisions, described above, underlies this theory. It turns on decision makers' ability to interpret and evaluate new information (evidence in a trial) and to integrate it with their general knowledge about human behavior into an evolving explanation about what happened and why. The resulting stories are economical, in that there are few loose ends, coherent, and to the point. It is, of course, but a short step from the construction of a reasonable story to reaching a decision about the innocence or guilt of the accused. This theory and its accompanying research were a major factor in the design of NBDM.

Argument Theory

Mini-theories that see arguments about the pros and cons of a decision also assume that scenarios and stories are part of the process.³⁵ In addition, they consider the thinking that goes into the decision, both before it is made and afterward, when it has to be defended. The central notion is that the decision maker assesses the decision situation and, drawing upon past experience and general knowledge, formulates a course of action that meets the demands of the situation. But this is just the first step, because the course of action is then tempered and reshaped in light of arguments for and against it—originating with the decision maker or with an outside advice-giver. Arguments are in the form of “Do A because R,” where A is an action and R is a reason.

The decision maker can entertain alternative actions at the same time, using arguments to refine and differentiate them until one emerges as superior as well as defensible. In some versions of the theory, differentiation continues after the action has been implemented in order to avoid post-decision regret (“buyer’s remorse”).

Argument theory is more about procedure than about the substance of decisions. In this sense it is something of a side-theory or supplementary theory for NBDM. It is about how the decision makers clarify narratives and plans using both their own knowledge and, in some cases, information provided by other people. Moreover, its goal is as much to produce defensible actions as to produce effective actions, which makes it suitable for political decisions. NBDM could account for this by including the need to be able to answer criticism as a feature of the desired future.

³⁵ Lipshitz (1993), Svenson (1992), Montgomery (1993).

Reflexivity

Reflexivity is a feedback mechanism offered by the financier and philanthropist, George Soros to explain how “bubbles” occur in markets.³⁶ The idea is that that decisions produce changes in the market which in turn changes subsequent decision behavior. This kind of feedback loop is familiar to psychologists but not, I gather, to economists. The prevailing economic view has been that that countervailing forces move markets toward equilibrium between supply and demand, which sets prices—but, there are small variations across time. If I understand Soros’ argument, individual and institutional investors seize upon upward variations (or, presumably, other encouraging information) to buy and, if enough of them do it, the effect is to move the market even further upward, prompting further buying, which moves the market even further upward, which prompts more buying, and so on. The result is a bubble that becomes larger and larger as investors seek to profit from an ever rising market. The problem is that the underlying assets haven’t really increased in value—production, product quality, or sales haven’t increased—so the bubble is based on nothing but exuberance. Eventually something pricks the bubble and the same reflexive process that produced it now operates to push the market increasingly downward, wiping out the apparent wealth that had been created by the bubble and, perhaps, real wealth as well. Soros, like many critics of the existing financial system, observes that bubbles are endemic to unfettered capitalism, causing cycles rather than equilibrium, so some sort of regulatory oversight is needed to dampen the cycles.³⁷

³⁶ Soros, 2008.

³⁷ This raises the question of what kind of oversight is needed. Current laws and regulatory measures are influenced by current economic theory, the heart of which (if somewhat obscurely) is the gamble analogy. That is, laws and regulations are designed to protect the investor-as-gambler. They take particular aim at “cheating,” meaning any attempts to reduce the risk associated with investment. This works for the ordinary

Reflectivity is compatible with NBDM because the latter assumes that the entire point of action is to change the environment and that feedback about its success or failure influences subsequent actions. In this sense, reflexivity is already part of the theory. However, Soros' analysis of bubbles is instructive; it suggests that much the same thing might happen to individual decision makers in the course of their daily lives. Presuming that feedback frequently is vague, it is possible to think of spurious but apparently favorable information about the success of a plan's implementation leading to changes in it that are aimed at capitalizing on its apparent progress. If everything remains vague enough, the decision maker can pursue this ratcheting up of the plan into something that resembles a market bubble. This may be what allows decision makers to think that failing plans are succeeding, encouraging them to pour in more and more resources in the belief that things are getting better and better—until the bubble bursts, usually in the form of clear, undeniable proof that things aren't going well after all. It also may be what allows people to believe that their efforts to produce good outcomes by engaging in superstitious rituals actually work. After all, vague information is vague precisely because it can be interpreted in various ways. The notion of personal bubbles is intriguing and deserves to be explored further.

investor; he or she actually is gambling when buying a financial product. But institutions are another matter. The financial industry has worked hard to invent financial instruments that pass risk on to someone else, often to unsuspecting purchasers—repackaged risky mortgages for example. This means that even those whose job it is to gamble within the framework of the global financial system don't want to gamble. They, like every decision maker, seek ways of reducing risk by taking charge of things and making the desired future happen. It is my opinion, not necessarily Soros', that it is time to re-think the aims of regulation. Continued reliance on the gamble analogy may prevent fresh and creative thinking about these problems.

Incremental Theories

There are two kinds of incremental theories, incremental evaluation and incremental implementation. Both of them are reflected in NBDM.

Incremental evaluation was “discovered” through the observation that despite the rhetoric of politicians, public policy is shaped less by what policy makers want to move toward than by what the public wants to move away from.³⁸ Even at that, change seldom occurs as a result of sweeping reform, it occurs by “nibbling” away at problems, observing the results of small changes, making adjustments and nibbling again, and so on. The emphasis is on the undesirable parts of the status quo and actions are aimed at remedying those undesirable parts. The result is that most decisions result in little change from the previous state of things, differing only in improving the most obvious flaws.

Incremental implementation means that the decision maker receives feedback during decision implementation that modifies his or her knowledge about the environment in which the implementation is taking place.³⁹ This prescribes changes in the action, which generates more feedback, and so on. But, it is more than just a feedback loop; the information obtained as a result of the decision makers' actions may have profound effects on his or her values, intentions, goals, and so on. Even so, the feedback loop provides a way for the decision maker to “feel his or her way along” toward a more desirable state of affairs than the one that gave rise to the decision.

NBDM borrows freely from both kinds of incremental theory. Incremental evaluation is reflected in the idea that plans address the undesirable (flawed) portions of a forecasted future (to “repair” the forecast), with the result that, when it arrives, the actual future usually is

³⁸ Lindbloom (1959)

³⁹ Connolly (1988)

not radically different from the forecast. Incremental implementation is reflected in the idea that feedback during plan implementation allows the decision maker to take stock and adjust what he or she is doing—sometimes changing direction if necessary or stopping if the result is good enough.

Deontology

Deontology refers to the influence of moral obligation and commitment on human behavior.⁴⁰ It is built on the observation that most humans are solidly anchored in a social context, which influences decision making in three ways. The utilitarian influence is the same as Economic Man's utility for goods, services and money. The social influence corresponds to both the codes of behavior for the decision maker's reference group and the cultural values of the larger community. The deontological influence reflects the moral and ethical considerations that guide the decision maker's behavior. The difference between social and deontological influences is that the former exercise their influence by threatening disapproval and ostracism if they are violated. In contrast, morals and ethics exercise their influence by a sense of obligation, commitment, and duty, with conscience rather than fear providing the motivation.

In some sense, deontological influences are at the very heart of NBDM. The desired future, around which every decision revolves, is dictated by the decision maker's values—which includes morals, ethics, and preferences. Utilitarian influences on decisions are exercised through preferences. Social influences are exercised through both internalized social codes in the form of ethics and through preferences to avoid social approbation. And deontological influences are exercised through moral imperatives.

⁴⁰ Etzioni (1988, 1993)

NBDM and Cognitive Errors

I recently read a review of yet another book that blames the follies of business and government—this time it was belated reaction to global warming—on decision makers' cognitive errors; not just ideology or plain foolishness, but ingrained flaws in their thinking. As I write, there are any number of similar books, articles, and editorials that blame these same cognitive errors for the 2008 financial meltdown. Indeed, cognitive errors have become the sophisticate's substitute for what most people regard as simple greed, stupidity and pig-headedness.

Biases

Cognitive errors, usually called “biases,” are deviations of decision makers' judgments or choices from some paradigmatic standard, usually deriving from probability theory or Economic Decision Theory. The literature on cognitive errors dates to the 1960's, and almost 50 years of research has resulted in an enormous list. Wikipedia lists 29 probability and belief biases, 22 social biases, 7 memory biases, and 35 assorted biases, for a grand total of 93, and I suspect the list is out of date. The question of interest is whether some of these cognitive errors make sense if they are viewed as the result of normal narrative thinking instead of as erroneous paradigmatic thinking. After all, the paradigms were developed in the first place to help when narrative thinking isn't enough, so we'd expect “errors” if decision makers engage in narrative thinking when experimenters expect them to engage in paradigmatic thinking.

Recall that underlying dynamic of narratives is causation and purpose; actors cause events to occur for a purpose. Moreover, good narratives are tidy—there are no loose ends—and they are resistant to all but minor change. And, because they exist to help us make sense of what is going on around us, they guide our attention to and interpretation of

events to ensure that we don't miss things that are relevant to the narrative.

Causal biases. Unfortunately, for all their usefulness, narratives don't work well in environments for which they are unsuited. For example, because they are founded on causality, narratives do not deal well with randomness (which, by the way is a rather recently invented concept). So, when confronted with an array of randomly occurring events—flashing lights, for instance—we tend to see order where none exists. For the same reason, we don't do well with correlated events; we infer causation, and even then the inference is flawed. That is, when we infer causation we focus on cases in which the cause and the effect both occur or both don't occur; we overlook the cases in which one occurs without the other, which argue against causation—or at least against a strict definition of causation. Moreover, we tend to have a rather shallow view of causal relationships in the world around us. We usually are satisfied if we understand simple what-causes-what relationships. If we have loose ends, we may try to figure out what caused the causes, but two levels of causation is about all we can manage. When events show real or imagined trends for which we have no readily understood cause, or the causal factor is third level or beyond, we tend to make up mega-causes—the hand of God, the Market (as in “the market is correcting itself”), conspiracies, and so on. Each of these allows us to tie together disparate events and multiple levels of causation in a nice bundle with what appears to be a single overarching cause.

The causal logic in a decision maker's narrative is governed by the rules and expectations that were described in Chapter 3. A common rule involves justice; good things result from well-intended actions and bad things result from ill-intended actions. Thus, when we observe something bad happening, our rule leads us to expect that someone did something

bad—either the victim, who is getting his just desserts, or someone else, who is motivated by malice. You see this when disasters (effects) occur; some people, particularly preachers, infer that the victims brought it on themselves by sinning (cause) or it happened because someone in authority failed to prevent it. By and large, narratives resist ambiguity, so the idea that “stuff happens,” no matter what the bumper stickers proclaim, is difficult to accept. We need causes to make our narratives neat and meaningful—there is no meaning in happenstance.

Information biases. Narratives tend to be robust. We resist making massive revisions in them even when we have information that argues for doing so. One reason is that it takes work to revise a narrative because so many events must be reinterpreted and the causal links that we believed to be clear become murky or change altogether. Sometimes events will prompt an epiphany, an insight that results in a thorough overhaul of a narrative, but this is so rare as to be remarkable. The result of our resistance to narrative change is that we constantly are on the lookout for congruent information while ignoring or distorting anything that is incongruent with the narrative. This may be dangerous, but it makes sense because incongruent information is discrepant, and discrepancies drive decision making; if we take incongruent information at face value, we have to do something about it. It is far easier to ignore or distort it—at least up to the point at which doing so starts to lead to trouble. After all, a faulty narrative is going to produce faulty forecasts; if the forecasts are falsely optimistic, we will be disappointed when the future arrives, and if they are falsely pessimistic, we will find ourselves working to produce a more desirable future when we don't have to.

Narratives are robust and resistant to revision because, while incongruent information tells us they are wrong, it doesn't always tell us what to do about it. Because it is better to have a weak or faulty narrative

than a patched together narrative that doesn't make as much sense as the old one, or, worse, no narrative at all, so we tend to hold on to our old narratives. Oddly enough, although we all resist giving up our narratives, we can be very critical when other people do it. We accuse them of being closed minded or set in their ways. But, in fact, we are all rigid when it comes to the core narratives that define who we are and what we truly believe about how the world works. We are more flexible about peripheral narratives—most of us will revise narratives in which Pluto is included among the planets because narratives involving planets aren't particularly important (but think about the guy who originally declared Pluto to be a planet—or his children or grandchildren—I suspect there is resistance). The result of all of this is that we tend to be very resistant to efforts to change our minds about anything that is important to us. Logic and reason and empirical evidence are all very nice, but they seldom prevail against a robust narrative.

Statistical biases. Many of the cognitive errors that are cited in the decision literature result from the use of statistical concepts as standards—randomness, the law of large numbers, correlation, base-rates, regression toward the mean, conjunctions of probabilities, additivity of probabilities, and so on. Statistics, including probability theory, is a paradigm for classifying information and making transformations on it in order to answer questions and make predictions without resorting to causal reasoning; what and how, but not why. Why is supplied by the user of the paradigm, it is not part of the paradigm itself.

Of course, the statistical paradigm was developed, in the first place, because narrative thinking doesn't do a very good job of non-causal reasoning. Small wonder that the discovery of the bell shaped curve was regarded by some observers as discovery of a natural law—the arrangement of observations merely by frequency, rather than by cause,

was a revolutionary idea. Narratives don't do it very well, so paradigms had to be developed for helping narratives incorporate what could be learned in this strange new way. Statistics and probability theory are not natural ways of thinking, as anyone who has ever taught a statistics class can testify. Moreover, the controls that are used in experiments and the rituals that have grown up around the use of statistics, exist to prevent researchers from imposing narrative thinking on their data—especially the confirmatory imperative that narratives impose. Even at that, those of us who are versed in the statistical paradigm tend to slip. For example, it is difficult not to interpret correlation as causation or to regard a small but statistically significant difference as more momentous than it actually is.

Narrative thinking doesn't account for all the cognitive errors in the Wikipedia list, but there is a study that suggests that at least some of them reflect narrative, rather than paradigmatic, thinking. The study was done by Valerie Barnes as part of her doctoral dissertation.⁴¹ She began with 15 word problems that have been commonly used to demonstrate judgmental errors: 5 problems for the “sample size bias,” 5 for the “base rate bias,” and 5 for the “conjunction fallacy”. The sample size bias is the tendency to treat information based on small samples as though it were as reliable as information based on large samples, which violates the law of large numbers. The base rate bias is the tendency to ignore information about differences among the frequencies of occurrence of various classes of events (their base rates) when making judgments about the class membership of a randomly drawn event, which violates the logic of probability theory. The conjunction fallacy is the tendency to think that the probability of the co-occurrence (the conjunction) of two events is equal to or greater than the probability of either event occurring independently, which violates the logic of probability theory.

⁴¹ Barnes (1984).

The problems were written on cards and presented to undergraduate students who had never had a statistics course. They were asked to sort them into piles of similar problems, and then to assign labels to each pile and describe what make the problems in each pile similar to one another.

It was found that the assigned labels and definitions did not at all reflect sample size, base rate, or conjunction—the three probability concepts that experiments presumed to be key. Instead, the students divided the problems according to content, such as “math problems” (not “probability problems,” just “math problems”) or “problems about people,” labels that reflect narrative thought rather than anything having to do with paradigms in general or probability theory in specific. Of course, none of the students had been taught about probability, which might make even the least skeptical among us question its use as the criterion for “rationality”.

Then the students were asked to solve each problem while talking aloud about what they were doing. In some cases, what was said revealed that they grasped the feature of the problem that most experimenters would regard as essential, although they didn't know what to do with it. In other cases, however, their approach was substantially different. In general, they attempted to use some kind of mathematical reasoning on problems that they regarded as mathematical and they used causal reasoning on the other problems. Mathematical reasoning requires knowledge of mathematical principles and procedures, the essence of paradigms. Causal reasoning requires knowledge of the reasons that people do the things they do and why events occur as they do, the essence of narratives. The students appeared to appreciate the distinction and approached the tasks accordingly even though they were hindered by not knowing the proper paradigm, probability theory.

This study suggests that at least some so-called cognitive errors may result from overeducated experimenters and undereducated participants operating at cross purposes. Experimenters design the problems for paradigmatic solution but, lacking the paradigm, participants stick with narrative thinking. Even when they suspect that a paradigmatic solution is appropriate, because they lack knowledge of the paradigm, they are bound to make errors. Because experimenters decide what the correct answer is, and because they publish the results, differences between the paradigmatically derived answer and the participants' judgments are labeled as biases.

Probably not all errors are the result of miscommunications between experimenters and participants, but many probably are. Because experimenters regard their paradigmatically derived answers as singularly correct and never consider what the narrative counterparts might be (or, if they do, they consider them to be wrong), we're unlikely to ever know which observed errors represent actual cognitive errors and which simply represent narrative thinking that differs from paradigmatic prescriptions. (Perhaps the real test is whether participants' judgments are systematically in error when they profess to be using the paradigm the experimenter thinks they should be using. But, to be fair, they must be given instruction about the paradigm and permitted to use pencil, paper, and calculators—just like the experimenters did when they were making up the problems. Of course, if you did this, the experiment would look like nothing more than the final exam in a statistics class.)

Heuristics

Heuristics are rules of thumb, the use of which produces erroneous, biased, judgments. Sometimes heuristics are talked about as though they were short cuts to reduce cognitive effort. The problem is that if decision makers are attempting to reduce cognitive effort, it implies that they

know the “right” way to make the judgment and how much effort it would take—a shortcut is defined by what it avoids as much as by what it entails.

Like the list of biases, the list of heuristics has grown over the years.

However, most of them seem to be variations on the original three:

representativeness, availability, and anchoring and adjustment.⁴²

Representativeness. Representativeness is a form of stereotyping in which the characteristics of an individual or event are assumed to mirror the assumed characteristics of the group of which the individual is a member. Thus, people think it is more likely that a randomly selected individual who is shy, withdrawn, and tidy is a librarian than it is that he or she is a lawyer or teacher, even though there are far fewer librarians in the world than lawyers or teachers—which makes it less likely that a *randomly selected* individual is a librarian (base rates). Or, people expect a sample of marbles from an urn containing 60% white marbles and 30% black to itself contain roughly 60% white marbles, which is true only for large, *randomly selected*, samples (law of large numbers).

The errors that are attributed to representativeness somehow make intuitive sense from a narrative viewpoint. Keep in mind that most research participants are untrained in statistics (although some studies suggest that training makes little difference), so they lack the paradigms necessary to derive the answers that experimenters require. Lacking these paradigms, and lacking familiarity with the different frequencies of occupations in the population, they might reasonably think that the person being described to them would, in the normal course of things, seek the company of like-minded persons and therefore resemble these other persons. Similarly, lacking familiarity with the variability of samples of marbles from an urn of marbles, a familiarity that few of us are likely to have, they might

⁴² Tversky & Kahneman (1974).

reasonably think that any sample from an urn should be more similar to the urn's contents than dissimilar.

The thing that makes these answers wrong, from the experimenter's point of view, is that the person being described and the sample from the urn are supposedly selected *at random*. Randomness is a basic concept in statistics, but very few people really understand it; most people think it merely means disorderly or haphazard.⁴³ Because they don't understand randomness or its use in statistics, participants in most of these experiments can't understand the implications of being told that the person, or the sample from the urn, was selected randomly. The result is that the experimenter is using one kind of reasoning and the participants are using another.

Box 8a: Playing the Lottery

Nearly every State has one or more lotteries and there are a few, like Powerball, that are nationwide. All of them appear to be thriving even though most lottery players are well aware that the odds against their winning are overwhelming—less than being struck by lightning. Still, they play. Why?

Scholars differentiate between risk and uncertainty. For them, risk pertains to situations in which the mathematical probabilities of alternative outcomes can be calculated; uncertainty pertains to situations in which probabilities cannot be calculated. The rest of us make an entirely different distinction: risk pertains to situations involving precariousness, jeopardy, or vulnerability, and uncertainty pertains to situations involving equivocality, ambiguity, or vagueness. Roughly speaking, risk equals danger and uncertainty equals doubt.

For scholars, lotteries are properly referred to as risky because it is possible to compute the mathematical probability of a single lottery ticket winning the jackpot. And, because this probability is usually very small, they, the scholars, conclude that it is irrational to spend money on a lottery ticket.⁴⁴ In contrast, for the rest of us, there isn't much risk (danger) because the amount to be lost is small. Moreover, for us the small print on the ticket stating that the odds of winning are, say, 1 in 250,000, merely means it is a long shot, nobody can actually comprehend a number as small as

⁴³ In fact randomness is a prescription about how a sample is to be drawn from a population and, strictly speaking, if the prescription isn't followed, statistical reasoning doesn't apply.

⁴⁴ That is, buying a lottery ticket is equivalent to making a bet so Economic Man's logic for analyzing bets (Chapter 1) applies; the expected value for a lottery ticket is negative, so it is unwise to buy it.

1/250,000. But, small as it may be, somebody is going to win, eventually and there's no reason why it shouldn't be me. And, if I win, just think of all the things I could do with the money!

The winning lottery numbers are determined randomly, but few people understand what random really means, probably because we almost never come across it in real life. As a result, few of us truly appreciate how indeterminate the outcome of a lottery actually is. Instead, we apply the causal reasoning that usually works for us; we assume that things don't simply happen, they are caused to happen, however obscure the cause may be. It follows, therefore, that the winning numbers must be caused by something and where there is a cause, there is the potential for influencing the result.

Of course, we can't physically influence the lottery drawing, but we can call upon indirect influences; prayer, positive thinking, prescience, shrewd guesses, good luck, fate, and so forth. Moreover, if we have played frequently in the past, we can figure that we're about due to win. This kind of reasoning may be regarded by scholars and other right-thinking people as rank superstition, but a mere glance at history reveals that superstition, rank or otherwise, has always been part of how humans tame uncertainty.

Having dealt with the rotten odds of winning by ignoring them or mitigating them, we deal with any residual uncertainty (doubt) by constructing a plausible, coherent narrative about how we'll spend the money—keep only a million or two and generously give the rest to worthy recipients, etc., etc., etc. A plausible, coherent narrative gives us confidence that we know what the future will bring and that we are in control, even if we don't and aren't.

In terms of NBDT, this fanciful narrative about how we'll spend our winnings is an action forecast for the potential plan of buying a lottery ticket. Like any action forecast, we compared it to our values and preferences—which it is bound to satisfy because, except for that niggling little bit about actually winning, it is virtually dictated by those same values and preferences. Thus, when it inevitably passes the test, we have, in effect decided to buy a ticket. If all of this sounds irrational, here's an argument in its favor: Only Economic Man could ever lose in a lottery—because he has no emotions. For the rest of us, even if we don't win the jackpot, at least we had the fun of thinking about how we would have spent it. Surely this, together with a few moments of hope, is worth a few dollars.

To use a more intuitive example: Suppose an experimenter were to *randomly* select a church in the United States and visit it one Sunday afternoon in June, the traditional month for weddings in America. He stands outside and waits for the wedding party to emerge. Then he approaches the Best Man, explains that he selected this church at random and asks “What is the probability that the newly married couple for whom you were Best Man will still be married to each other ten years from now?”

For the experimenter, the question is based on his appreciation of randomness and statistical inference—he randomly selected this couple and he knows absolutely nothing about them as individuals. For him the probability their marriage will last ten year is the base rate for divorce in the United States, which, if I recall, is something like .50. For the Best Man, however, the information about the church being selected at random implies nothing because the question is about his friends. His answer will reflect his knowledge about them as individuals and his narrative about their relationship and about what makes for an enduring marriage. In short, the experimenter uses statistical, paradigmatic thinking and, because he mentioned randomness, expects the Best Man to use it too. However, the Best Man, having little or no understanding randomness (which, in any event, he wouldn't regard as relevant to questions about his friends), uses narrative thinking. The experimenter, being the score-keeper, declares the Best Man's answer to be wrong.

Availability. Availability means that a judgment about the probability of an event is based on how easily exemplars come to mind or can be imagined to occur. Cognitive research shows that frequent events are, in fact, more easily recalled than infrequent events, unless the latter are particularly striking; so ease of recall isn't a wholly spurious clue about probability. Thus, people think that airplanes are more dangerous than automobiles in spite of the fact that there are far more deaths each year in automobile accidents than in plane crashes. The availability explanation is that because plane crashes most frequently make headlines and highway accidents rarely even make the back page, the former are more easily brought to mind and therefore are judged to be more likely. Moreover, people seem to more easily imagine dire outcomes when others (the pilots) are in control than when they're in control (driving

their car), which also contributes to overestimates of the dangers of flying and underestimates of the dangers of driving.

In narrative thinking, decision maker's evaluations of the plausibility of plans depend, in part, on how easy it is to imagine successful implementation, even if it will require heroic effort. (And plans must be plausible, in that they don't rely on fantastic or unlikely events.) This suggests that availability is the same or similar concept. People regard plausible, easily recalled or imagined events to be more probable than implausible events.

I don't think that there is any question that availability, or something akin to it, contributes to over- and underestimation of probabilities, but the fact is that most decision maker's have no alternative to using it. We don't normally know the objective frequency data about every aspect of each of our many plans, so we have no alternative to using availability. However, this doesn't account for the fact that availability is robust in the face of contradictory information; having been repeatedly been told that driving is more dangerous than flying, few of us have an excuse for believing the opposite. Nonetheless, people are more likely to pray during takeoff than when they're pulling their car out of the garage.

_Anchoring and adjustment. Anchoring and adjustment means that judges often use a reference point, called an anchor, and adjust their judgment upwards or downwards from that point in light of other information. Thus, for example, for the same house, a higher asking price may elicit higher offers than a lower asking price would, suggesting that buyers are more influenced by the asking price as a reference point (anchor) for formulating their offers than they are by how much the house actually is worth to them.

Anchoring is a commonly reported phenomenon in perception, social psychology, and cognition. For example, whether an object is

judged to be “heavy” depends what it is. A pen that weighs more than 4 oz is judged to be heavy but a baseball bat has to weigh more than 40 oz to be judged heavy. In this case, the anchor is the typical weight of the objects belonging to the category to which the object in question belongs. Similar phenomena are found in social situations in which social norms serve as anchors in judging the normality, or deviance of observed behaviors. Before cell phones came on the scene, someone talking loudly on a street corner was presumed to be mentally ill or a preacher; now it's so common we (unfortunately) regard it as normal. These and similar phenomena demonstrate what is generally known in psychology as contextual effects on judgment.

It is the second part of the anchoring and adjustment heuristic, insufficient adjustment away from the anchor, which makes it interesting. The mechanism is rather mechanical—the judge focuses on the anchor and moves his or her judgment away from it in light of information that suggests that the thing being judged is more than or less than the anchor on some dimension. Whether this is done consciously or not isn't usually addressed, but the mechanical metaphor is retained. All in all, in the absence of an alternative, it probably isn't such a bad strategy, even though it frequently results in either more or less movement away from the anchor than the information may warrant. Errors tend to reflect our hopes and misgivings perhaps as much as they do inability to make accurate judgments.

NBDM doesn't address anchoring and adjustment, although it is likely that anchors influence the judgments of discrepancy magnitude and importance that decision makers provide for decision sheets. For example, when trying to provide judgments that represent the organization's priorities rather than their own, single agent decision makers may anchor on their own priorities and adjust them to reflect their view of

the organization's priorities. Conversely, they may anchor on what they think their peers think and adjust toward “good sense” (their own priorities). Either way, it is reasonable to expect that the adjustment is frequently too little or too much and the resulting judgments err in the one direction or the other.

Progress Report

In this chapter we have reviewed the theory of Narrative-Based Decision Making (NBDM) that was developed in the previous chapters. In addition, we have seen how the kernels of a number of non-economic mini-theories—recognition, scenarios, stories, arguments, reflexivity, incrementalism, deontology—have been incorporated into NBDM. And, we saw how narrative thinking may account for some of the cognitive errors that have gained such notoriety as the causes of human folly.

A parting thought: Narrative thinking may be the basic way of thinking, but it has its limits. Paradigms and paradigmatic thinking are purposely designed to take over when those limits are reached. By supplying the required rigor and precision, paradigms expand our cognitive capabilities far beyond what narratives can do alone. Together, narratives and paradigms make human thought a powerful force, especially when coupled with such tools as computers and telescopes and all the paraphernalia of science and technology. With them, we can explore atoms and the heavens, discover what makes us sick and makes us well, construct micro-chips and skyscrapers, forecast the future and change it to improve the quality of our own and others' lives. Instead of thinking of ourselves as hopelessly biased and mired in cognitive error, we should focus on (and perhaps be awed by) what can be achieved when narrative and paradigmatic thinking work together.

Sources and Further Reading

- Barnes, V. E. (1984). *The quality of human judgment: An alternative perspective*. Unpublished doctoral dissertation, University of Washington, Seattle.
- Beach, L. R. (1990). *Image theory: Decision making in personal and organizational contexts*. Chichester, UK: Wiley.
- Beach, L. R., & Connolly, T. (2005). *The psychology of decision making: People in organizations*. Thousand Oaks, CA: Sage.
- Beach, L. R., & Mitchell, T. R. (1990). Image theory: A behavioral theory of decisions in organizations. In B. M. Staw & L. L. Cummins (Eds.), *Research in organizational behavior* (Vol. 12). Greenwich, CT: JAI.
- Beach, L. R., Smith, B., Lundell, J., & Mitchell, T. R. (1988). Image theory: Descriptive sufficiency of a simple rule for the compatibility test. *Journal of Behavioral Decision Making*, 1, 17-28.
- Connolly, T. (1988). Hedge-clipping, tree-felling and the management of ambiguity: The need for new images in decision-making. In L. R. Pondy, R. J. Boland, Jr., & H. Thomas (Eds.), *Managing ambiguity and change*. New York: Wiley.
- Etzioni, A. (1988). *The moral dimension: Toward a new economics*. New York: Free Press.
- Etzioni, A. (1993). *The spirit of community*. New York: Crown.
- Jungermann, H. (1985). Inferential processes in the construction of scenarios. *Journal of Forecasting*, 4, 321-327.
- Jungermann, H., & Thüring, J. (1987). The use of causal knowledge for inferential reasoning. In J. L. Mumpower, L. D. Phillips, O. Renn, & V. R. R. Uppuluria (Eds.), *Expert judgment and expert systems*. New York: Springer.
- Klein, G. A. (1989). Recognition-primed decisions. *Advances in Man-Machine Systems Research*, 5, 47-92.

- Klein, G. A. (1993). A recognition-primed decision (RPD) model of rapid decision making. In G. A. Klein, J. Orasanu, R. Calderwood, & C. E. Zsombok (Eds.), *Decision making in action: Models and methods*. Norwood, NJ: Ablex.
- Klein, G. A. (1996). *Sources of power: The study of naturalistic decision making*. Mahwah, NJ: Erlbaum.
- Lindbloom, C. E. (1959). The science of "muddling through." *Public Administration Review*, 19, 79-88.
- Lipschitz, R. (1993). Decision making as argument-driven action. In G. A. Klein, J. Orasanu, R. Calderwood, & C.E. Zsombok (Eds.), *Decision making in action: Models and methods*. Norwood, NH: Ablex.
- Montgomery, H. (1993). The search for a dominance structure in decision making: Examining the evidence. In G. A. Klein, J. Orasanu, R. Calderwood, & C.E. Zsombok (Eds.), *Decision making in action: Models and methods*. Norwood, NH: Ablex.
- Pennington, N., & Hastie, R. (1988). Evidence evaluation in complex decision making. *Journal of Personality and Social Psychology*, 51, 242-258.
- Pennington, N., & Hastie, R. (1988). Explanation-based decision making: Effects of memory structure on judgment. *Journal of Experimental Psychology: Learning, Memory and Cognition*, 14, 521-533.
- Pennington, N., & Hastie, R. (1992). Explaining the evidence: Tests of the story model for juror decision making. *Journal of Personality and Social Psychology*, 2, 189-206.
- Soros, G. (2008). *The new paradigm for financial markets: The credit crisis of 2008 and what it means*. New York: Public Affairs.
- Svenson, O. (1992). Differentiation and consolidation theory of human decision making: A frame of reference for the study of pre- and post-decision processes. *Acta Psychologica*, 80, 143-168.

Tversky, A., & Kahneman, D. (1974). Judgment under uncertainty:
Heuristics and biases. *Science*, 185, 1124-1131.

Appendix 1: Research on the Discrepancy Test

[Note: In the following, the “compatibility test” is the same as the discrepancy test and discrepancy analysis, a “violation” is the same as a discrepancy on a feature, and a “characteristic” of an option is the same as a feature of an option.]

Starting in 1989, a series of laboratory studies examined the usefulness of the compatibility test for understanding *unaided* decisions. The general paradigm was to ask college students to make decisions about hypothetical options such as entry-level jobs, rental rooms, or time-share condos. They were instructed to examine the options' characteristics and form a choice set, and then, in some cases, to choose the best option from the choice set.

The early research (Beach & Strom, 1989) examined the role of violations and nonviolations in decisions. Participants were asked to assume the role of a newly graduated student who was looking for a job. The job seeker's preferences for 16 job characteristics were provided so all participants would be using the same standards. Then they were presented with the 16 characteristics of each of 16 different jobs. The characteristics were listed on successive pages of a small booklet and there was a booklet for each job. The characteristic on each page of a job's booklet did or did not violate one of the job seekers preferences (e.g., the jobseeker wanted to live in a large city; the violation was that the job was in a small city). Participants went through the pages of the book until they had seen enough to decide to reject the job or place it in the choice set. The order of violations and nonviolations for the various jobs was contrived to permit inferences about how the information was used to make the decision.

The results showed that rejection of jobs regularly occurred after observation of an average of roughly four violations; that is, the average rejection threshold was four violations. There was no comparably

consistent number of nonviolations for deciding to retain jobs for the choice set. In fact, nonviolations played no role at all in the decision, a result that has been observed in subsequent research (e.g., Beach, Smith, Lundell, & Mitchell, 1988; Potter & Beach, 1994a,b; Rediker, Mitchell, Beach, & Beard, 1993). Taken together, these studies demonstrate the primacy of violations in the decision process and the existence of the rejection threshold, two key concepts in the compatibility test. This is why nonviolations are scored 0 (zero). Incidentally, follow-up studies showed that it is the number of violated characteristics that determines the decision, not the proportion of the characteristics that are violated.

The psychological difference between screening and choice was investigated by Potter and Beach (1994b). One group of participants was given descriptions of time-share condos and asked to rate their attractiveness and then screen them to create a choice set. A second group was given the same descriptions and asked to choose the best. For both groups, the descriptions contained information about the likelihood that the condo would be available when the decision maker wanted to use it; either .25 or .75. The finding was that participants who were asked to screen regarded the .25 likelihood as a violation of the feature (desired availability) of the condo and .75 as an indication of a nonviolation. Participants who were asked to choose used the likelihood to discount the importance of the feature. This shows that information is used differently when the task is to screen and when it is to choose.

Potter & Beach (1994a, b) also examined what happens when information about a feature of an option is missing (it is assumed to be a violation), what happens when there is a paucity of information about some of the options (the rejection threshold goes down, making it easier to reject options about which you don't have sufficient information), what happens when the choice set ends up empty (participants prefer to start

over with new options; if forced to re-screen the previously rejected options, they reduce the importance of the features and they increase their rejection thresholds, both of which made rejection more difficult so at least one option could survive).

Research also has demonstrated that participants factor in differential feature importance (Beach, Puto, Heckler, Naylor, & Marble, 1996) and that they respond to time constraints by increasing the speed with which they do the test while decreasing the carefulness with which they do it. In addition, they raise their rejection thresholds so fewer options are rejected, lest good options get overlooked in the rush (Benson & Beach, 1996). Finally, a recent study shows that participants ignore small discrepancies between desired features and option's features but that above some critical point they treated discrepancies as violations, which is evidence for a violation threshold comparable to the rejection threshold. Moreover, that violation threshold decreased as the importance of the feature increased and as the number of other violations increased. That is, what might not qualify as a violation is more likely to do so if the feature's importance increases or if other violations have already been observed (Benson, Mertens, & Beach, 2008).

The compatibility test also has been examined in the context of monitoring the progress of plan implementation toward achievement of the desired future. Called "progress decisions," monitoring requires the decision maker to repeatedly test the difference between the plan's progress and the progress that it should have made at this point in its implementation. Dunegan (1995) found that decision makers were more likely to continue implementation and to commit additional resources to implementation when progress was compatible with expected movement toward the desired future. In a subsequent study (Dunegan, Duchon, & Ashmos, 1995) it was found that as the compatibility between

progress and goals increases, decision making becomes more automatic. That is, more careful monitoring is reserved for “crises;” decreases in compatibility resulting from a lack of progress.

The compatibility test also has been extended to real-world unaided decisions. For example, Beach, Smith, Lundell, & Mitchell, 1988) examined the decisions of executives of two sports clothes manufacturing firms and found that the decisions could be accounted for by a simple compatibility test. Other real-world work includes, audit decisions (Beach & Frederickson, 1989), planning decisions (Beach, DeBruyn, & Jungermann (1966), satisfaction with supervision (Bissell & Beach, 1996; Richmond, Bissell & Beach, 1998), job search and job selection (Stevens & Beach, 1996), selection of clients by audit firms (Asare, 1996), the effects of organizational culture on decisions (Weatherly & Beach, 1996), employee turnover (Lee & Mitchell, 1991; 1994), consumer decisions and social responsibility (Nelson, 1996), decisions about marketing and communication strategies (Puto & Heckler, 1996), career decisions (Stevens, 1996), and the constraints of cultural values on group decisions (Walsh, 1996).

Bibliography for the Compatibility Test

Two anthologies contain many of the studies cited above (Beach, 1996; 1998). The following is a bibliography of theory and research in which the compatibility test has been a key point of interest. The citations are listed chronologically within category.

Books

Beach, L. R. (1990). *Image theory: Decision making in personal and organizational contexts*. Chichester, UK: Wiley.

Beach, L. R. (1993). *Making the right decision: Organizational culture, vision, and Planning*. Englewood Cliffs, NJ: Prentice Hall

Beach, L. R. (Ed.) (1996). *Decision making in the workplace: A unified*

- perspective*. Mahwah, NJ: Erlbaum
- Beach, L. R. (Ed.) (1998). *Image theory: Theoretical and empirical foundations*. Mahwah, NJ: Erlbaum.
- Beach, L. R., & Connolly, T. (2005). *The psychology of decision making, Second edition*. Thousand Oaks, CA: Sage.
- Articles from My Laboratory*
- Mitchell, T. R., Rediker, K. J. & Beach, L. R. (1986). Image theory and its implications for organizational decision making. In H. P. Sims & D. A. Gioia (Eds.), *The thinking organization*. San Francisco: Jossey-Bass.
- Beach, L. R. & Mitchell, T. R. (1987). Image theory: Principles, goals and plans in decision making. *Acta Psychologica*, 66, 201-220.
- Beach, L. R. & Mitchell, T. R. (1987). More on image theory: A reply to Montgomery and Vlek. *Acta Psychologica*, 66, 231-235.
- Beach, L. R., Smith, B., Lundell, J. & Mitchell, T. R. (1988). Image theory: Descriptive sufficiency of a simple rule for the compatibility test. *Journal of Behavioral Decision Making*, 1, 17-28.
- Beach, L. R. & Frederickson, J. R. (1989). Image theory: An alternative description of audit decisions. *Accounting, Organizations and Society*, 14, 101-112.
- Beach, L. R. & Strom, E. (1989). A toadstool among the mushrooms: Screening decisions and image theory's compatibility test. *Acta Psychologica*, 72, 1-12.
- Beach, L. R. & Mitchell, T. R. (1990). Image theory: A behavioral theory of decisions in organizations. In B. M. Staw & L. L. Cummings (Eds.), *Research in organizational behavior (Vol. 12)*, Greenwich, CT: JAI Press.
- Mitchell, T. R. & Beach, L. R. (1990). "...Do I love thee? Let me count..." Toward an understanding of intuitive and automatic decision

- making. *Organizational Behavior and Human Decision Processes*, 47, 1-20.
- Beach, L. R., Mitchell, T. R., Paluchowski, T. F. & van Zee, E. H. (1992). Image theory: Decision framing and decision deliberation. In F. Heller (Ed.), *Decision making and leadership*. Cambridge, UK: Cambridge University Press.
- van Zee, E. H., Paluchowski, T. F. & Beach, L. R. (1992). The effects of screening and task partitioning upon evaluations of decision options. *Journal of Behavioral Decision Making*, 5, 1-23.
- Beach, L. R. & van Zee, E. (1992). Close but no cigar: A reply to Pitz. *Journal of Behavioral Decision Making*, 5, 22-23.
- Beach, L. R. (1993). Four revolutions in behavioral decision theory. In M. Chemers & R. Ayman (Eds.), *Leadership theory and research: Perspectives and directions*. New York: Academic Press.
- Beach, L. R. (1993). Image theory: Personal and organizational decisions. In G. A. Klein, J. Orasanu, R. Calderwood & C. E. Zsombok (Eds.), *Decision making in action: Models and methods*. New York: Ablex.
- Rediker, K. J, Mitchell, T. R., Beach, L. R. & Beard, D. W. (1993). The effects of strong belief structures on information-processing evaluations and choice. *Journal of Behavioral Decision Making*, 6, 113-132.
- Beach, L. R. (1993). Image theory: An alternative to normative decision theory. In J. Muncy & L. McAlister (Eds.), *Advances in Consumer Research*, vol. 20. Provo, UT: Assoc. for Consumer Research.
- Beach, L. R. (1993). Broadening the definition of decision making: The role of pre-choice screening of options. *Psychological Science*, 4, 215-220.
- Beach, L. R. & Jungermann, H. (1994). Generating, selecting and monitoring plans of action. *Chinese Journal of Applied Psychology*, 9, 1-23.

- Potter, R. E. & Beach, L. R. (1994). Decision making when the acceptable options become unavailable. *Organizational Behavior and Human Decision Processes*, *57*, 468-483.
- Potter, R. E. & Beach, L. R. (1994). Imperfect information in pre-choice screening of options. *Organizational Behavior and Human Decision Processes*, *59*, 313-329.
- Connolly, T. & Beach, L. R. (1994). The theory of image theory: An examination of the central conceptual structure. In J.-P. Caverni, M. Bar-Hillel, F. H. Barron & H. Jungermann (Eds.), *Contributions to Decision Making I*. Amsterdam: North Holland.
- Beach, L. R., DeBruyn, E. E. J., & Jungermann, H. (1996). The role of imagination in planning decisions. In L. R. Beach (Ed.), *Decision making in the workplace*. Mahwah, NJ: Erlbaum.
- Bissell, B. L., & Beach, L. R. (1996). Satisfaction with job supervision. In L. R. Beach (Ed.), *Decision making in the workplace*. Mahwah, NJ: Erlbaum.
- Stevens, C. K., & Beach, L. R. (1996). Job search and job selection. In L. R. Beach (Ed.), *Decision making in the workplace*. Mahwah, NJ: Erlbaum.
- Weatherly, K. A., & Beach, L. R. (1996). Organizational culture and decision making. In L. R. Beach (Ed.), *Decision making in the workplace*. Mahwah, NJ: Erlbaum.
- Beach, L. R., Puto, C. P., Heckler, S. E., Naylor, G., & Marble, T. A. (1996). Differential versus unit weighting of violations, framing, and the role of probability in image theory's compatibility test. *Organizational Behavior and Human Decision Processes*, *65*, 77-82.
- Benson, L., III, & Beach, L. R. (1996). The effects of time constraints on the prechoice screening of decision options. *Organizational Behavior and Human Decision Processes*, *67*, 222-228.

- Richmond, S.M., Bissell, B. L. & Beach, L. R. (1998). Image theory's compatibility test and evaluations of the status quo. *Organizational Behavior and human Decision Processes*, 73, 39-53.
- Ordonez, L. D., Benson, L., & Beach, L. R. (1999). Testing the compatibility test: How instructions, accountability, and anticipated regret affect prechoice screening of options. *Organizational Behavior and Human Decision Processes*, 78, 63-80.
- Benson, L., Mertens, D.P. & Beach, L. R. (2009). When a difference makes a difference in the screening of decision options. Working paper, Department of Management, University of Arizona, Tucson, AZ.

Articles from Other Laboratories

- Lee, T. W. & Mitchell, T. R. (1991). The unfolding effects of organizational commitment and anticipated job satisfaction on voluntary employee turnover. *Motivation and Emotion*, 15, 99-121.
- Dunegan, K. J. (1993). Framing, cognitive modes, and image theory: Toward an understanding of a glass half full. *Journal of Applied Psychology*, 78, 491-503.
- Lee, T. W. & Mitchell, T. R. (1994). An alternative approach: The unfolding model of voluntary employee turnover. *Academy of Management Review*, 19, 51-89.
- Asare, S. & Knechel, W. R. (1995). Termination of information evaluation in auditing. *Journal of Behavioral Decision Making*, 8, 163-178.
- Dunegan, K. J. (1995). Image theory: Testing the role of image compatibility in progress decisions. *Organizational Behavior and Human Decision Processes*, 62, 79-86.
- Dunegan, K. J. (2003). Leader-image compatibility: an image theory view of leadership. *Journal of Business and Management*, 9, 61-77.
- Dunegan, K. J., Duchon, D. & Ashmos, D. (1995). Image compatibility and the use of problem space information in resource allocation

- decisions: Testing a moderating effects model. *Organizational Behavior and Human Decision Processes*, 64, 31-35.
- Dunegan, K. J., & Hrivnak, M. W. (2003). Characteristics of mindless teaching evaluations and the moderating effects of image compatibility. *Journal of Management Education*, 27, 280-303.
- Nelson, K. (1996). Consumer decisions involving social responsibility. In L. R. Beach (Ed.), *Decision making in the workplace*. Mahwah, NJ: Erlbaum.
- Puto, C. P. & Heckler, S. E. (1996). Designing marketing plans and communication strategies. In L. R. Beach (Ed.), *Decision making in the workplace*. Mahwah, NJ: Erlbaum.
- Pesta, B., Kass, D., & Dunegan, K. (2005). Image theory and the appraisal of employee performance: To screen or not to screen? *Journal of Business and Psychology*, 19, 341-360.
- Stevens, C. K. (1996). Career Decisions. In L. R. Beach Ed.), *Decision making in the workplace*. Mahwah, NJ: Erlbaum.
- Walsh, K. R. (1996). Mitigating cultural constraints on group decisions. In L. R. Beach (Ed.), *Decision making in the workplace*. Mahwah, NJ: Erlbaum.
- Seidl, C., & Traub, S. (1998). A new test of image theory. *Organizational Behavior and Human Decision Processes*, 75, 93-116.
- Pesta, B., Kass, D., & Dunegan, K. (2005). Image theory and the appraisal of employee performance: To screen or not to screen? *Journal of Business and Psychology*, 19, 341-360.

Appendix 2: Interpreting D'

Whether or not a particular D' exceeds your rejection threshold is a judgment you must make. Unfortunately, most people find that, for complex analyses, by the time the numbers have been written down and all the computations completed, the obtained D' has lost all intuitive meaning. Meaning can be restored by providing an understandable measurement scale for D' . Such a scale requires a meaningful zero point and a meaningful unit of measurement.

Basic Discrepancy Model

For the basic discrepancy model, the lower end of the scale, zero, is defined by the case in which a forecasted future has no discrepancies. In the case in which there are discrepant features, the unit is a single discrepancy, labeled 1, and D is a multiple of the unit, up to the upper limit of the measurement scale, n , the number of features in the analysis. Thus, when there are no discrepancies, D is zero and when all of the features are discrepant, D is n . When there are in-between numbers of discrepancies, D equals that number.

Elaborated Model

In what follows, we will construct a measurement scale for the elaborated model that has the same properties as the scale for the basic model, a meaningful zero and a unit of measurement, but it will take a little work.

Zero. As before, the zero point on the scale represents a forecast without discrepancies. For example, consider what would happen for the decision in Figure 6.2. A forecast without discrepancies would result in every entry in the magnitude column, m_c , being, zero. This would make all products and sums of products zero, irrespective of feature importance, and D' would

therefore be zero. This is not only the case for the analysis in Figure 6.2, it also is the case for the more complicated analysis in Figure 6.3 and, it would be the case for any analysis because multiplying by zeros yields products equal to zero.

Unit. The unit of measurement is a single discrepant feature. The question is what this means when discrepancies are measured as magnitudes rather than all-or-none. Complicating things further, as you can see in Figures 6.2 and 6.3, for complex decisions, importance of each feature also can be nuanced. As a result, for complex analyses, it isn't immediately clear what it means to define the unit as a single discrepant feature. What follows is the solution to this problem, which requires us to derive a *scaling unit*, a number, \mathbf{x} , that translates the results of our computations into units of our scale.

In every analysis, the magnitude of a discrepancy is fairly straightforward, a number from 0 to 3, for example. But, for reasons that are explained in the text, complex level 2 and level 3 decisions are most easily managed by segregating the individual features into sets of similar features and, sometimes, segregating the sets into groups of similar sets. The resulting *hierarchies* can have different numbers of tiers: a 1-tier hierarchy is merely a list of features; a 2-tier hierarchy is the list of features segregated into sets; a 3-tier hierarchy is the list of features segregated into sets which are in turn segregated into groups, and so on. Figures 6.1 and 6.2 contain 1-tier hierarchies and Figure 6.3 contains a 3-tier hierarchy.

Because the calculation of the unit of measurement for \mathbf{D}' is slightly different for different tiered hierarchies, we'll begin with the 1-tier hierarchy in Figure 6.2 and then go to the 3-tier hierarchy in Figure 6.3. In both cases, the goal will be to obtain a scaling unit, \mathbf{x} , that translates the end results of our computations, $\sum \mathbf{i}_c \mathbf{m}_c$, into units of our scale, thus

providing a meaningful \mathbf{D}' . From these examples we will infer a general rule for calculating the scaling unit for any discrepancy analysis, basic or elaborated, single-tiered or multiple-tiered.

For a 1-tier hierarchy (a list of features) the equivalent of a discrepant feature in the basic model is a maximally important, maximally discrepant feature in the elaborated model. Evaluation of such a feature would result in a 3 for its importance, \mathbf{i}_c , and a 3 for the magnitude of its discrepancy, \mathbf{m}_c , making the product, $\mathbf{i}_c\mathbf{m}_c = 3 \times 3 = 3^2 = 9$. That is, the counterpart of the scaling unit, $\mathbf{x} = 1$ for the basic model, is $\mathbf{x} = 3^2 = 9$ for the 1-tier elaborated model in Figure 6.2. Note that the exponent on the 3 is 1 more than the number of tiers.

To make \mathbf{D}' a multiple of 9, divide the sum of the products of importance and magnitude, $\Sigma \mathbf{i}_c\mathbf{m}_c$, by \mathbf{x} , so for a 1-tier hierarchy, $\mathbf{D}' = \Sigma \mathbf{i}_c\mathbf{m}_c / 9$. This yields a scale for which, if every feature in Figure 6.2 were nondiscrepant, \mathbf{D}' would be equal to 0 (zero), and if every feature were maximally important and maximally discrepant, \mathbf{D}' would be equal to \underline{n} (which is the total number of features in the analysis; in this case, 9). In between degrees of importance of the features and in between magnitudes of the features discrepancies will yield in between values of \mathbf{D}' .

Now, turn to the 3-tier hierarchy in Figure 6.3. Once again, the equivalent of a discrepant feature in the basic model is a maximally important, maximally discrepant feature. Evaluation of such a feature in a 3-tier hierarchy would result in 3 for the group to which it belongs, 3 for the set in the group to which it belongs, and 3 for the feature itself, which would result in an entry of $3 \times 3 \times 3 = 3^3 = 27$ in the blank in the importance column, \mathbf{i}_c and 3 in the magnitude column, \mathbf{m}_c . When the entry in the importance column (27) was multiplied by the entry in the magnitude column (3), the product entered in the $\mathbf{i}_c\mathbf{m}_c$ column would be $(3)^3(3) = 3^4$

= 81, which is the scaling unit for the analysis in Figure 6.3. Note that, once again, the exponent on the 3 is 1 more than the number of tiers. Thus the calculation of D' would be $\sum i_c m_c / 81$.

In all of the examples given in this book, measurement of both importance and magnitude are done using scales for which the upper extreme is 3, but other scales are admissible. Depending on your ability to make discriminations, it might make sense, for example, to measure magnitude on an 11 point scale: no discrepancy = 0, and 1 through 10 are discrepancies of increasing magnitude. For importance, you might, for example, use a 5 point scale: 1 = low, 2 = medium low, 3 = medium, 4 = medium high, 5 = high. Or, if discrimination is difficult, perhaps you'd merely use: 1 = low importance and 2 = high importance or 1 = low magnitude and 2 = anything other than low magnitude. The only constraint is that the same scale for importance must be used at every tier of the hierarchy.

General Rule

The general rule for computing the scaling unit for any discrepancy analysis is $x = (\max i)^t (\max m)$, where $\max i$ is the maximum value on the scale you use to measure importance, $\max m$ is the maximum value on the scale you used to measure magnitude, and t is the number of tiers in the hierarchy. If the maximum scale value is the same for both importance and magnitude, the rule devolves to $x = (\max_{i/m})^{t+1}$, which means that the scaling unit is equal to the maximum scale value raised to the $t+1$ power.

Interpreting D'

To interpret an obtained D' , simply draw a scale on a piece of paper, giving it a lower limit of zero and upper limit of n , the number of features in your analysis. Then mark the point on the scale corresponding to the D' obtained from your analysis. An obtained D' of,

say, 7, is equivalent in undesirability to any forecast you might perform, this one or any other, that has 7 highly important features that are wholly discrepant from your desired future. It is still your judgment call about whether that **D'** exceeds your violation threshold, **T**. I don't know about you, but I would regard any forecasted future to be seriously flawed, no matter how many features it had, if as many as 7 of its features were both highly important and wholly discrepant from what I felt they should be (my desired future).

Subject Index

Ambiguity	121
Anger	75
Anticipated Regret	75
Argument theory	194
Ashmos, D.	218
Automatic intervention	78
Barnes, V. E.	203
Beard, D. W.	217
Behavioral Decision Theory	21
Behaviorism	11
Benson, L. III	218
Biases	199
casual	200
informational	201
statistical	202
Bissell, B. L.	219
Bruner, J.	22
Candidate entity	101
Categorization	96, 217
Choice	118
Cognitive errors	199
Cognitive overload	112
Cognitive Revolution, First	11
Cognitive Revolution, Second	12, 22
Compatibility test	216
Complex decisions	139
Connolly, T.	197
Consciousness	32
Conscious Experience	33
Costs	138
Counterfactual reasoning	49
D	101
D'	108, 225
De Bryn, E. E. J.	219
Decision agents	152
Single	152
Collective	154
Decision aids	120, 129
Decision rules	101
Decision sheet	131, 137, 144, 169
Decision simplification	115, 155
Decisions	95

aided	120
rudimentary	95
unaided	119
Deliberate intervention	78
Deontology	198
Desired future	46
Discrepancies	97, 98
Discrepancy Test	98
basic	100
enriched	107
Duchon, D.	218
Dunegan, K. J.	218
Economic Decision Theory	14
Economic Man	14
Emotions	73
Endure	78
Entities	100
Escape	77
Estranged paradigms	26, 27
Etzioni, A.	198
Expected value	17, 18
Explanation theory	193
Fear	75
Feature	107
Importance	107
Magnitude	107
Features	100
comparable	101
defining	101
Fechner, G.	10
Forecasts	45, 52
action	45
extrapolated	45
plausibility	47
Frederickson, J. R.	219
Fundamental puzzle	9
Freud, S.	10
Hastie, R.	185, 193
Heckler, S. E.	218
Heuristics	205
anchoring and adjustment	210
availability	209
representativeness	206
Idea generation	165
Identification	95, 118
Identity assumption	98, 101

Image theory	28, 189
Imagination	48
Implementation	91
Incremental theories	197
Intervene	78
James, W.	10
Jungermann, H.	192
Kahneman, D.	21, 206
Klein, G. A.	186, 192
Knowledge	60
Lee, T. W.	219
Levels of decision making	109
level 1	109
level 2	111
level 3	111
Lindbloom, C. E.	197
Lipshitz, R.	194
Lottery	207
Lundell, J.	188, 217, 219
Marble, T. A.	218
Maximization of expected value	17
Memory	58
Memory, Declarative	60
episodic	60
function	58
procedural	59
retentive	59
semantic	60
structure	58
Mertens, D. P.	218
Mitchell, T. R.	28, 188, 189, 217, 219
Monitoring	92
Montgomery, H.	194
Morgenstern, O.	15
Narrative thinking	23
Narratives	33
coherence	41
consistency	41
cultural	44
current	34
episodic	38
good	38
hierarchy	34
in psychotherapy	39
origins	34
others'	40

plausibility	41
Naylor, E.	218
Negotiations	161
Nelson, K.	219
Non-routine decisions	160
Non-violations	216
Null hypothesis	100
Occam's Razor	100
Organizational culture	153
Organizations	155
Paradigmatic thinking	24
Pascal, B.	15
Pavlov, I	10
Pennington, N.	185, 193
Persuasion	43
Plans	77
incremental	90
master	90
recycled	83
structure	79
Potter, R. E.	217
Progress decisions	218
Puto, C. P.	218
Rational choice theory	16
Recognition	95, 118
Recognition theories	192
Rediker, K. J.	217
Reflexivity	195
Rejection Threshold	97, 101, 216
Research strategies	183
Richmond, S. M.	219
Routine decisions	159
Rules	61
contingent	61
non-contingent	62
normative	65
Scaling unit	133, 225
Scenario theory	192
Screening	217
Search stop	99
Sensitivity analysis	143
Sensory events	9
Simon, H. A.	21, 192
Skinner, B. F.	11
Smith, B.	188, 217, 219
Smith, V.	21

Soros, G.	195
Steller Art Frames	163
Stevens, C. K.	219
Strength of expectation	64, 67
Strom, E.	216
Svenson, O.	194
Tactics	79
sequence	81
sources	82
Theory of Narrative-Based Decision Making	178
Theory versus reality	123
Thresholds	103
Thüring, M.	192
Time horizons	51
Tools	80
Truth	42
conceptual	42
empirical	42
Tversky, A.	21, 206
Uncertainty	121
Universal cognitive abilities	13
Values	67
and meaning	72
cues to	70
origins	69
primary	68
secondary	68
Veto rights	114
Violation threshold	218
Violations	216
Von Neumann, J.	15
Walsh, K. R.	219
Watson, J.	11
Weber, E.	10
Welles, H. G.	42
Wells, O.	42

Reference Index

Alexander, L. D.	93
Barnes, V. E.	213
Basu, K	30
Beach, L. R.	30, 55, 76, 93, 126, 213, 176
Betsch, C.	76, 127
Betsch, T.	76, 127
Bruner, J.	30, 55
Byrne, R. M. J.	76
Capps, L.	56
Connolly, T.	30, 126, 213
Crandall, B. W.	56
Denning, S.	55
Etzioni, A.	213
Gerrig, R. J.	55
Hacking, I.	30
Hastie, R.	56, 57, 214
Herman, D.	30
Hoffman, R. R.	55, 56, 150
Hollnagle, E.	150
Jungermann, H.	55, 57, 93, 94, 213
Kahneman, D.	30, 76, 126, 214
Klein, D. E.	150
Klein, G. A.	55, 56, 150, 213, 214
Lindbloom	214
Lipshitz, R.	56, 214
Lowenstein, G.	30
Lundell, J.	213
Mar, R. A.	56
McAdams, D. P.	56
Meyers, D. G.	76
Miller, G. A.	126
Mintzberg, H.	94
Mitchell, T. R.	213
Mitchell, W. J. T.	56
Montgomery	214
Moon, B. M.	55, 56, 150
Myers, D. G.	126
Ochs, E.	56
Peeters, G.	98, 127
Pennington, N.	56, 57, 214
Pin, C. L.	56
Plessner, H.	76, 127

Polkinghorne, D. E.	30
Quinn, J. B.	94
Ross, K. G.	150
Smith, B.	213
Snowden, D.	56
Soros, G.	214
Strom, E.	126
Thüring, M.	55, 57, 93, 94, 213
Tversky, A.	30, 214