Three Roads to Commitment:
A Trimodal Theory of Decision Making

Marvin S. Cohen

Perceptronics Solutions

4200 Lorcom lane

Arlington, VA 22207

703-577-5678

mcohen@percsolutions.com

and

Raanan Lipshitz

Department of Psychology,

University of Haifa,

Haifa, Israel 31905

raanan@psy.haifa.il

Revised manuscript

17 November 2011

Three Roads to Commitment:
A Trimodal Theory of Decision Making

ABSTRACT

Researchers in different paradigms have concentrated on different narrow slices of the decision making domain, such as *choice* in classical decision theory or *recognition* in naturalistic research. We distill a more inclusive conception of decisions as dynamic, graded commitments to action at different levels of specificity and temporal scope. We identify three prototypical pathways to commitment –*Matching*, *Reassessment*, and *Choice* – which differ in their starting points (i.e., initially active commitments), the questions they use to frame planning (about appropriateness, efficiency, and reliability, respectively), and the normative backing they confer on action. The framework identifies three dimensions of action uncertainty: issues relevant in each mode, types or patterns of uncertainty about those issues, and methods used to handle it. It provides an ecological approach to action uncertainty, as a decision maker's distance (in effort and time) from actionable commitment, and draws some implications for the management of time. The three modes, crossed with different levels of action uncertainty, help organize the field of decision making research and define a space of cognitive possibilities to which decision aids and training should adapt.

Keywords: Decision making, intention, commitment, naturalistic decision making, decision aiding, decision analysis, bounded rationality, cooperation, self-control, confirmation bias, escalation of commitment, critical thinking

INTRODUCTION

Rethinking Decision Making

A decision is conventionally defined as a choice between two or more options, or acts, based on reasoning about the desirability of their consequences (e.g., Jeffrey, 1965/1983). A large amount of research has been formulated and carried out within the framework of this definition, with little or no critical consideration of its implicit assumptions. What we will call a standard story has emerged from this work, addressing *normative* issues (how choices ought to be made) by means of decision theory (Savage, 1972; von Neumann & Morgenstern, 1944; Luce & Raiffa, 1989), *descriptive* issues (how in fact they are made) by behavioral decision making research focused on systematic errors (Kahneman, Slovic & Tversky, 1972: Gilovich, Griffin, & Kahneman, 2002), and *prescriptive* issues (how to help real decision makers make better decisions) by means of decision analysis (Raiffa, 1968; Keeney & Raiffa, 1976: Brown & Paschoud, 2005) and cognitive psychology (von Winterfeldt & Edwards, 1986; Bell, Raiffa, & Tversky,1988).

The story is interesting and important, but its plausibility hinges on the assumptions embedded in its starting point – i.e., *selecting* the *best* of *multiple options* by *deliberating* about *consequences*. We will question almost every piece of this definition. We will also question the more basic presupposition that decision making can be subdivided into *stages* corresponding to what happens before, during, and after a discrete event called *choice*: viz., that *before* this event, options, criteria, and outcomes must be identified, and that only *after* a choice has been made (by application of the criteria to the outcomes of the options) can action be initiated. Our ultimate aim is constructive: to sketch a more empirically valid, theoretically insightful, and pragmatically useful framework, which dispenses with these assumptions.

Choice, when it occurs, is not usually the most important part of the decision process. Simon (1960) observed that managerial decision makers spend a considerable amount of time (i) identifying problems and opportunities and (ii) assessing the situation. They spend *most* of their time in what he referred to as (iii) the *design* of suitable actions or plans. They spend the *least* amount of time in (iv) comparison and selection of options, or choice. Other researchers (March, 2000; Schön, 1983) have studied processes of (v) adjustment, learning, and innovation that occur *after* plans are adopted, often in the midst of action. (i), (ii), (iii), and (v) are neglected, perhaps, because conventional decision theory applies only to choice, guided by an overly narrow view of what counts as normative (Lipshitz & Cohen, 2005; Cohen, 1993).

As Alexander (1979, p. 383) noted three decades ago, a myopic focus on choice leads to a normative paradox:

> If planners, administrators, and managers are being equipped today with tools and techniques for articulating goals and evaluating projects and programs, how should they develop those alternative solutions which they are to evaluate…? *Are there any normative techniques for systematically designing alternative solutions in a given situation?* If there are, how appropriate are they in the light of actual behavior and perceived constraints in real environments? How crucial these questions are is clear when we realize that *the best evaluation techniques can only be applied to those alternatives which are given*. [italics added]

Compelling as Alexander's point is, it does not state the paradox in full generality. He could have posed the same questions not just about option generation but also about identifying fundamental goals, adhering to or abandoning decisions already made, or even more fundamentally, framing a situation in terms of a particular set of objects and concepts. Standard choice methods must make do with given goals, given ways of describing the situation, and given options, even though agents remain capable of thought even as they evaluate and execute actions in a changing and uncertain world. Are there normative principles and methods that apply to *these* processes?

Empirical research in naturalistic environments casts doubt not only on the importance of choice, but on whether, in many cases, it occurs at all. Experienced decision makers often recognize appropriate actions quickly in familiar environments (Klein, 1998); when they are uncertain, they often explore options sequentially rather than concurrently, stopping when they find a satisfactory solution rather than looking for the best (Lindblom, 1959; Mintzberg, 1975; Simon, 1987). Conversely, supposedly distinct processing stages (e.g., situation assessment, goal specification, option generation, option evaluation-and-selection, and action) occur in parallel, or in complex embedded cycles, rather than in a linear series (Schön, 1983; Braybrooke & Lindblom, 1963; Witte, 1972; Mintzberg, Raisinghani, & Théorêt, 1976). As Witte (1972, p. 180) observed, "human beings cannot gather information without in some way simultaneously developing alternatives. They cannot avoid evaluating these alternatives immediately, and in doing this they are forced to a [revisable] decision. This is a package of operations and the succession of these packages over time constitutes the total decision making process."

The standard story attributes this real-world disarray to human capacity limitations, irrationality, or ignorance, to be rectified by decision aids and training. Other researchers suspect that there is an order in naturalistic decision making which the conventional story overlooks or undervalues. Naturalistic Decision Making research (e.g., Klein, Orasanu, Calderwood, & Zsambok, 1993; Hoffman, 2007; Mosier & Fischer, 2011) has rejected the standard story, but has not provided a comprehensive theory to take its place. Yet the search for order across a full range of human

decisions should be of paramount concern to cognitive engineers. Decision support or training may disrupt rather than improve decision making unless it is based on an apt understanding of decision making itself.

Overview

This paper introduces a theoretical framework, Trimodal Decision Making, or TDM. The discussion proceeds in five steps. In the next section, "Decision Making as Change in Commitment," we develop a more inclusive conception of decisions, discuss an empirically testable approach to intentions, and describe some dimensions along which they vary. In the section, "Three Modes of Decision Making," we argue that three distinct *paths* to change in commitment can be usefully distinguished – by starting points, questions that guide attention, deeply divergent justifications that they construct for actions, and different functionalities and priorities. In "Decision Strategies as Paths through the Framework," we use diagrams and examples to show how TDM relates to, and accommodates insights from, other decision research paradigms, including both rational choice theory (the standard story) and naturalistic alternatives proposed by Klein (1993), Beach, (1990), Connolly & Wagner (1988), Orasanu & Fischer (1997), Pennington & Hastie (1993), and Montgomery (1993). The next section defines subjective and objective senses of "Action Uncertainty" in terms of questions, answers, and methods for answering questions that move decision makers toward actionable commitment with varying degrees of efficiency; it also provides a *decision modes* x *methods* taxonomy of decisions and existing decision research paradigms. In the Conclusion, we briefly sketch high-level prescriptive implications for aiding and training. One of the most important tasks of cognitive engineering is to identify the different sources of rationality associated with commitments in particular domains, to identify the most critical obstacles to those commitments, and develop ways of overcoming them.

We do not claim that TDM is the only way to carve up the domain of decision making, or that it says everything there is to say about it. Our goal is nevertheless ambitious: to propose a substantive and insightful (or at least thought-provoking) theoretical framework for use in decision making research and prescription in real-world environments. Theories should be judged by efficiency (i.e., how much they explain relative to how much they assume), and fruitfulness (e.g., novel empirical predictions and practical applications). With regard to efficiency, TDM will hopefully be recommended by a distinctive combination of empirical coverage and simplicity. A small number of concepts centering around an empirically grounded notion of *commitment* is sufficient to generate a richer and more inclusive view of decision making than offered by the standard story, to distinguish three decision making modes along normative and descriptive lines, delineate different types of action uncertainty, accommodate the findings of important existing decision research paradigms, and point to future directions in descriptive research

and prescriptive tool design. With respect to future research, footnotes throughout the paper call attention to some interesting, experimentally testable hypotheses. Understanding the empirical implications of TDM should clarify its meaning (and, for hard-nosed empiricists, show that it *has* meaning), and stimulate empirical research and criticism that should ultimately improve theory. There are many practical applications in cognitive task analysis, and we touch briefly on high-level implications for training and decision aid design in the conclusion. Nevertheless, most of that story is still to be told.

DECISION MAKING AS COMMITMENT CHANGE

If decision making is not choice, then what is it? We concur with Yates, Veinott, & Patalano (2003), whose definition highlights "committing to a certain course of action." More specifically, we regard decisions as *graded commitments of mental, affective or material resources to courses of action*. Decision making includes any *cognitive proces*s that can create, reject, or modify such commitments – regardless of the cognitive resources it may or may not consume in doing so. This definition highlights the invariant output (i.e., increasing, decreasing, or sustaining commitment) while avoiding restrictive assumptions about how it is produced.

When a decision concerns action in the future, the commitment is an *intention* (Brattman, 1987). Intention should be understood very broadly, to include long-term, high level *goals* and *policies*, *plans* to achieve the goals or apply the policies, and specific *tactics* or *actions* to implement the plans. (Thus, an intention may be a constituent of any of the various "images" discussed by Beach, 1990.) Commitments can be unmade as well as made. Therefore, a commitment must be *less* than the physical or even psychological impossibility of doing anything else. At the same time, commitment to a future action must be *more* than a mere gesture (e.g., a sub-vocal "I'm going to be a doctor") or magical mental event (an unobservable construct called "deciding"). Intentions are attributed to ourselves and others based on evidence between these extremes, involving cognitive, physical, and social behavior. Drawing in part on Bratman (1987), we define the intention or commitment to do A (which may be a goal or action at any level of generality and temporal scope) as a set of dispositions: (a) to stop looking for or thinking seriously about alternatives to A unless there is specific reason for reassessing it; (b) to notice, seek out, retain, and think about aspects of the situation that are relevant for A's implementation or success; (c) to specify and plan A in more detail; (d) to take preparatory steps for A (e.g., by allocating resources, creating opportunities for A, rehearsing A, or enlisting others' cooperation in A); (e) to experience negative affect if A is blocked; and (f) to do A at suitable time(s) or place(s). These dispositions – and associated intentions – may vary in strength both among themselves

and across time.[1]

Intentions, because they are dispositions to act in future, can be thought of as implicitly conditional, i.e., an intention to do A *if C*. Both actions A and conditions C may be fuzzy and even somewhat indeterminate (for example, a vague commitment to the value of helping others whenever the opportunity arises), and need not be easily articulated. Both the acts and the conditions for acting may vary in specificity and in temporal scope. For example, the activity of *being a doctor* is significantly more general and slightly more extended than the activity of *practicing a particular medical specialty*. Both of these, in turn, are more general than the action of *using treatment method A*, whose temporal scope is distributed among discrete events whose occurrence is conditioned on specific intervals (e.g., while the doctor is *in charge of a patient with condition X*). This intention, in turn, is more general than the action of *using a customized variant of treatment method A*, which begins and ends within the single short interval of care for a single patient. Specificity and temporal scope need not be negatively related. An intention may include *parts* that are also intentions, e.g., a plan may include a series of treatments for a particular patient contingent on results at each stage, which is both more specific and more temporally extended than a plan for a single treatment. An action may also be an intended *aspect* or feature of another action. For example, Dr. House might intend to minimize unnecessary tests and interventions throughout his medical practice. Intentions such as these can spread either formally or informally, becoming medical protocols or standards to be followed throughout a profession. Cultural traditions and norms evolve over longer periods of time and spread even more widely. In more uncharted and non-traditional regions of practical decision making, however, specificity and scope do tend to trade off: The larger the temporal scope of an intended activity, and the more general or temporally extensive its

[1] This implies empirically testable claims: that commitment induces characteristics (a) through (e), and that, as a result, these characteristics co-occur. These predictions can be tested, for example, by methods developed in research on *prospective memory* (McDaniel & Einstein, 2007; Kliegel, McDaniel, & Einstein, 2009). Prospective memory is defined by Goschke & Kuhl (1996) as the representation of an action to which the decision maker is committed and that plays a causal role in performance of the action. They developed an experimental paradigm in which participants learn two scripts under exactly the same conditions; they are *then* told which one of the two they will have to act on. When tested after intent formation but before execution, memory is better for the script corresponding to the active intention, despite equal exposure, learning incentives, and distractions. Moreover, the advantage disappears when memory for the two scripts is compared after the intended script has been implemented. To test item (b) above, the paradigm could be modified to test recognition of *new* information that is *relevant* to two plans, one intended and one not, rather than verbatim script text. For item (c), the competence and completeness of spelling out plans, e.g., for contingencies, could be compared for two tasks, one intended for execution and one not. To test resistance to revision (a), willingness to change a plan in light of disconfirming information could be compared in the two cases.

conditions of implementation, the less specified the intended action is likely to be. Instead of being determined in one fell swoop, details of actions will be filled in or modified in future cycles of planning or in the midst of execution (Mintzberg, 1994: Schön, 1983).

As illustrated by dispositions (a) through (e) above, intentions not only guide action but provide a *background frame* for subsequent decision cycles. Decision making never starts from scratch. Every decision is framed by a relevantly active subset of the agent's pre-existing commitments (Bratman, 1987; Beach, 1990, 1998; McDaniel & Epstein, 2007), which actively recruit new decision making cycles to shape new intentions or generate immediate action. Existing commitments combine with long-term knowledge and current information to *pose questions* – i.e., they sensitize attention to information and information acquisition opportunities at levels of specificity and temporal scope suited for *their own elaboration, critical evaluation, selection, or execution.* A plan or activity is usually the dynamic product of multiple decision cycles distributed over time, rather than the output of a discrete choice stage. New questions become active as old ones are answered – by filling gaps, resolving problems, or making choices – with respect to commitments that have not yet hardened into irreversibility. External actions occur opportunistically as *needed* to collect information, prepare, plan, persuade, test the waters, take care of other matters, buy time, implement, or if necessary change course, not as endpoints *after* an "option" is "chosen."

A distinction between front and back stages of processing – i.e., "situation assessment" and "decision making," respectively – may seem necessary, because perceptual observations are driven by the external environment and actions appear to occur as a result. Neisser (1976) and Gibson (1979) challenged this frame: actions operate in and on the environment, and new actions respond to their effects in a continuous cycle. Moreover, situation assessment hardly ever occurs in real-world contexts without orienting action. TDM reinforces the point that neither situation assessment nor decision making "comes first" in a fundamentally cyclic process. It also recasts the distinction between them. *Situation assessment* is an activity that can be an intrinsic part of some decisions, combining with previous commitments to directly affect action commitment at any level of specificity and scope – just as *comparison of options* can directly affect action commitment as an intrinsic part of choice, when it occurs. Cognitive processing goes on continuously, with frequent or infrequent inputs caused by environmental changes. Changes in commitment (i.e., *decisions*) are not phases of this processing at all; they may emerge from the stream of cognition at any time, or they may take place continuously under its influence – with no *a priori* or fixed relationship to the topics of cognitive processing (e.g., features of the situations versus features of outcomes). The relationships between the two streams is a subject for empirical investigation, and we discuss three possibilities in the next

section.

In the research literature, *choosing* lingers on as a generic description, even in naturalistic analyses that explicitly acknowledge other processes. (It is sometimes said, for example, that a decision maker "chooses" to perform a *rule-based* response, or "chooses" *no go* over *go*.) TDM enables us to say, more accurately, that *the decision maker commits to an action* and *rejects commitment to an action* – or with even less baggage, that an *action commitment is created, increases, decreases,* or *is rejected.* This is more than linguistic. A broader conception of decisions highlights a crucial question in naturalistic research: *How and why do commitments change?*

THREE MODES OF DECISION MAKING

This conceptualization of decision making prompts recognition of three *maximally simple* ways commitments can change: by *increasing*, *decreasing*, or *focusing* strength – or, in discrete terms, by *adding*, *rejecting*, or *selecting among* constraints on action. Corresponding to these operations, a decision cycle may ask, respectively: *What should I do?* (where the answer must be a degree of commitment noticeably *more* specific or comprehensive than the starting state), *Is this course of action reliable?* (where the answer is to *reduce* the specificity or scope of the starting commitment if it does not meet challenges, otherwise to retain it), or *Which course of action is best?* (where the answer is to *select*, or reallocate preference to, one from the set of alternatives, whose specificity and scope are unaffected). The three operations apply to different initial states of commitment, which we will simplify by subtracting commitments that are unaffected. The first applies before any commitment exists, i.e., *null,* and relative to which increases are occurring; the second applies after commitment is established, i.e., *full*, and relative to which decreases may occur; and the third occurs when the decision maker is committed to a mutually exclusive set, i.e., a *disjunction*, relative to which a selection must occur before action.

We argue that these transitions are more than a formal scheme. By regarding them as distinct prototypes or *modes* of decision making, we gain extensive empirical coverage and richness by relatively parsimonious means. Three lines of argument underpin this claim. First, these three modes are associated with qualitatively distinctive phenomenologies of thinking and acting, i.e., deep differences in the ways they justify actions and in the actions they justify, and these are associated in turn with importantly different functionalities and priorities. This argument is the focus of the present section. Second, they provide a set of building blocks for the description and analysis of a comprehensive range of real world decisions; for example, planning need not be regarded as a distinct mode because it is accounted for a mix of more fundamental processes. As a corollary, they shed light on relationships, strengths, and weaknesses of existing decision research paradigms. The following section, "Decision Strategies as Paths to

Commitment," focuses on this use of the framework, as a parsimonious language for describing both simple and complex real-world decisions. Third, the same three modes apply, at a meta-level, to the management of the processes used to make decisions in all three modes, accounting for the selection of strategies, management of time, and cognitive effort expended. That is the topic of the section on "Action Uncertainty." TDM is not the only way that elements of decisions can be understood. Yet we are not aware of any other scheme that combines these three advantages.

The path to commitment change in each mode can thought of as the construction of an *argument* – to persuade oneself or others, either before or after the fact.[2] We can describe such arguments by reference to Toulmin's (1958; Toulmin, Rieke, & Janik, 1984) dialogue-based model: If a *Claim* (i.e., an intention or action) is challenged, it may be defended by presenting *Grounds*, or reasons, which focus attention on particular parts of the decision problem (Barber, Heath, & O'dean, 2003; Lipshitz, 1993; Shafir, Simonson, & Tversky, 1993). If the relevance of the Grounds to the Claim is then challenged, it may be defended by presenting a *Warrant*, i.e., a general rule or practice, which explains why the Grounds justify the action. Exceptions to this apparent linkage – i.e., possible errors – are reflected in *Rebuttals* to the Claim in unusual cases, where the Grounds hold but the Claim may not. The Warrant in turn may be defended by citing its normative *Backing*, based on some body of data, theory, logic, or intuition. Table 1 summarizes how an ideally reflective decision maker would justify Claims in the three modes – by distinctive Grounds, Warrants, and Backing, subject to different types of Rebuttals.

Insert Table 1 about here

Matching

Matching seeks actions that are obligatory, appropriate, or permissible for a particular person in a particular situation. It starts with *no* foreground commitment at target levels of specificity and scope; existing relevant

[2] Lipshitz (1993) described decisions as *argument-driven actions*, which are justified by reasons and by successfully rebutting objections (i.e., reassessment) (for similar ideas, see Barber, Heath, & O'dean, 2003; Shafir, Simonson, & Tversky, 1993). Cohen et al., (2006) described *critical thinking* as a *challenge-defense dialogue with oneself or others*, consisting of four stages: detection of a problem with a current commitment, selecting a strategy (e.g., whether to continue with a critical dialogue, with whom, and on what topic), an exchange of challenges and defenses regarding the commitment (i.e., reassessment), and a resolution in which the commitment is retained, rejected, or modified.

commitments are a baseline that can be strengthened or added to. Matching prompts the implicit question, *What should a person in my role do in a situation of this kind?* (cf. March, 1993, p. 58). The answers are provided by Warrants, or rules (we will call them *practices*), whose conditions are found to match the situation and decision maker role, and whose actions may then be executed or adopted as intentions. Practices are themselves intentions that are general across situations and persistent in temporal scope. [3] They reflect pre-existing social norms (Bicchieri, 2005; Bicchieri, Jeffrey, & Skyrms, 2009), organizational routines (Levitt & March, 1988), standard operating procedures and best practices (Betsch & Haberstroh, 2004), expert solution patterns in disciplines like medicine or chess (Chase & Simon, 1973), patterns recognized by proficient decision makers (Klein, 1993), scripts for familiar behavior (Shank & Abelson, 1977, 1995), previous cases and precedents (Schank, 1990, 1999; Kolodner, 1993), argumentation schemes (Walton,1996), principles and strategies for conversation (Grice, 1989; Sperber & Wilson, 1995; Sacks, 1995), heuristics (Gigerenzer et al., 1999; Simon & Newell, 1958), competitive or cooperative strategies in social enterprises (Axelrod, 1984), culturally modulated emotional dispositions such as anger, guilt, honor, and love (LeDoux, 1996), personal policies (Bratman, 1999; Beach, 1990), cultural maxims and themes (Lakoff, 1987), religious doctrine (Atran, 2002), shared intentions for joint activity (Gilbert, 1996), moral principles and political ideologies (Etzioni, 1988; Thompson, Ellis, & Waldavsky, 1990), or goal trajectories and high-level values (Beach, 1990). Practices are often associated with different social identities or roles that a single individual might take on, related to kin, work, profession, voluntary association, religion, nationality, region, political ideology, or personal self-image (Etzioni, 1988, 2000).

When a practice (i.e., a Warrant) is activated by a current or anticipated situation (Grounds), it focuses attention on the situation type and the agent's role in it. However, matching is not merely a front end for choice. Multiple mutually exclusive options need not be, and most frequently *are* not, made directly available by a practice or rule. An action generated when the conditions of a practice match a situation and role can be committed to and implemented without stepping outside the practice to generate and compare other possibilities. In short, matching generates intentions and actions by processes other than choice. If the situation and role that match the practice are current, the practice can give rise to immediate action (the Claim). When the situation is anticipated, the practice

[3] Terms like *schema, rule,* or *norm* often fit just as well. We do not, however, equate intentions with internal or external representations of the intended action (although such representations often do play a role in planning). Intentions exist as an interconnected set of *dispositions* ((a) through (e) above) – even if they are merely implicit in neural network connections and facts about the situation (Hutchins, 1995).

may generate new intentions (Claims) that are more detailed, hence, more nearly actionable, than previous intentions. Constraints continue to be added across successful episodes of matching, resulting in commitments to actions that are progressively more specific or more temporally complete, until details and timing are fully settled by execution.

Old paradigms die hard. It might be thought, for example, that matching decisions are, at the least, *go-no go* choices: i.e., the alternative is always the status quo, no commitment. But this is a bridge too far. The supposed *no-go* option cannot be identified by the decision maker, and therefore cannot exist for the decision maker as a choice option, until *after* matching has told her what *go* means. And the *go* intention generated by matching need not be reassessed before being executed. In some cases, e.g., routine actions or responses to immediate danger, not performing the action is literally unthinkable. The logical *possibility* that no-go might have occurred in some counterfactual universe does not raise it to the level of a choice option.

Matching decisions are implicitly justified by reasons of the form, *Do A because it is the appropriate type of action for someone in my role in situations of this type*. The underlying normative rationale (i.e., Backing) for decisions in the matching mode depends on the practice – but they differ significantly in form, function, and origin with no guarantee of mutual consistency. There is no *a priori* overarching "meta-rule" to establish priorities and resolve conflicts, for example, between the reputation of a domain expert, scientific ethics, organizational culture, family responsibilities, religious prohibitions, personal policies of critical thinking, and universal ethical principles (Fiske & Tetlock, 1997). Matching is ubiquitous because of this variety; it is responsible for some of the strongest human commitments, despite its lack of unity. The bare bones Backing that instances of matching share is that *some actions are appropriate or inappropriate according to accepted practices for a person in a particular role in a particular type of situation* (Alexander & Moore, 2008).

Reassessment

In Reassessment the decision maker monitors for or actively probes for problems with an intended action (i.e., Rebuttals to previous arguments), in response to the implicit or explicit question, *Is my course of action reliable?* Grounds for reassessment include signs of trouble that emerge from monitoring or testing, and which support an argument to reject (or at least more intensively reassess) an intention (the Claim). The Warrant for reassessment is the specific association between such signs and lack of reliability. It may be based on an empirical sample of performance that has been observed personally or reported by others, or it may be based on a process of critical thinking (Cohen, Adelman, Bresnick, Freeman, Salas, & Riedel, 2006), mental simulation (Markman, Klein, &

Suhr, 2009; Klein, 1993), inference from theory, or responsiveness to dissent expressed by other informed parties (Walton, 1998; van Eemeren & Grootendorst, 1992). Whether reassessment is based on practice or deliberation, decisions are justified by reasons of the form, *Continue A because it survived challenges*, or *Do not continue [aspects or parts of] A because it failed to survive challenges*. If elements of an intention are rejected, the commitments are *less* specific or *less* comprehensive than they were before reassessment. If a replacement is not waiting in the wings or has not become obvious, what follows is usually a renewal of matching to find or create one.

Reassessment is the backup strategy for commitments made on the basis of incomplete information – as most commitments are in real world environments. Fortunately, it has impeccable normative Backing. The rationale for continued commitment to a practice or plan is that it has been tested – i.e., exposed to the possibility of relevant new information or Rebuttals – and has so far survived; i.e., it *works*. It has been argued along Darwinian lines (e.g., Dennett, 1996) that genuinely new knowledge (i.e., the appearance of design when there is no omniscient designer) can *only* emerge from some version of trial and error, or generate and test. Differential empirical success (which is the basis for reassessment) appears to influence the survival and shape of organisms (Gould, 1977), behavior (e.g., Sutton & Barto, 1998; Skinner, 1953), and ideas (Popper 1994, 1969; Dawkins, 2006), respectively.

Reassessment may begin at any time after the decision maker commits to an intention, as long as commitments to future parts of an on-going activity, or future repetitions of a practice, are not yet irreversible. Commitment excludes baseless reassessment (see criterion (a) above), but when the course of action is a general practice or long-term plan, the need for small improvisations and adjustments may be anticipated. If major difficulties occur or if minor ones accumulate, reassessment kicks into higher gear and the core commitment may be in jeopardy. Reassessment presupposes what Schön (1983) has called *double vision*: A decision maker can both commit to a course of action and simultaneously be aware, at least implicitly, of the *possibility* of problems of any size.

Choice

Choice seeks the most efficient means to desired ends. It begins with a commitment to select one action from a set of two or more mutually exclusive alternatives, framed by the question, *Which of these options is the best means to my ends?* The Warrant for choice is maximization of the expected desirability of future states. The Grounds, and the focus of attention during decision making, are alternative courses of action and their estimated consequences. Using these Grounds, the choice process seeks to reduce the mutually exclusive disjunctive intention (e.g., *to do A or B – but not both*) to a singular intention to perform *one* of the given options (e.g., *to do A*). Selections are justified by reasons of the form, *Do A because it is [a part or aspect of] the most efficient available means to my ends.*

The normative rationale (i.e., Backing) for choice is uniquely systematic and well-studied. Instrumental rationality was first conceived in economics (e.g., Jevons, 1871; Walras, 1874) as what Simon called *objective* (1956) or *substantive* (Simon, 1976) rationality. By contrast, decision theory formalized modern instrumentalist rationality (Ramsay, 1926; De Finetti, 1937; Savage, 1954/1972; Jeffrey, 1992/1965; Luce & Raiffa, 1957) in terms of *subjective* rationality (Simon , 1956), which localizes optimality not in the best possible actual outcomes, but in the *coherence* of beliefs, desires, and behavior (or more precisely, in the "self-evidence" of axioms that define coherence among probabilities, utilities, and choices, respectively). Decision analysis provides a highly developed set of tools for quantifying probabilities and utilities and recommending choices in ways that satisfy these coherence requirements. Proponents of decision theory often make exclusivist normative claims for it (e.g., Binmore, 2009; Baron, 1993; Lindley, 1982). Any sources of action other than instrumental optimality, such as custom, personal policy, or religious prohibition, are flat-out superstitious, and violations of coherence are flat-out irrational (e.g., Kahneman, Slovic, & Tversky, 1982; Gilovich, Griffin, & Kahneman, 2002). By contrast, *bounded rationality* research investigates shortcuts, or heuristics, that are not necessarily "coherent" but produce acceptable results with less effort in the specific contexts to which they are adapted (Gigerenzer & Goldstein, 1996; Goldstein, 2001; Gigerenzer, 2000; Payne, Bettman, & Johnson, 1993).

Rebuttals and Relationships

Unless an intention is executed as soon as it is formed, reasons may arise to reassess it. Table 1 (last two rows) shows that each mode is vulnerable to different types of errors, or Rebuttals, which can motivate reassessment. There are two general types of Rebuttals: Those that are consonant with the mode in which the intention was formed, and those that apply from the perspective of a different mode (which may become active prior to implementation). The first type of rebuttal involves the possibility that key elements of the problem, i.e., Grounds for the mode in question, may have been incorrectly perceived, interpreted, inferred, or predicted. Reassessment of matching focuses on potential misinterpretations of the situation and one's role in it; reassessment of choice focuses on potential errors in predicting or valuing consequences; when reassessment applies to its own Grounds, it focuses on the thoroughness to date of testing, and the representativeness of the testing with respect to a current or anticipated situations. [4]

[4] An experimental test can start with the same three groups described in note 9. After each group has committed to a decision in an experimental problem, they would be presented with disconfirming evidence that either fits their

A second type of rebuttal is relevant when multiple practices (i.e., distinct Warrants) are applicable to the same problem, are deemed appropriate by matching at the same or different times, and produce conflicting Claims. For example, satisficing (Simon, 1987) is a hybrid of modes: Appropriate actions are generated by matching and then reassessed by the desirability of their consequences. The last row of Table 1 shows that rebuttals can be found to Claims of any decision making mode if reassessed from the perspective of any other. In principle, such rebuttals might generate a series of challenges and defenses that pit competing warrants and backing against one another. When they apply in the same contexts, different modes may generate normatively clashing results just as different practices do.[5] In the next three sections, we provide some perspective on these clashes, for two reasons: First, because they have important implications for the prescriptive use of different modes in training and decision aiding; and second, because they highlight the distinctiveness of the three modes as forms of argument.

The Primacy of Matching

Choice and reassessment are *practices* whose applicability is limited by a web of cultural obligations and rights and individual differences, and which may fit some situations badly (Etzioni, 1988). Choice depends on cultural attitudes or personal policies that license calculation of outcomes and deliberate weighing of tradeoffs. Reassessment depends on cultural attitudes or personal policies that license critical reflection, dissent, autonomy, change, or progress. Societies differ in the leeway they afford individuals to negotiate their roles and relationships or choose their own goals and tactics (Douglas & Wildavsky,1982; Douglas, 1999; Thompson, Ellis, & Wildavsky, 1990). Matching itself determines when it is permissible to engage in choice and reassessment in lieu of the default, which is always continued matching. Whenever choice and reassessment processes occur, they are implicitly *embedded* within prior processes of matching.[6]

---

decision mode or does not. Ethically committed participants should be more likely to reassess and reject their initial commitment based on ethical counter-evidence; instrumentally committed participants should be more likely to reassess and reject their initial commitment based on instrumentalist counter-evidence; reassessment oriented participants should respond to information about testing and performance. Similarly, decision aid users will actively reject decision aid recommendations that are generated and justified by an inappropriate decision making mode.

[5] A number of studies have examined how people resolve conflicts among ethical principles, or between ethical principles and instrumentalist incentives. Material inducements to abandon an ethical or identity-based decision may result in anger or outrage (Ginges, Atran, Medin, & Shikaki, 2007), while conflicts among ethical principles is experienced as emotionally stressful and difficult (Hanselmann & Tanner, 2008).

[6] TDM implies that personal and cultural values might influence decision making mode in situations where more than one might apply. For example, drawing speculatively on the elements of Schwartz' (1992; Schwartz & Bilsky, 1987a,b) value model, matching might be more dominant for persons or cultures that score high on security,

An important function of matching is its ability to screen decisions off from instrumentalist rationality and reassessment. Ethical principles (Etzioni, 1988; Schauer, 2009), cooperation in social dilemmas (Etzioni, 1988; Elster, 2000), and self-control (Prelec & Bodner, 2003; Herrnstein & Prelec, 1992; Nozick, 1993) rely on such protection. Abiding by a policy (e.g., to stay on a diet or to behave ethically) or being a particular type of person (e.g., trustworthy, altruistic, or self-controlled) are not easily treated as desirable *outcomes* of choice (Raz, 1990; Prelec & Bodner, 2003). For example, the desirability of an outcome should be independent of the action it is the outcome of (Savage, 1954). It is inexplicable from the perspective of choice that a clear preference among states of affairs (e.g., *four people live and one dies* preferred to *one person lives and four die*) can be reversed depending on whether or not they are the outcomes of actions by an agent in a particular role (e.g., Dr. House allows one patient to die so that four others can receive organ transplants!). How can it be wrong to bring about a *desirable* state of affairs? Moreover, society associates negative utility with the very method (i.e., weighing tradeoffs involved in killing one to save four, cheating in a commercial transaction, or breaking a diet) that is used to make choices (Nozick, 1993, p. 55). Experimental research and naturalistic observation suggests that some tradeoffs are seen as "taboo," and attempts to bargain are counterproductive when "sacred" or "protected" values are at stake (Tetlock et al., 2000; Fiske & Tetlock, 1997; Ginges, Atran, Medin, & Shikaki, 2007). Note that bounded rationality (i.e., limits on capacity and time) does not help. Rules decree that an agent should *not* reassess some commitments, even when it costs little to do so. Moreover, violations of rules for ethics and self-control are not looked upon simply as "errors"; they often evoke guilt in the agent and anger in those who observe the transgression (Fiske & Tetlock, 1997;Gignes et al., 2007; Gibbard, 1990). As Dosi et al. (1998) put it, quoting Kenneth Arrow, "a system of literally maximizing norm-free agents '…would be the end of organized society as we know it' (Arrow 1987, p. 233)." [7]

Matching is the only one of the three modes that can operate self-sufficiently: Reassessment retains or rejects a pre-existing course of action, choice selects among pre-existing options, and both apply pre-existing criteria, while

---

conformity/tradition, benevolence, or universalism. Reassessment might be at a maximum for persons or cultures that put high value on stimulation (i.e., change) or autonomy. The role of Choice might be related to hedonism, achievement, or power. Integrating TDM with a model of values offers a more differentiated prediction about cross-cultural differences in cognition than the familiar research distinction between analytically-oriented Western societies and tradition-oriented non-Western societies (e.g., Nisbett, 2003.).

[7] A testable hypothesis regarding the priority of matching over choice would compare contexts where ethical principles are at stake with contexts where they are not, e.g., in terms of the effect of choice aids whose primary function is to reduce workload. More decision makers will be influenced by the aid in the non-ethical context.

matching generates action (including plans, tactics, and values) in the absence of commitment at desired levels of specificity and scope. Matching is fundamental, then, both because the other two modes are *practices* that must fit a particular situation while continued matching is the default, and because matching is the only non-random source of the raw material for their operation. The alternative in both cases is relatively unguided, i.e., "random," search.

The priority of matching in these two senses does not mean that matching is or should be the initial "phase" in every episode of decision making. First, that would diminish, not enhance matching's role. The contribution of matching is accumulated over lengthy stretches of cultural and individual history; it cannot be squeezed into a "phase" of decision making. Choice and reassessment are the visible tip of an iceberg while matching is its far vaster submerged body. Second, explicit decision making is always *in media res*: Decision makers may therefore find themselves with a commitment they need to reassess, or with a set of options from which they need to select, in environments that license the practice. Unless there is specific reason to reassess more fundamental assumptions (e.g., about the situation, goals, or permitted practices), they need not be concerned with where the initial commitments came from. While matching is "prior" in some respects, there is no *a priori* sequence of modes that applies to every decision problem, either descriptively or prescriptively.

The Flexibility of Reassessment

Reassessment and matching are deeply complementary. While matching seeks to maximize the number of truths, reassessment seeks to minimize the number of errors. Thus, matching seals commitments and reassessment unlocks them. Matching, if taken by itself, assumes that personal and collective experiences embodied in rules are sufficient for the problem at hand. Reassessment assumes the possibility of mistakes and stands ready to learn from them. The rules and practices that underlie matching have largely evolved through the kind of trial and error involved in reassessment – whether biological, cultural, organizational, or individual – rather than by the deliberate comparison and calculation of benefits exemplified in choice (Smith, 2008). Matching and reassessment are both necessary: one to generate commitments based on accumulated knowledge and the other to vett commitments and stimulate improvements, which are then used in future matching.

The more connected a commitment is to other commitments, the more resistant it is to reassessment. Regardless of its origins, an intention becomes more central through continued successful use in the background of decision making, resulting in an expanding web of connections to the new goals, plans, and tactics that it has helped to generate, as well as to other goals, plans, tactics, and beliefs that combined with it to do so. As centrality increases through use, resistance to reassessment must also grow, due to the psychological, social or physical effort required

to make a large number of simultaneous changes. This web of commitments produces the stability necessary for effective planning, coordination, and timely action. Yet reassessment is necessary to avoid traps, adapt to change, resolve disagreement, and exploit new information. Reassessment minimizes loss of existing knowledge by making small adjustments to *peripheral* assumptions; if repairs are not available, too numerous, or cumulatively unconvincing, it expands the search space outward, bringing more *central* assumptions into consideration. The result is *double loop learning* (Argyris & Schön, 1974, 1996; Rudolph, 2003), in which first order difficulties (e.g., unavailability, growing implausibility, or unworkability of peripheral adjustments) motivate the surfacing of more fundamental, central premises from the background frame (Cohen et al., 1996). If a central assumption can be saved only by adopting a large number of *unreliable peripheral assumptions*, the balance of plausibility shifts to increasingly fundamental reassessment. (We will see an example of this below in Figure 4.) [8]

The pattern of double loop learning predicted by TDM is analogous to Kuhn's (1996) contrast between normal science, which fills in the gaps and solves puzzles within a prevailing paradigm, and scientific revolutions, which replace more fundamental assumptions, methods, and criteria of success. It also resembles Gould's (1977) punctuated equilibrium hypothesis for natural selection, according to which long periods of relative stasis are interrupted by periods of very rapid wholesale change.

Reassessment privileges the current plan and looks to its past performance; choice evaluates multiple alternatives on a roughly equal basis and looks only to the future. Not surprisingly, choice proponents view behavior based on reassessment as irrational. Experiments have shown that committed decision makers who experience a negative outcome or disconfirming information (1) are reluctant to change (the *sunk cost fallacy*: Arkes & Ayton, 1999), (2) ignore or distort the negative information (*confirmation bias*: Wason, 1968; Nisbett & Ross, 1980; Poletiek, 2001; Klayman, 1995; Evans, 2007), and (3) may even take actions that deepen commitment (*escalation of commitment*: Staw, 1976; Staw & Ross, 1989). These phenomena violate path independence (Garud & Karnoe 2001): people act differently on the same information depending on past commitments or the order in which the

[8] TDM makes an empirically testable claim regarding *causes* of centrality. An intended action that has been used to frame subsequent planning will become more central than an equivalent intended action that has not. Moreover, this effect will depend on whether the subsequent planning is intended for implementation or merely hypothetical. Centrality can be assessed directly or indirectly. A direct approach is to apply standard network measures of centrality to results of knowledge elicitation: Central elements are on a larger percentage of the paths that connect elements of knowledge. Indirect approaches include hypothesized effects of centrality, such as resistance to reassessment. This can be measured as the quantity of new evidence against a planned action (e.g., indicators that key assumptions about the situation or about outcomes are false) that is sufficient to prompt a change of mind.

information was received. This must be mistaken if there is one right answer to a decision problem.

But the latter assumption, along with the allegation of bias, can itself be mistaken – especially in complex, novel situations where reassessment is most needed. In such cases, reassessment "searches" in a landscape that it must construct as it goes along, without knowledge of the variables and values that will turn out to be relevant, and which is in any case too large or irregular for anything close to exhaustive exploration. Reassessment adapts to such problems by starting to search in a relatively familiar neighborhood of causes, predictions, plans, and outcomes, changing only what looks wrong, and moving to a new neighborhood only when the current one appears beyond repair (Lindblom, 1959; Braybrooke & Lindblom, 1963; Connolly & Wager, 1988). Reassessment's bias in favor of current commitments reflects a reasonable initial judgment about where solutions are likely to be found, hence, it provides an anchor that improves the prospects of generating new knowledge (Politeik, 2001). More pragmatically, it is less disruptive of on-going practice, reduces decision making costs, preserves more existing knowledge, and frees resources for innovation in less risky or disruptive ways. In the context of *non-exhaustive search in an ill-defined search space*, rational explorers may come to different conclusions because they have different starting points. Allegations of irrationality and bias need to be reassessed.

The Non-Universality of Choice

In the standard story, choice is not just the dominant type of decision making but it is the *only* kind, by definition. In fact, conscious choice may itself be the least indispensable decision mode (a better case may be made for unconscious allocation of time, as we discuss later). Its functions can be emulated by serial generation and evaluation of individual actions by matching and reassessment respectively. For example, satisficing, which retains a course of action if it clears aspiration levels on all evaluative dimensions, is a hybrid of matching and reassessment that approximates choice, especially if reassessment focuses on consequences. Satisficing performs virtually as well as optimal choice models when there is a stable set of satisfactory options that do not differ drastically in expected utility (Gigerenzer & Goldstein, 1996). Genuine conscious choice in TDM requires awareness of a pool of concurrently available options that start out on a roughly equal footing with respect to the decision (as opposed to satisficing, and reassessment more generally, where results are highly path-dependent). Decisions tend to be framed as choice when (a) multiple options are prepackaged for comparison in the environment (e.g., a web site that provides a table of products by features), (b) cultural or organizational constraints require justification by comparison with other options (Tetlock, 1996), or (c) goals compete with one another, which is to say that salient options trade off advantages and disadvantages or compete for scarce resources, and some approximation to

optimality is desirable. Choice stretches matching by allowing the generation of multiple concurrent candidates before committing. It stretches reassessment by permitting relative rather than absolute evaluation of concurrent options. Tradeoffs may be dealt with formally, by aggregating diverse evaluative dimensions into a single measure, or informally – e.g., by trying out criteria in order of importance to see if they discriminate one option from all the rest (*one-reason strategies*, Gigerenzer & Selten, 2001), or successively eliminating options that fall below thresholds on criteria in order of importance (*elimination by aspects*, Tversky, 1972). The results can be considered valid only with respect to the options and criteria utilized in the model or strategy, which were generated by matching – and conditional on the assumptions inevitably embedded in the model or strategy, which are subject to reassessment (Lipshitz & Cohen, 2005; Klein, 2001). The price of validity, even in this restricted sense, is more shallow sampling of information regarding each possibility, in place of the more detailed and comprehensive understanding often afforded by serial matching and reassessment.

Summary Regarding Decision Modes

Our order of priority for the decision making modes – matching, then reassessment, then choice –coincides, we believe, with their frequencies of occurrence in naturalistic contexts: Matching is pervasive, reassessment less frequent, and conscious choice relatively rare. They have compensating strengths and weaknesses, corresponding to different ways search can go wrong. Matching, is slow to change, hence, it is most vulnerable to rapid changes or instability in the search landscape. Choice tends to spread efforts out over multiple possibilities and to readily abandon past commitments, hence, it may be misled by limited samples in a stable landscape. Reassessment counterbalances the conservatism of matching and the shallowness of choice, but can become trapped in local optima on a hilly, irregular surface (Levinthal, 1997, 1998, 2008). Offsetting this risk, reassessment may occur on multiple levels, looking out for opportunities or risks on a larger scale and temporal scope while simultaneously fine-tuning and correcting the current plan – a process that Etzioni (1988) refers to as *mixed scanning*.

According to TDM, an episode of decision making is part of a continuous cycle of thinking while acting and acting while thinking. Figure 1 shows successful transitions from the initial commitment state for each decision mode to a single action – generated by matching, selected by choice, or verified by reassessment. To summarize our view of these, matching starts with a *null* commitment at relevant levels of specificity and scope, attends to the *present* situation type and the decision maker's role, and *adds* a commitment to *appropriate* action. Reassessment starts with a *singular* commitment, attends evidence associated with its performance in similar situations, and either *retains* it as *reliable* or *rejects* it as *unreliable*. Choice starts with *multiple* mutually exclusive possibilities, attends to

the desirability and likelihood of their *future* consequences, and *selects* the one that is *best*. The three modes rely on distinctive normative foundations, based on accepted practice (matching), variation and selective retention (reassessment), or coherence between actions, consequences, and desires (choice).[9]

Insert Figure 1 about here

A newly generated, selected, or validated commitment has three possible fates, as shown by arrows emerging from *One COA* in the upper left quadrant of Figure 1. It may be acted on immediately; it may be tested and reassessed immediately; or it may remain as an intention until implementation – in the meantime framing further decision making and perhaps also subjected to reassessment. Moreover, in all cases of *increased* commitment, elements of the background frame that contributed to it become more connected, hence, more central to the web of beliefs, goals, and actions; as a result, they are themselves more resistant to subsequent reassessment.

The traditional story of decision making conceals the fact that many, and probably most, decisions concentrate on the present situation (matching) or past performance (reassessment) rather than future consequences (choice). Given the systematic differences among the modes, and the functional priority of matching and reassessment, it is a mistake to dismiss these decisions as poor approximations to choice. The standard story fails to explain where goals, situation frames, and options come from (matching); why choice and reassessment are sometimes not permitted (matching); and how learning and innovation take place (reassessment). In the next section, we look at relationships between TDM and a variety of naturalistic models.

DECISION STRATEGIES AS PATHS TO COMMITMENT

Cognitive strategies are well-travelled sets of paths through the framework outlined in Figure 1, often spanning multiple cycles in multiple modes and at different levels of specificity and scope. In this section, we exercise the

[9] A testable implication of TDM is that the distinctive features of each decision mode will co-vary. For example, an experiment might present the same decision problem to three groups of participants. The groups may be formed on the basis of prior proclivities for matching, choice, or reassessment, respectively; alternatively, they may be formed by priming with material emphasizing ethics, profit maximization, or flexibility and experimentation. The matching group would be expected to attend more to the situation and their role in it, and to justify decisions in terms of appropriateness. The choice group should attend more to options and their outcomes, and justify decisions in terms of finding the best. The reassessment group should attend to recent testing of a course of action and ways to improve it, and justify decisions in terms of what works.

framework by exploring some of these paths. Our purpose is to illustrate how TDM-style diagrams can be used to describe actual decision episodes, to clarify TDM by comparing it to other decision research paradigms – in particular, classical rational choice and prominent alternatives – and finally, to motivate a conception of *action uncertainty* as distance in effort and time to actionable commitment (which we turn to in the section following this).

Beach & Connolly (2005; also see Lipshitz, 1993a) mention four categories of research that pose alternatives to classical rational choice: *recognition* theories such as Klein's (1993) Recognition-Primed Decision Making(RPD); *narrative* and *argument* theories such as Pennington & Hastie's (1993) Explanation-based Decision Making and Montgomery's (1993) Dominance Structuring; *incremental* theories such as Connolly's (1988) Decision Cycles and hedge clipping strategy; and *moral/ethical* theories such as Etzioni's (1988) moral-affective factors. TDM has been influenced by all of these. Beach's (1990) Image theory also attempts a synthesis, and not coincidentally, bears the strongest resemblance to TDM. Figure 2 diagrams rational choice, RPD, hedge clipping, and Image theory as successful trajectories through the TDM framework, hopefully shedding light on their relationships to TDM and to one another. Another important synthesis emerged from Orasanu & Fischer's (1997) empirical investigation of aviation decisions; examples from their taxonomy are presented in Figure 3. Explanation-based reasoning and dominance structuring, can also be represented as paths through the TDM framework, as shown in Figure 4. A fuller account in all these cases, but especially the last, relies on our discussion of action uncertainty.

Insert Figure 2 about here

Decision Paradigms

Rational choice

Figure 2 (upper left) is based on the traditional normative theory of choice, which begins with a set of mutually exclusive options (1a), selects the one with the most desirable expected outcomes (1b), and implements it (2). Our intent is not to present a strawman, but to represent the idealized normative core of rational choice theory., *In practice* (e.g., decision analysis) choice is far less linear and more multifaceted (von Winterfeldt & Edwards, 1986; Bell, Raiffa, & Tversky,1988). The rational choice model, however, theoretically combines all subjective beliefs and preferences of the decision maker, so there is no explicit role for matching to filter out inappropriate information; it uses theoretically optimal algorithms, so there is no explicit role for reassessment to verify model implications. The resulting strategy boils down to an isolated moment of choice, which fits Connolly & Wagner's (1988)

characterization of a one-shot all-or-nothing *tree felling* strategy. The notion of *intentions* and *commitment* are also superfluous in idealized normative rational choice, because it ignores cognitive capacity limitations and regards proximity to time of implementation as in principle irrelevant (McClennen, 1990). Any choices, no matter how momentous, can be recalculated at any time with an invariant result unless there is genuinely unanticipated new information .

Recognition-Primed Decision Making (RPD)

Figure 2 (upper right) represents RPD (Klein, 1993, 1998) by three TDM decision cycles. The first cycle (1a) matches the situation against a learned repertoire of familiar patterns, which specify cues, expectancies, actions, and goals, and retrieves the most typical action (1b), which becomes the decision maker's intention (2a). In relatively novel situations the decision maker reassesses it by mentally simulating its implementation (2a); if it survives this test (2b), the intention (2a) is implemented (3). In more familiar contexts, the first retrieved action may be executed immediately (1a,1b,3), without mental simulation.

Hedge clipping

Connolly & Wagner's (1988) hedge clipping strategy involves a series of small commitments that target specific problems in the status quo rather than large decisions that target a globally desirable future state of affairs. In the former case, errors are easily corrected; in the latter case, they may be disastrous. Figure 2 (lower left) represents hedge clipping by means of four TDM decision cycles, with the possibility of continuing iterations. The first cycle (1a, 1b) involves either recognizing or simply guessing an appropriate action, which becomes the decision maker's initial intention (2a). A criterion of appropriateness is that the action address a specific problem with the *status quo*, but allow for corrections if things go wrong. Instead of reassessing the action by mental simulation (as in RPD), the decision maker implements it (2a, 2b) and observes the outcome (2c). The success or failure of the action is then evaluated (3a). If it was a success (3b), the decision maker continues by further small steps in the same direction (cycle 4, which iterates cycle 2).

Image theory

In Figure 2 (lower right), which depicts processes specified in Image theory, six TDM cycles combine elements of all the previously discussed strategies. The first and second cycles are components of what Beach (1990) calls an *adoption* decision: Options must clear a *compatibility* test that resembles matching (1a), to screen out alternatives that conflict with the decision maker's principles, goals, or on-going plans (1b). If more than one option survives (2a), the best is chosen by a *profitability* test (2b). (Choice, however, is decidedly less important than initial

screening based on compatibility; most decisions involve only a single option.) Once a course of action has been arrived at, two kinds of *progress* decisions – or reassessments – may occur, also based on compatibility with principles, goals, or on-going plans. The first resembles mental simulation: it occurs before implementation and projects the action forward in time (3a) to identify and mitigate potential problems (3b). The second progress decision occurs after implementation (4a, 4b, 4c) and resembles hedge clipping. Based on observation of the action's impact (4d), reassessment (5a) determines the extent to which it had the intended effects (5b). If it did, implementation continues (6). Image theory accommodates moral and ethical factors as filtering criteria during compatibility testing, while TDM relates them to matching (to generate intentions or action) and reassessment (to filter them). Both theories emphasize the active subset of a decision maker's knowledge, values, goals, and plans that frame every decision.

Decision Effort Taxonomy

Rather than focusing on a single strategy in commercial airline pilot decision making, Orasanu & Fischer (1997; also, Orasanu, 1993, 1994, 1998) emphasized the "flexible application of a varied repertoire of strategies" by effective crews. They usefully distinguished among six decision types that they observed in simulator performance or identified in transcripts of incident self-reports: *go/no-go*, *condition-action*, *choice*, *scheduling*, *procedural management*, and *creative problem solving*. They also introduced a Decision Process model, with two dimensions corresponding to situation assessment and decision making, respectively: Cues regarding the nature of a decision problem may be clear or ambiguous; the situation may afford a single response (condition-action, go-no go), multiple responses (choice, scheduling), or no response (procedural management, creative problem solving). Orasanu & Fischer propose that a taxonomy based on these dimensions might predict the cognitive work required by the six decision types. In this section and in Figure 3, we will show what five of the six decision types in their taxonomy look like from the TDM perspective. The discussion has two purposes: to illustrate how empirically observed decision patterns in a specific domain (e.g., aviation) can be represented by particular combinations of cycles in the TDM framework, and to expose some theoretical differences.

Insert Figure 3 about here

Condition-action decisions fit Klein's RPD model, in particular, a single, simple *matching* cycle, as illustrated in Figure 2: a pattern of cues is recognized (1a) and an appropriate response, A, is retrieved (1b) and implemented

(3). In *Go-No Go* decisions, on the other hand, "an action is in progress [e.g., takeoff or landing], a pattern is recognized that signals dangers [e.g., insufficient thrust during takeoff, dangerous cross-winds on approach to landing], and the response is pre-set: stop the action" (1993). From the perspective of TDM, go-no go decisions presuppose a *prior* process, such as the condition-action matching cycles just discussed, in which the decision maker has committed to an action and initiated implementation; we include these initial cycles in Figure 3 (upper left, 1a,b,2,a,b,c). The actual go-no go decision is more properly identified with the subsequent *reassessment* process, in which decision makers actively monitor for signs of trouble (3a), reject future components of the intention when it appears unworkable (3b), and halt on-going activity (3c). TDM therefore agrees with Orasanu & Fischer's distinction between condition-action and go-no go decision types. In fact, TDM distinguishes them more fundamentally, and we think more clearly, than Orasanu & Fischer's (1997) model. (According to the latter, both apply rules that prescribe "a single available response," but go-no go involves "bifurcation" of both conditions and actions, or a "binary option." [10])

Orasanu & Fischer (1997) illustrate *choice* by a diversion decision enacted in a simulated flight scenario. After crews have aborted a landing due to crosswinds (e.g., the no-go decision modeled in Figure 3, upper left), they are back to matching (3b). They may now reject the intention to land at the original destination (a no-go decision at a higher level of generality and scope than rejecting a single attempt) because weather is not improving. After these two negative reassessment cycles, the normal action is to fly to airport B, the alternate listed on the flight plan, which is recognized as appropriate by RPD (this is where Figure 3, lower left, picks up, at 4a,b). However, crews reassess this alternate (5a) and reject it (5b), based on a short runway and bad weather in combination with an aircraft hydraulic failure. They may now look up another potential landing site and evaluate it in the same way, finding that it has other problems of its own (e.g., no passenger handling facilities). This cycle of matching and negative reassessment continues until crew members realize that no landing site within range meets all requirements. (To simplify, we will assume that the original destination A and alternate B are the only possibilities, yet both have

[10] For TDM, *No go* is not an option or response with the same status as *Go*; it is simply the rejection of an existing commitment to *Go*, leaving a state of null commitment. A replacement action (e.g., *going around* after rejecting a landing, or *heading to another airport* after rejecting the original destination) is often necessary, but it is not supplied by the go-no go decision, because no go is not *equivalent* to any specific alternative action. A replacement must be supplied by matching (via automatic recognition or more extensive deliberation). Notice that if *No go* were admitted as an option, then *condition-action* decisions would also have to be regarded as *Go-no go* "binary" decision between doing the action and not doing it. The distinction between condition-action and go-no go decisions collapses; more importantly, it becomes harder to sustain a distinction between rule-based decisions and choices.

been found faulty.) At this point, crews initiate a process of simultaneous comparison, or choice, to make a selection (6a,b), which they then implement (7). Orasanu & Fischer suggest, and TDM agrees, that genuine choice does not begin until 6a,b, when there is a pool of concurrently available options that trade off in their advantages and disadvantages. The diversion example illustrates an important type of transition: Conditions favorable for choice can be produced by serial generation (via matching) and (tentative) rejection by reassessment for divergent reasons.

These three decision types – condition-action, go-no go, and choice – resemble TDM's matching, reassessment, and choice, respectively. TDM decision modes are more inclusive because they place no restrictions on specificity and scope of the actions generated, evaluated, or selected from (of course, TDM decision modes are not meant exclusively for the aviation domain); each mode can be executed under varying degrees of action uncertainty. Orasanu & Fischer require, at least for rule-based decisions, that both conditions and actions be prescribed in detail. This difference has implications for the way the remaining elements of their taxonomy are dealt with. In particular, *procedural management* involves the use of standard emergency procedures (e.g., "get down fast") in high risk but poorly understood situations. According to Orasanu & Fischer, "These decisions look like condition-action rules but lack prespecified eliciting conditions. The response also is generalized…" TDM does not see this as a difference in mode. Procedural management involves a series of *matching* cycles that progressively transform a vague, general practice into actionable commitments for a particular situation (illustrated schematically in Figure 3, middle column). The first cycle (1a,b) matches the specifics of the current situation (e.g., airframe vibration, unknown cause) to the broad class of conditions where the practice applies (ambiguous, high risk), and produces a general intention (2a), e.g., to "make things safe." Another cycle of matching, prompted by active reassessment, tries to fill in the details (2b). The pilot now makes the causal connection between vibration and aircraft speed (3a), and generates an appropriate action (3b) – "reduce cruise speed" – which can then be implemented (4a,b,c). Additional reassessment (5a) and additional matching to flesh out details (5b) may occur after implementation has begun. For TDM, the difference between condition-action decisions and procedural management is the number of cycles needed to fill in the blanks and produce actionable commitment in a new situation, i.e., the amount of *action uncertainty* (distance in effort and time from commitment) at the start of the process. Even in apparently simple condition-action decisions, the details of the action are not fully settled until implementation is complete.[11]

[11] This explains an apparent contradiction in the way action availability (in the decision effort model) and initial commitment state (in TDM) are described. According to Orasanu & Fischer, a single response is afforded by the situation in rule-based decisions; in procedural management, *no* response is afforded because it must be created by

Scheduling (Figure 3, right column) is necessary when several time and resource-consuming tasks must be accomplished within a restricted window of time with limited resources. Although the tasks themselves are afforded by standard practices, according to Orasanu & Fischer, the actual implementation of them in combination with one another and under time pressure is unfamiliar; there are no rules or standard procedures for determining which are most important, when to do them, and by whom. From TDM's point of view, both scheduling and procedural management require multiple decision cycles, In procedural management, these cycles focus on particularizing the rule to a specific situation; in scheduling, the cycles focus on extending the temporal scope of the plan. Another important difference is that procedural management (as we interpret it) permits commitments to be fleshed out in the course of implementation, as in hedge clipping (Figure 3, middle column). Scheduling requires planning *before* implementation to resolve potential conflicts, e.g., in resource demands and temporal relationships, as in tree felling. Figure 3 (right column) illustrates scheduling as a mix of matching and choice decisions accomplished over a series of cycles that must be completed before action begins. In this example, the initial cycle of matching (1a) turns up more than one candidate for the first time slot (1b, broken line), leading to choice (2a). Once a selection is made (2b), it becomes a part of the plan (3a). Continued processing (3b), e.g., for the next time slot, produces another cycle of matching (4a), which this time lead directly to a single option (4b), adding B as the next element (5a). The plan is now to perform A followed by B, i.e., A—B. Reassessment of this plan (5a) discovers, however, that they cannot both be performed without advance preparatory effort, leading to plan rejection and a return to matching (e.g., to insert a preparatory action before A, or to drop either A or B or both). Reassessment can revisit and revise earlier commitments at any time in this process, in order to resolve resource conflicts, satisfy sequential constraints, or meet time deadlines. This process continues until the plan is complete; only then does implementation begin.

When Things Go Wrong

The diagrams in Figure 1 and Figure 2 did not show what happens when decision makers encounter *obstacles* in

---

the decision maker. TDM describes the initial commitment state in *both* cases as null commitment, because they both require at least one decision cycle to generate an action. TDM thus separates action uncertainty, which varies in quantity, from decision mode. A similar difference emerges in the characterization of choice. According to Orasanu & Fischer, multiple options are afforded because no rule *determines* the tradeoff that should be made among competing goals. But such a rule could be constructed, based on boundedly rational choice heuristics or decision analysis plus an airline or FAA specified importance ranking; in the case of boundedly rational heuristics, such a rule might be applied very quickly and intuitively. TDM would still regard the decision as a choice because the *initial* state of commitment was disjunctive and tradeoffs are present, even if the rule makes the decision easy.

arriving at as single course of action. Additional cycles are required if, in RPD or hedge clipping, reassessment uncovers problems in mental simulation or practice, respectively. In Image theory, initial compatibility testing (i.e., matching) may produce ties that are resolved by choice; progress testing (i.e., reassessment) may uncover problems that cause an on-going action to be rejected or modified. To describe these possibilities, TDM uses more complex trajectories. Figure 3 illustrates cases in which (a) reassessment rejects an *on-going activity*, as in hedge clipping and progress testing (upper left), (b) reassessment rejects a *candidate action* based on mental simulation, as in RPD (lower left), and (c) choice occurs because more than one action is deemed appropriate by matching, as in compatibility testing ties (upper right). Other decision patterns identified in Figure 3 are not a addressed explicitly by those models: successive rounds of matching (d) to specify a general intention (middle column) or (e) to extend temporal scope (right column), (f) choice among imperfect options that arise from successive rejections (lower left), and (g) resolution of conflict by reassessing and revising intentions (right column). Figure 3 also indicates that two uncertainty types – *lack of information* and *conflicting arguments* –guide processing during these transitions, which will be discussed at greater length below. TDM provides a rich framework for classifying specific decision making patterns.

In this section, we investigate even more complex trajectories and some of the processing that goes on within them, which is prompted by different types of uncertainty. What makes these cases interesting is that proficient decision makers do not reject their original conclusions at the first signs of trouble and immediately circle back to matching. Nor do they simply ignore problems, clinging to the original conclusions no matter what, and continue with pointless reassessment. What they do is look for plausible *explanations* or *fixes*. Elements that were previously in the background may be introduced into mental models, providing finer discrimination of situation types, roles, actions, or outcomes in a path-dependent search for a plausible overall account. These tactics may postpone a change in commitment, but they lay groundwork for changes should they become necessary in the future. Narrative and dominance structuring arguments demand this additional complexity in their representation.

Narrative-based Decision Making

Pennington and Hastie (1993) describe decision making by jurors as a two-step process: story construction, to organize and interpret the multitude of interdependent facts in testimony, and a comparison of the best story with the legal conditions of the alleged crime (i.e., a rule). The main components of a story episode, according to Pennington and Hastie, are relationships among initiating events (which elicit) goals (which motivate) actions (which result in) consequences. Story patterns of this kind are general schemata for understanding human intentions based on “folk

psychology" (Malle, 2004). Cohen et al. (1996, 1998, 2001) generalized this account. They found that when experienced decision makers are directly responsible for actions (e.g., military commanders), they use a story telling strategy that *integrates* situation assessment and mental simulation. According to the Recognition / Metacognition theory, the story makes sense of observations by explaining those that have occurred and predicting new ones (e.g., disease C caused observed symptoms S1 and S2; if we test for another symptom of C, S3, we will find it as well) and justifies actions by showing how they will succeed given those explanations (treatment A is appropriate for disease C, and will be followed by observable patient recovery).

According to Pennington & Hastie, stories enable jurors to identify gaps in the explanatory structure – where additional information is needed or, in its absence, where assumptions might have to be made (e.g, test for symptom S3, or just assume it is there). The Recognition / Metacognition theory generalizes this aspect of story-telling as well. Proficient decision makers identify and correct not only gaps in stories, but also *conflicts* (e.g., S3 is tested for but not observed; A is applied but the patient does not recover) and *unreliable assumptions* (e.g., some patients do not respond to treatment A even though they have condition C). They attempt to correct these problems by information collection when it is feasible (e.g., testing first to see if a patient is the type for whom treatment A works), or by adjusting assumptions when it is not (e.g., the doctor assumes A will work until she finds that it doesn't).

Insert Figure 4 about here

Figure 4 diagrams an example of the story-building strategy. It shows a temporal sequence of TDM cycles and how the story changes within and between them.[12] In Figure 4, fictional Dr. House, M.D., matches a patient's

[12] The example is based on Rudolph's (2003) experimental research with anesthesiologists. Rudolph studied error-correction performance in a difficult medical scenario, in which initial diagnoses and treatment decisions were invariably incorrect. She found a strong relationship between problem solving strategies and eventual clinical success. One group, whom Rudolph describes as *fixated*, changed treatments after a failure but never changed diagnoses – no matter how many attempts had been made to treat the same ailment (*single loop* learning only). Another group utilized an approach she called *vagabonding*, in which they changed diagnoses after a single failed treatment attempt (single loop and double loop learning at the same rate). In contrast, the most effective problem solvers did *not* abandon a diagnosis after a single treatment failure. They explained away initial treatment failures as idiosyncratic patient reactions or implementation error. They *did* change diagnoses after multiple failures – when the accumulation of such assumptions had become less plausible. They combined rapid single-loop learning about

symptoms against learned medical practice (1a), and retrieves and instantiates an appropriate causal schema (1b). The schema provides a diagnosis that explains the observed symptoms, condition C, and depicts a causal relationship between treatment A and recovery from condition C. He implements treatment A (2a,b,c) and, as in hedge-clipping, monitors the results (3a). A habit of low intensity reassessment is what Schön (1983) referred to as *mindfulness*. If the treatment is working as expected, House lets it continue. In our example, as in Rudolph's (2003) experiment, the patient fails to respond (3b), prompting reassessment to shift into higher gear. However, reassessment does not directly target the diagnosis, C, or even the treatment, A, but more peripheral assumptions (4). Dr. House asks, in effect: H*ow could it happen that (i) a patient has condition C, (ii) A is the appropriate treatment for condition C, (iii) the patient receives treatment A, but (iii) the patient does not recover?* Premises (i), (ii), and (iii) are in decreasing order of centrality in the mental model. This process of assumption-based reasoning (ABR) turns up a speculative explanation, P1: bad implementation of the treatment (5a), i.e., the failure of A is explained away by assuming that premise (iii) is false. Dr. House makes inquiries in order to verify P1 (5b,c) and determines that it is not correct (6a).

After P1 is falsified, Dr. House is no longer inclined to explain away A's failure, but he is also not yet prepared to question the diagnosis C. ABR resumes (6b) with the same question: H*ow could it happen that (i) a patient has condition C, (ii) A is the appropriate treatment for condition C, (iii) the patient receives treatment A, but (iii) the patient does not recover?* This time the explanation rejects premise (ii), appealing to individual variation in patient response (P2). Matching shows that B is an alternative treatment for C, which should work if P2 is correct (7b). Dr. House now applies treatment B (8a) and monitors the results (8b,c). Unfortunately, he finds that treatment B also fails (9b). The accumulation of failed adjustments (represented as *overall implausibility*) finally persuades Dr. House to challenge a more fundamental assumption, premise (i), that the patient has condition C (10). He will now return to matching to find a new course of action, based on a new explanation for the original symptoms as well as what has been learned from the treatment itself. Schön (1983) referred to such cycles of active questioning and explaining as "reflection in action."

Any part of an evolving justification (in this case in the form of a story) may be elaborated or modified in the light of failed expectations. In this short example, the focus of explanation shifts from peripheral assumptions about

---

treatment-oriented action with a *slower* process of double-loop reflection about their central understanding of the problem – a result that is consistent with the Recognition / Metacognition model.

implementation of A (P1), to the action commitment (A replaced by B and P2), and finally all the way to the diagnosis (C). There is no logical error in explaining (away) unexpected results by peripheral adjustments. Where changes are made is a matter of judgment based on experience. In TDM, this increasingly central questioning is driven by overall implausibility of the account to which the diagnosis C is central, which increases each time a problem is explained away and then influences the level at which subsequent explanations are attempted. [13]

Dominance Structuring and Argument-based Decision Making

Montgomery (1993) describes a choice strategy called dominance structuring, which bears a striking resemblance to the story-telling strategy in matching. The decision maker starts by intuitively selecting a particular option, and reassesses it by attempting a justification. Problems with the favored option are dealt with by adjusting peripheral assumptions: for example, decision makers take a goal on which the option does poorly (or a rival does well) and dismiss it as unimportant, combine it with other goals into a more comprehensive objective (on which the favored option comes out best), or reassess scores. (Story telling may even be enlisted as part of dominance structuring, to explain why the favored option has better outcomes than its rivals.) The aim is to construct a potent argument for the initial course of action, i.e., one in which that option is as good as or better than all other options with respect to all goals. As in story building, the exercise may be more than mere rationalization. Montgomery (1993) cites evidence that if a justification cannot be plausibly constructed, decision makers tend to abandon their initial intuitive selection. The best option, at the end of this process, is the one whose rationalization is most convincing (Tetlock, 1996) – i.e., involves the least revision of initial intuitions about values and goals.

There are many kinds of action arguments other than stories and dominance structures, with different organizing principles that lead to different details in the handling of uncertainty. For example, *sketches* play the role in some fields, e.g., architectural design, that stories do in others (Cross, 2006, 2007; Schön, 1983). According to Cross (2006, p. 116), "Sketching is tied in very closely with features of design cognition, such as the generation and exploration of tentative solution concepts, the identification of what needs to be known about the developing concept, and especially the recognition of emergent features and properties… Concepts drafted in design sketches are there to be criticized, not admired; and they are part of the activity of discovery, of exploration, that is the activity of designing." In a typical design task, according to Cross (2006, p. 114), a "solution conjecture," i.e., a

[13] TDM predicts that more experienced or expert decision makers will behave like Rudolph's (2003) most effective group (note 12): They will not reject a plausible action commitment at the first sign of trouble, but will do so when problems accumulate. Other decision makers may reject commitments either too soon, or too late (or never).

general course of action, appears early. "The nature of the problem can only be found by examining it through proposed solutions…" rather than by prior analysis (p. 35). A solution concept is needed to induce order in the data and place at least tentative boundaries around what will be considered relevant.

Different elements appear to play the central role in different mental models of the decision problem: explanatory hypotheses in stories (which are frequently used in *matching* decisions), high level values in dominance structuring (which is a *choice* strategy), and a solution concept in design sketches (which makes frequent use of *reassessment*). These differences have important implications for processing. TDM predicts that central elements in knowledge structures will change over temporally longer reasoning cycles – they will be generated earlier and dropped later – than more peripheral elements, because they bring more order to observations and commitments, and their rejection brings more disruption. Elements of a mental model may be generated all at once in familiar, intuitive problem situations, as demonstrated in some connectionist architectures (Cohen, 2011). But if they appear over time in some order, that order is likely to reflect their relative centrality in the current decision representation, which in turn is influenced by decision mode, not by their content as situation factors, actions, or values.[14]

How these processes work is important because it determines whether stories, dominance structures, or sketches provide a genuine test of initial intuitions for proficient decision makers, rather than a simple rationalization (Cohen, et al., 1996, 1998, 2001). Effective reasoning demands that decision makers keep track of the overall plausibility of the justification – which is decremented by the causal assumptions, criteria adjustments, or design kluges necessary to make it work. In story-based decision making, reassessment generally begins with peripheral assumptions about the situation, but if the story demands too much of a stretch, it may end with new explanatory hypotheses. Similarly, dominance structuring generally begins with criteria, but if no course of action can be plausibly justified, more skilled decision makers may generate new courses of action, or re-examine more fundamental goals or values (Levi, 1986). Skilled designers may ultimately modify their solution concept in what Cross describes as "a [parallel] search

[14] These hypotheses about *effects* of centrality complement the hypotheses about *causes* of centrality discussed in note 8. There is already evidence that successive difficulties with a commitment tend to produce increasingly central reassessments. Cohen and colleagues (Leddo et al., 1987; Freeman & Cohen, 1995) found this effect with a devil's advocate type decision and training aid. In a decision exercise, the trainee is told that an infallible crystal ball indicates that an assessment is false or a planned action will not produce expected results. The crystal ball clouds up when he tries to find out why, so the trainee has to explain it himself. When the trainee produces an explanation, the crystal ball clears up enough to declare that the explanation is wrong; the trainee is asked to provide another way the assessment could be wrong or the planned action could fail; and so on. The result is a series of explanations that tends to probe more and more deeply into underlying causal assumptions about the situation. Klein's (1998) *pre-mortem* technique is a variant based on this method.

for a matching problem-solution pair, rather than a propositional argument from problem to solution." These issues are addressed more formally in the discussion of assumption-based reasoning in the next section.

ACTION UNCERTAINTY

Uncertainty plays a central role in both the standard story and in alternative models (Smithson, 1989); the broken lines and ovals in Figure 3 and Figure 4 illustrate some of the problems that uncertainty poses in real-world decision making. Statistical approaches generally utilize measures derived from probability (e.g., *resolution* is lowest when probabilities are furthest from zero or one; *entropy* is maximum for *n* possibilities when the chance of each is *1/n*). Probabilities, however, have an explicit application only to *choice*, where uncertainty about outcomes figures in calculations of the optimal course of action. Even in choice, probabilities do not reflect the amount of *knowledge* or *lack of knowledge* underlying the specification of problem elements: courses of action, values, or roles, the events and outcomes to which probabilities are applied, or the probabilities themselves (Langlois & Cosgel, 1992; Cohen, Schum, Freeling, & Chinnis, 1984; Cohen, Laskey, Chinnis, Ulvila, 1986). Moreover, reducing uncertainty to a single quantitative dimension overlooks qualitative differences in uncertainty types and in the methods enlisted for handling them (Cohen et al, 1996; Lipshitz & Strauss, 1997). The kind of uncertainty we are interested in is *pragmatic* (i.e., relevant to difficulty in arriving at commitment), *differentiated* (i.e., provides qualitative guidance for decision making strategies), and *general* (i.e., not tied exclusively to choice). To distinguish it from other senses, we will refer to this concept as *action uncertainty*.

Affordance for Commitment

Uncertainty in the standard story (i.e., probability) has both subjective and objective senses (Gillies, 2000). If subjective, it is a property of the decision maker (e.g., personal betting odds in Bayesian statistics); if objective, it is a property of the environment (e.g., symmetries or relative frequencies in a sample space). In either case, uncertainty has a particular locus (in the decision maker's head or in the environment, respectively), describes a state that exists at a particular time, and has a value that can change only as the decision maker's knowledge or the physical arrangement of the environment, respectively, changes. We intend to challenge this conception: it cannot satisfy the pragmatic condition on action uncertainty, that it correlate with the difficulty of decision making. In its place, TDM supplies an *ecological* concept, which (i) includes both the decision maker and the environment, and (ii) captures prospective outcomes of decision making over time extending from the present. An example will illustrate.

Suppose Dr. House intends to administer a lengthy treatment process, A, to a particular patient, with many steps

contingent on results of previous steps. Consider the following scenarios. (i) There is a standard medical protocol for A which Dr. House will apply, with no foreseeable need for changes. He doesn't remember the details, but can simply ask the head nurse before implementing each step. (ii) There is no standard medical protocol for A, and every case is somewhat different. But Dr. House can rely on his medical knowledge to make good decisions at each stage. This is a somewhat more difficult strategy that resembles hedge clipping (Figure 2, lower left; Connolly & Wagner, 1988) and procedural management (Figure 3, middle column; Orasanu & Fischer, 1997). (iii) There is no standard medical protocol for A. Because there are numerous possible complications (e.g., some medications, if used early, may have side effects that will preclude treatments that may become necessary later), Dr. House will have to work his strategy out in careful detail before he begins. This laborious planning process resembles tree-felling (Connolly & Wagner, 1988) and scheduling (Figure 3, right column; Orasanu & Fischer, 1997).

If we equated action uncertainty for treatment process A with the amount of missing problem-specific information in Dr. House's head, or his current willingness to bet on how details of A will ultimately be settled, we would see almost exactly the same *high* degree of uncertainty at the start of all three scenarios! Yet useful distinctions can and should be made: There is virtually no lack of assurance or conviction in scenario (i), because the information about A is accessible from known external sources. Note that action uncertainty is not a matter of time until the solution is arrived at. In scenario (i), it would not change appreciably if Dr. House had to walk 50 miles to ask the head nurse – as long as cognitive effort (e.g., measured in decision time addressed to primary and secondary questions) is not significantly affected. There is more pragmatic uncertainty in scenario (ii), because Dr. House will have to recognize and respond to unique combinations of events in the midst of treatment – consuming more decision cycles than he would if applying a standard approach. Nevertheless, it might be counterproductive in scenario (ii) for Dr. House to work out a detailed plan in advance, since the environment will provide information to which he will be able to react at the appropriate time with less cognitive effort. By contrast, scenario (iii) involves a great deal of pragmatic uncertainty about A, because of the many issues Dr. House will have to anticipate and think about while constructing a plan. The key point is not the *number* of issues, but the cumulative difficulty of identifying and resolving them. If, at some later point in scenario (iii), Dr. House has worked out the entire plan except for one potentially fatal drug interaction which he sees no way to resolve, his pragmatic uncertainty *at that point* – with a nearly complete but unacceptable plan – may be greater than at the beginning of scenarios (i) and (ii), because of what still lies ahead.

TDM defines *objective* action uncertainty as lack of *affordance for commitment* – or more explicitly, *the*

cognitive effort or processing time required by an agent to answer the primary question associated with a decision mode. This concept is *ecological*, because it spans features of both the decision maker and the environment.[15] Different cognitive strategies are made objectively possible (i.e., afforded) by informational features of the environment in conjunction with the decision maker's relevant cultural, organizational, and personal knowledge and capacities for perception, recollection, reasoning, and interacting with others. Action uncertainty is not an instantaneous time slice of the agent's mental contents, or any other instantaneous measure regarding the decision maker (e.g., momentary betting odds) or of the environment (e.g., availability of the head nurse). First, information that reduces action uncertainty may be in the decision maker's head (e.g., scenario ii), but it need not be (e.g., scenario i). It may be accessed just as readily when stored in the environment – e.g., standard procedures for handling an emergency may be looked up and then utilized in procedural management (Orasanu & Fischer, 1997). Second, action uncertainty is not subjective in the sense of being *up to* the decision maker. It is an objective feature of the decision problem. (It is relative to a decision maker, however, because it can vary for persons with different skills and knowledge facing the same problem.) Third, action uncertainty at a particular moment is a function of what occurs beyond that moment; it refers to future decision making cycles that must occur during the interval – however long or short – before commitment is final and action occurs. Ideally it declines as information is acquired or generated, over fractions of a second or months – until it reaches a vanishing point after complete implementation. Thus, TDM's concept of action uncertainty satisfies part of the pragmatic condition: it is defined so that it corresponds to difficulty. The requirement for an unanswered primary question completes the pragmatic condition, by ensuring that difficulty pertains to decision making rather than idle curiosity. The fact that primary questions frame the effort within a decision mode extends the concept beyond choice, and thus ensures generality.

The differentiation condition for action uncertainty is addressed by two further elements: *secondary questions* and *uncertainty types* defined in terms of their status. If the decision maker has some familiarity with the decision setting, the primary questions interact with the current frame and with the environment to generate secondary questions, which direct attention to specific parts of the decision problem (features of the situation and decision maker role in matching, features of options and the desirability of their consequences in choice, and relevant

[15] We use *affordance* in what we take to be analogous to Gibson's (1979) original sense: "physical possibilities for action" (e.g., stair climbing) construed as a specific objective relationship between features of the environment (e.g., the height of a step) and an agent's physical capacity (e.g., his stride length) (Chemero, 200). TDM extends the idea to include strategies for making and modifying commitment.

indicators of the success or failure of the current course of action in reassessment). Primary and secondary questions associated with each decision mode thus determine the *issues* whose uncertainty is most important -- i.e., the constituents of the argument that the decision maker is constructing to support her decision.

Uncertainty *types* are perceived problems in the current arguments that prevent either posing secondary questions or answering them. *Inadequate frame* applies when it is unclear what secondary questions to ask, hence, what factors are most important in the current problem (if questions *are* posed, they are not well understood, i.e., it is unclear what the *possible* answers might be; if the possible answers are known, it is not clear what their implications would be for the primary issue). A second uncertainty type pertains to the existence of answers. *Gaps in arguments* occur when meaningful, specific questions have been posed but not yet answered. Other uncertainty types, further along in the decision cycle, are relevant only after answers have been given. *Conflicting arguments* occur when questions have been answered, but in inconsistent ways; i.e., some arguments are perceived as confirming and others as disconfirming key claims. *Unreliable arguments* occur when an answer is based on changeable assumptions, i.e., the argument for a claim is seen as vulnerable to potential rebuttal. A final uncertainty type, further along in the cycle, applies only after multiple questions have been answered. Unreliability may be assessed one argument at a time. However, as assumption-based explanations accumulate, or if a decision making strategy is not generally trustworthy, proficient decision makers may react to the *overall implausibility* of the resulting story, model, or plan. It may seem unreliable *as a whole* even if the arguments it depends on are individually acceptable.[16]

[16] To identify uncertainty types, we draw from the authors' research on decision making by real world practitioners. Lipshitz (1997; Lipshitz, & Strauss, 1997) studied retrospective narrative descriptions of over 100 decisions by military officers, coding the primary reason given for the decision's difficulty and the first tactic used to address it (Lipshitz & Strauss, 1997). The RAWFS heuristic (Lipshitz & Strauss, 1997) identifies three types of uncertainty (i.e., lack of understanding, lack of information, and conflict) and six uncertainty handling tactics (i.e., Reduce, Assumption-based reasoning, Weigh pros and cons, Forestall, and Suppress), which give the heuristic its name. The model was subsequently tested in more detail in the domain of incident management by firefighting unit commanders (Lipshitz, Omodei, & Wearing, in press).

Similarly, Cohen and his colleagues conducted retrospective critical incident interviews and experiments with Navy and Army officers (Cohen et al., 1996, 1998, 2001) and commercial airline pilots (Cohen, Adelman, & Thompson, 2000; Freeman, Cohen, & Thompson, 1998). They coded multiple types of uncertainty successively alluded to in a single incident time line and the methods used for handling them. According to the Recognition/Metacognition model (Cohen et al., 1996, 1998, 2001; Cohen, 2011), When stakes are high and time is available, proficient decision makers probe stories for qualitatively different types of uncertainty (i.e., gaps in arguments, conflicting arguments, and unreliable assumptions) and correct them by tactics that shift attention to potential causes of the

The subjective counterpart of objective affordance is *perceived affordance for commitment*. It is a based on the decision maker's experience with the objective affordances of methods she has attempted to use, combined with her use of secondary questions and identification of known uncertainty types in the current decision problem. Uncertainty types have two complementary effects on decision making: First, they are *perceived cognitive obstacles to commitment*, associated with increased cognitive effort in answering the primary question posed by a decision mode (this is consistent with Lipshitz & Strauss's (1997) definition of uncertainty as *a doubt that blocks or postpones commitment*). While action uncertainty is the objective difficulty of the road to commitment taken as a whole, types of uncertainty are the washed out and blocked segments encountered along the way. In addition, they are signposts or markers that guide the allocation of cognitive effort by the decision maker to specific issues and uncertainty handling methods.

These definitions enable us to ask certain questions and provide a framework for the answers: First, what obstacles to commitment do practitioners perceive in real-world decision settings, and what do they do about them? Second, are their perceptions accurate, and do they use cognitive effort efficiently? The two kinds of questions correspond to a subjective and an objective dimension, respectively, of action uncertainty. Naturalistic researchers have studied the first question, viz., perceived qualitative distinctions among types of uncertainty and among tactics used to handle them (including the authors: Lipshitz & Strauss, 1997; Cohen et al., 1996, 1998; see also, Orasanu & Fischer, 1997; Klein et al., 2007). But they have paid less attention to the *quantity* of uncertainty and the difficulty of answering the question posed by a decision mode.

Methods for Handling Action Uncertainty

Figure 5 depicts the components of TDM's concept of action uncertainty. We have already discussed two key ingredients: *questions* whose answers may change commitment, including at least one active primary question associated with a decision mode; and *types of uncertainty*, which are defined by the current status of secondary questions and answers. Figure 5 adds a third component: *methods* for arriving at answers, which can be classified by the *efficiency* with which they enable a decision maker to achieve actionable commitment. Figure 5 (bottom) models what happens *within* a decision cycle in support of any decision mode. The solid line (Figure 5 center) ascending from *Method* shows the cycle's success in arriving at a singular action commitment. Ovals representing uncertainty

---

problem (e.g., either potential new evidence or existing evidence) and by either retrieving relevant information from long-term memory, collecting it externally, or adopting and revising either peripheral or core assumptions.

types and the broken lines emerging from them show where obstacles may be encountered, entailing additional effort, and in some cases, enlisting a different decision making mode.

Insert Figure 5 about here

Strong and Weak Methods: The Relative Prominence of Matching and Search-Reassessment

Agents respond to decision problems with strategies like those depicted in Figure 2, Figure 3, and Figure 4. Within those strategies, they may employ more specific tactics to obtain needed information, e.g., search memory, make external observations, ask someone, or guess. TDM groups decision making strategies and the tactics they use into classes of *methods* (rounded rectangles in the lower part of Figure 5) that are hypothesized to behave in similar ways based on their *efficiency*. For this purpose, we borrow a distinction from research on problem solving (e.g., Patel & Groen, 1991): *Strong* methods produce reliable answers in relatively short amounts of time, because they are constrained by and adapted to familiar problems and tasks. *Weak* methods take longer than strong methods to achieve comparable levels of reliability, because customized strategies are not available. This difference reflects the balance between matching versus search and reassessment in processing. Strong methods are fast because matching primes a familiar practice to guide processing in the relevant decision mode; e.g., story schemas generate questions to guide matching, boundedly rational heuristics generate questions to guide choice; inspection protocols and checklists generate questions to guide reassessment. In each case, the questions have known possible answers that define a set of implicit branching paths to action commitment. Weak methods are slower because the process is not tailored to the domain. They are used in response to an *Inadequate frame*, which fails to provide questions, possible answers, and/or a path from answers to action; hence, relatively unguided search followed by reassessment must pick up the slack.

Within each category, we distinguish relatively fast and slow versions. In the fastest strong method, *adaptive intuition*, answers are known as soon as secondary questions are posed; commitment occurs so quickly that it corresponds effectively to *No uncertainty*. *Adaptive deliberation* involves more active scanning, interpretation, explanation, and prediction to find answers. It responds to relatively well-defined *gaps in arguments*, as shown in Figure 5. Because the domain is unfamiliar, weak methods consist of relatively general cognitive tools, strategies, or mechanisms that apply across domains or problems. The more efficient weak method, *top-down analysis*, is guided by secondary questions, but they are based on general-purpose templates (*What are the options? What are the*

possible outcomes of each option? How likely and how desirable is each outcome?) and lack known possible answers. *Bottom-up search* falls back on non-specific questions (e.g., *Why are you doing that? What happens next?...*), reflecting open-ended *search* for information, examples, advice, or solutions, or ground-level acquisition of basic concepts, facts, and practices, without advance knowledge of where the process is going. Top-down analytic approaches are more efficient than bottom-up methods if their assumptions are satisfied, because they leverage abstract theoretical knowledge, rules, frameworks, or techniques to constrain search and build understanding of an unfamiliar situation.

Objective affordance for commitment, i.e., action uncertainty, is determined by the efficiency of the methods afforded to a particular decision maker in a particular decision setting. What determines a method's efficiency – i.e., the cognitive effort it requires to answer the primary question associated with a decision mode – is the relative role of matching versus search and reassessment at a second-order level, i.e., in the determination of *process*. This varies, as shown in Figure 5, from adaptive intuition (questions primed by matching and immediately answered; no search), to adaptive deliberation (process and *possible* answers primed by matching; limited search), to top-down analysis (process provided by matching, but no advance knowledge of possible answers; exhaustive search of possibilities), to bottom-up search (matching provides little constraint; search space must be created while it is explored).

A *Modes x Methods* Taxonomy of Decision Strategies

Decision strategies can now be represented in more detail, as sets of paths through the more complex framework depicted in Figure 5 – comprising primary questions and iterated cycles of posing and answering secondary questions, by methods that respond to different types of uncertainty. Both strong and weak methods may apply in any decision mode. Table 2 cites examples of research paradigms and strategies corresponding to each combination of decision mode (matching, choice, and reassessment) and uncertainty handling method (*Strong, Weak*), in both fast and slow versions. For example, with strong intuitive methods, the primed process may be matching (e.g., committing to the action recognized as appropriate in the situation; Klein, 1993), choice (e.g., an intuitive strategy of taking the option recognized as most familiar; Gigerenzer & Selten, 2001), or reassessment (e.g., intuitively deciding to abort takeoff based on a recognized pattern of cues; Orasanu & Fischer, 1997).. Similarly, with strong deliberative methods, the primed process may be matching (e.g., taking the action associated with the most plausible story; Cohen et al., 1996), choice (e.g., showing that one option dominates all the others; Montgomery, 1993), or reassessment (e.g., mentally simulating implementation of a course of action; Klein, 1993). Top down weak methods include legal, ideological, and scientific reasoning for matching, decision analysis for

choice, and scenario-based planning for reassessment. Finally, bottom up weak methods include associative learning for matching, reinforcement learning for choice, and novice problem solving for reassessment.

Insert Table 2. about here

Boundaries are not sharp, and many if not all of the paradigms spill over into other categories. However, the taxonomy– *decision mode* x *method* x *fast/slow* – provides some explanatory and prescriptive leverage, as we discuss below. To the extent that membership in different categories reflects empirical differences in the shape of efficiency curves roughly like the ones in Figure 6 below, TDM predicts that strategies in the same class will behave in roughly similar ways – in relation to different types of uncertainty, in competition with other methods, and in control strategies that manage time in decision making.

Assumption-based Reasoning

Assumption-based reasoning supplements strong and weak methods when answers are not fully satisfactory (lower right, Figure 5). Answers may remain incomplete because the information needed to fill perceived *gaps in arguments* is either not available or too costly to obtain, i.e., *lack of information* (Lipshitz & Strauss, 1997). In such cases, decision makers often fall back on assumptions (Pennington & Hastie, 1993; Lipshitz & Strauss, 1997; Cohen, 1996, 1998, 2001). The likelihood that such assumptions will come into conflict with observations, inferences, prior beliefs, or other assumptions is minimized, whenever possible, by filling gaps with relatively peripheral default assumptions about how things normally happen.

A second potential problem is apparent inconsistency of answers. As shown in Figure 5, *conflicting arguments* can be handled in three ways : (a) By continuing with a strong or weak method, including either more information collection or more careful reasoning, or both. (b) By framing the problem as one of *choice* – i.e., weighing pros and cons among well-defined options (Lipshitz & Straus, 1997), using boundedly rational heuristics (a strong method) or more formal analysis (a weak method). (c) By framing conflict as *an anomaly to be explained*. In this case, conflict is regarded as a symptom of a problem in understanding, and is taken as an opportunity to improve knowledge by *reassessing* the arguments that led to the conflict. For example, when conflict occurs between expectations and observations during reassessment (e.g., treatment A failed to work in Figure 4), decision makers find plausible tentative explanations by adjusting peripheral assumptions (e.g., the treatment may not have been effectively implemented) rather than rejecting the action commitment. Explaining (away) conflict can be part of a legitimate

strategy for detecting errors in the absence of solid information.

Nevertheless, there is no free lunch. A problem arises from the use of ABR itself to fill gaps or explain conflicts. To the extent that an argument has been elaborated or bolstered by questionable assumptions rather than reliable information, it counts, to some degree, as an *unreliable argument*. Unreliable arguments may be addressed, if at all, in two ways: (a) By using a strong or weak method, hoping to confirm or disconfirm the assumption either by more information or more careful reasoning, or both. (b) By using ABR, in hopes of finding a more plausible or robust explanation. But there is limited time for such efforts, and assumptions are never completely eliminated.

The greater the number of assumptions adjusted or adopted in order to make the case for commitment, and the less plausible or more central they are, the less convincing the story becomes. Assumptions may be individually plausible and peripheral, yet an accumulation of small deviations from expectations raises a red flag. Figure 5 shows that unreliable arguments add to o*verall implausibility*, i.e., a rough cumulative sum of deviations to date from "normalcy" in an evolving action justification. A second source of overall implausibility is perceived unreliability of the method used to answer questions. According to our definitions, weak methods are less reliable than strong methods, at least over short periods of time. (We can think of this as a generalization about the number of assumptions typically required. Top-down methods depend on abstract analytical assumptions; additional assumptions are required to link theory and observations. Bottom-up methods depend on very large, representative samples to discern interdependencies. Strong methods use less strenuous assumptions, e.g., about the similarity of the current problem to the situations in which the methods evolved.)

Overall implausibility has two regulative effects, which may operate in parallel: (a) When new conflicts occur, overall implausibility influences the level at which ABR seeks explanations. We saw earlier the widening scope of reassessment in Figure 4, as Dr. House advances from peripheral adjustments (P1) when treatment a first fails, to rejection of treatment A when it fails again, to rejection of the underlying diagnosis (C) after the failure of treatment B. Although the more central revisions in such a sequence are *prima facie* more disruptive, they become more plausible after the most likely peripheral explanations have been ruled out; also, if multiple peripheral assumptions are required, central revisions can reduce *overall* implausibility by explaining conflict in a simpler, less ad hoc way (Cohen et al., 1996). A primary function of ABR in TDM is to identify an appropriate level of specificity and scope for adjustment, to achieve a plausible overall situation picture and plan (see note 14). (b) The second regulative function of overall implausibility is its influence on the likelihood and intensity with which conclusions are *reassessed*. For example, it may determine whether or not a decision maker mentally simulates the effects of an

action before implementing it, discusses it with others, or closely monitors results during execution.

ABR, and overall implausibility in particular, can function as a bridge between strong and weak methods, in both directions. For example, a solution based on a weak method (e.g., textbook learning, creative use of previous cases and analogies, or a complicated, assumption-laden analysis) may be sufficiently implausible overall that it is treated as a tentative hypothesis, i.e., an assumption to be reassessed in practice. Over time, it becomes adapted to specific problems, gains credibility, and improves in efficiency, i.e., it evolves into a new strong method. Conversely, the product of a strong method may grow too high in overall implausibility, with no obvious way to fix it by adjusting assumptions (it may be unclear what core assumptions are at fault, or what the alternative core assumptions would be). If time is available and stakes are high, this may prompt the decision maker to revisit first principles by means of a weak method.

The uncertainty handling model in the lower part of Figure 5 zooms in on processing that might occur within a single decision cycle in any decision mode. It describes an inner loop of cognitive processing that can precipitate, or continuously drive, changes in action commitment – which are themselves registered in the outer loop corresponding to decision cycles. Figure 3 and Figure 4 mark some of the more well-travelled paths. Figure 3 shows where *lack of information* and *conflicting arguments* play a role in mode transitions. Figure 4 describes the role of *unreliable arguments*, *conflicting arguments*, and *overall unreliability* in storytelling. In principle, inner loop processing can be prompted by any mode and prompt transition to any other mode. In particular, *conflicting arguments* may originate from matching when more than one response is appropriate (Figure 2, lower right; Figure 3, upper right), from reassessment when courses of action are serially rejected for different reasons (Figure 3, left column), or from choice itself, when criteria fail to eliminate all but one option; in each case, the conflict may be formulated as a (new) choice problem or explained by ABR. Similarly, lack of information may occur because matching fails to generate a solution, reassessment rejects the current course of action (Figure 3, upper left and lower right), reassessment finds gaps in a plan (Figure 3, middle and right columns), or choice rejects all options; in each case, lack of information may prompt matching to discover additional actions. Unreliable assumptions may be identified about the situation in matching (Figure 4, middle), about past performance in reassessment, or about criteria and goals in choice; in each case, the resulting overall implausibility may lead to more intensive reassessment of outputs, or produce shifts in specificity and scope of subsequent matching to find alternative explanations, additional criteria, or new goals.

Time and Strategy Management

Tactics for Control

TDM identifies three control tactics that decision makers can use to manage time: *Forestall*, *Suppress*, and *Buy time* (shown as Actions in Figure 1 and Figure 5). Forestalling mitigates the consequences of errors, e.g., by planning for contingencies, improving readiness and resilience, and avoiding irreversible commitment (Lipshitz & Strauss, 1997). The reduction in risk enables decisions to be made more rapidly. One such tactic, for example, is *incremental decision making* (Lindblom, 1959; Braybrooke & Lindblom, 1963), i.e., making small commitments that permit errors to be easily corrected via continuously adaptive hedge clipping (illustrated in Figure 2 and Figure 4 above; Connolly & Wagner, 1988). The second strategy, Suppression, cuts decision making short despite residual uncertainty; the decision maker acts on whatever the current best solution is, whether it is a hunch, a gamble, a thoughtful plan, or denies uncertainty altogether (Lipshitz & Strauss, 1997). The third strategy, Buying time (Orasanu & Fischer, 1997), reduces the costs of delay rather than the costs of errors (in contrast to Forestalling); hence, unlike Suppression, it extends the time available for decision making rather than cutting decision making short to fit the time. All three of these strategies become more relevant at lower levels of objective affordance, i.e., decision problems that demand more time and effort to achieve given levels of accuracy. An important project for TDM, and the goal of this section, is to shed light on these time management decisions and, more generally, to make the notion of decision making efficiency more precise.

Time management cannot invariably occur by conscious decision making, on pain of an infinite regress of deliberation about deliberation. One mechanism that undoubtedly plays a role is *matching* to situational and role constraints, such as fixed deadlines, quality standards, or customary practices. But well-defined external constraints are not always present. Another proposed mechanism is a stopping rule built into boundedly rational decision strategies (e.g., Gigerenzer & Selden, 2001; Payne, Bettman, & Johnson, 1993); for example, stopping when the first satisfactory option is found, or when a cue discriminates one option from all the rest. But this merely pushes the problem back to strategy selection. It seems clear that multiple methods are often afforded that vary in the amounts of time required (Payne, Bettman, & Johnson, 1993; note, in Figure 5, that *lack of information*, *unreliable arguments*, and *conflicting arguments* can be dealt with by strong or weak methods or by assumption-based reasoning; *conflicting arguments* can also be dealt with by framing the problem as choice). Perhaps more importantly, it is not so clear that individual methods have built-in rigid stopping rules. The strategies identified by bounded rationality researchers, like *satisficing* and *take the best*, are idealizations (Goldstein, 2001). As Simon

himself (1983) points out, satisficing rarely actually stops when the *first* adequate option is found. If finding acceptable options is too easy, decision makers continue to explore options beyond what they originally considered "good enough," with new evaluative criteria or raised aspiration levels. *Weak methods* in particular, may be pursued for highly variable periods of time. An analogous fixed stopping rule proposal, by Klein et al. (2007), although couched in terms of knowledge representations (i.e., "data-frame congruence") rather than strategy completion, appears to raise similar issues of inflexibility: it is flexible only if knowledge structures have no fixed boundaries, i.e., can be elaborated and expanded in the course of decision making; but in that case, it does not provide a "stopping rule" (Cohen, 2011). Rigid stopping rules are especially implausible for experts and proficient decision makers. Recent studies show that proficient decision makers learn to adapt the time they take for decision making to the costs of errors and the costs of time (Cohen, 2011; Khoo and Mosier, 2008).[17]

TDM predicts that if only one method is afforded for decision making, *and* if there is a built-in or externally imposed stop rule, time management will occur by matching. When those conditions fail, TDM suggests that an automatic *choice* process sometimes plays a role – both in strategy selection and in regulating the time taken to make the decision. It accepts as inputs what we earlier referred to as *perceived affordances for commitment*. In the next section, we look at how subjective affordance emerges from objective affordance by experience.

Efficiency of Methods

Uncertainty handling methods vary in objective efficiency, i.e., in the reliability gains they deliver for given investments of effort or time. Figure 6A and Figure 6B clarifies the concept of efficiency by plotting time or effort devoted to a decision on the horizontal axis, and the resulting chance or degree of success in the relevant decision mode on the vertical axis. Each figure notionally compares four methods: weak *vs* strong, and fast *vs* slow variants of each. In both charts, strong methods arrive at maximum reliability more quickly than weak methods from the same starting point. Adaptive intuitive methods are faster than adaptive deliberative methods; and top-down weak methods are faster than bottom-up weak methods.

In addition, there is an important qualitative difference between the curves for strong methods, which are

[17] For example, in a simulated weather and fuel related diversion decision, experienced commercial airline pilots took more time when it was available and less time when it was not; in the same problem, less experienced pilots took the same amount of time in both conditions (Cohen, 2011). Orasanu & Fischer (1997) report that high performing crews took less time than low-performing crews in go-no go decisions, but took more time in diversion decisions (i.e., choice).

negatively accelerated, and the curves for weak methods, which are S-shaped. Strong methods show decreasing returns at relevant time scales as deliberation continues. In strong methods, the information known to be most relevant or easiest to obtain is likely to be acquired early, because they are guided by specific secondary questions; information acquired later is likely to have less impact. By contrast, returns initially *increase* over time at relevant scales when weak methods are used, as decision makers learn how to make decisions in the relevant domain. For example, top-down weak methods require analytical modeling, in which elements and relationships are identified and perhaps quantified. No results are available until the process is complete and results are computed. Bottom-up methods require learning new concepts and relationships in a new domain – effort that does not pay off until a critical mass of useful knowledge is obtained. In both cases, the primary value of early phases is to find out what questions to ask; benefits are realized after relevant factors are identified and information about them is obtained. With repeated use of the same method in the same problem domain , what is learned may be reapplied on subsequent occasions, gradually increasing the "strength" of the method.

Insert Figure 6 about here

The two parts of Figure 6 reflect another qualitative difference, between situations in which the available methods can deliver similar levels of success after sufficient investment of time (Figure 6A), versus situations in which they cannot (Figure 6B). Figure 6A provides an unambiguous ordering of methods in terms of efficiency, which is unaffected by the amount of time invested or the target level of success. In Figure 6B, by contrast, strong methods can be surpassed by weak methods, and fast versions by slow versions, over long periods of time. In this extreme case, the most efficient methods at short time intervals are the least efficient over longer periods of time, and we can rank order methods by efficiency only relative to a fixed decision deadline or target accuracy level.[18]

The efficiency curves for assumption-based reasoning may resemble those for strong methods when peripheral

[18] Hypotheses about the shapes of affordance functions can be experimentally tested by posing decision problems, constraining participants to respond within experimentally varied response windows, and plotting accuracy (e.g., appropriateness of the response, expected utility of consequences, or performance record) against time spent making the decision. TDM predicts that negatively accelerated curves (i) are more likely to be observed in familiar problems, and (ii) are characteristic of methods that do not transfer readily to different problems or domains (i.e., they represent strong methods); and that S-shaped curves (i) are more likely to be observed in unfamiliar problems, and (ii) are characteristic of methods that do transfer to different problems or domains (i.e., they represent weak methods).

assumptions are addressed, e.g., in familiar domains where minor improvements are expected, are made relatively quickly, and there is a limit to both their value and risk. The efficiency curves for ABR may resemble those for weak methods when central assumptions must be addressed and evaluated in terms of *overall implausibility*. Concern with central assumptions causes elements of a story to depend on one another for their meaning and plausibility, like pieces of a single puzzle. A change in any one can have implications for all the others, via their influence on the central assumption. Value is not realized until the process stabilizes, i.e., when it is less likely that new information will prompt the emergence of an entirely new pattern. This is a characteristic of less familiar or less predictable domains. S-shaped efficiency curves graphically depict the idea of change by punctuated equilibrium (Gould, 1977): a period of little or no progress toward commitment ends with a burst of rapid progress, which gradually tapers off.

Strategy selection and time management in TDM depends on decision makers' *awareness* of the strategies afforded to them by the problem situation, and the efficiency of those methods for arriving at successful decisions. *Perceived* affordance for commitment may be based on experience, i.e., from sampling specific durations of a method (i.e., an interval on the horizontal axis in Figure 6) and learning the results (i.e., the corresponding change in value on the vertical axis). Perceived affordance in each sample would be the decision maker's *experienced success rate*, i.e., the gain (in ultimate success, or in some surrogate like perceived uncertainty) divided by the amount of time the decision maker took to achieve it. A subjective affordance estimate for a strategy corresponds roughly to the slope of a straight line drawn between the start and end points of an episode involving that strategy on the applicable *objective* affordance function (e.g., a curve in Figure 6). The more experience a decision maker has with a method, the more such segments she collects, bringing the overall shape of the affordance curve into play.

Strategy selection and time management must also take external time constraints into account. These can be represented in the same format as efficiency curves for methods. Recall that the curves in Figure 6A represent the gain in (chance or degree of) success in a decision mode as a function of time using a particular uncertainty handling method. Similar curves can represent the cost of delayed commitment, i.e., the percentage loss in (chance or degree of) success as a function of time spent delaying commitment or action. If these loss curves are negatively accelerated (i.e., if they look like the curves for strong methods), delay costs kick in immediately; depending on parameters, costs may increase more or less steeply with increasing delay until reaching a maximum of 100% or less of the value of the action. On the other hand, if loss curves are S-shaped curves, there is a period of respite in which losses are very small, after which they increase at an accelerating rate before tapering off at some maximum. The costs of delayed commitment can be learned from experience in the same way as affordance curves: by sampling different

durations of delay and experiencing the differences in realized success (i.e., appropriateness of action, desirability of outcomes, reliability of practices).

Adaptation of Time Management and Strategy Selection

We do not propose that agents optimize efficiency, either objectively or subjectively. To illustrate a more plausible mechanism, we discuss a process called *melioration* (Herrnstein, 1997; Herrnstein & Prelec, 1992; Herrnstein, 1982). The advantages of melioration (as distinguished from *optimization*) are: it can be executed without discernable cognitive effort (it fits the presumably automatic behavior of herbivores who allocate their foraging time among different patches); it does not assume omniscience but offers a plausible route for learning; and it applies to situations where some proposed optimization methods do not (e.g., it accommodates affordance or cost curves that are not concave; Cohen, 2011; cf. Pirolli, 2007). Most importantly, it can explain how afforded uncertainty handling strategies simultaneously compete for time both with one another and with earlier commitment to action.

In any new decision problem, the decision maker will perceive certain types of uncertainty plus a set of subjectively afforded uncertainty handling methods. With the latter, she may have a number of experiences, each with an associated efficiency and time invested. For each afforded method, the experience most similar to the present situation (e.g., in number and type of perceived uncertainties as well as substantive issues) is retrieved. The average efficiency of these methods (i.e., total gain over total time invested), scaled to reflect the specific stakes in the present problem relative to the historical average, can be thought of as the decision maker's *aspiration level* for gain per unit time from uncertainty handling in the new situation. The decision maker will also have an impression of the value that was sacrificed by delaying commitment for different durations in problems similar to the current one. For each subjectively afforded method, if the method's efficiency is higher than the aspiration level for uncertainty handling and the estimated loss for delay, the decision maker increments the time she uses the method (relative to similar past experiences). If its efficiency is less than either the aspiration level or expected loss, she decrements the time for using the method. In exactly the same way, if the loss for delay is greater than the aspiration level for uncertainty handling methods, action will be moved forward at the expense of deliberation time; if the loss is less than the aspiration level, delay will be increased and more time spent deliberating. The more experience a decision maker has with different methods and their affordances, and the better able she is to recognize different types of uncertainty, the more adapted her time management and strategy selection will be. Graded decrements will eventually drive out methods that do not gain efficiency as less time is devoted to them. If no methods gain

efficiency as time is reduced, commitment will eventually occur without discernable deliberation. Conversely, if little gain is experienced as a result of more timely action, commitment will be delayed a bit more. Together with real-time adjustments for perceived uncertainty, this simple mechanism results in a dynamically adaptive Suppression strategy, or stopping rule, for decision making (Cohen, 2011). It also predicts the effects of Forestalling and Buying time. In the former, reduced risk (due, for example, to incremental decision making) decreases the value of uncertainty handling relative to the cost of delay, leading to earlier commitment and more timely action. In the latter, reduced cost of delay has the opposite effect, allowing a greater range of uncertainty handling strategies to clear threshold. [19]

By contrast, decision makers with less experience will be unable to differentiate methods by efficiency or to adapt the time they take to varying constraints. Even worse, errors may occur due to limited sampling. For example, brief experience of methods that require time investment before they pay off (i.e., S-shaped affordance curves) will result in even less time being allocated to them, until they disappear – even if their ultimate efficiency, had they been given enough time, would have been superior (like those in Figure 6B). Conversely, if decision makers act only after long delays, they never experience the benefits of more rapid responses. [20]

---

[19] This implies a testable hypothesis regarding training (Cohen, 2011). Time management is *not* best taught, in general, as an explicit "executive control" strategy. Decision makers become more adaptive in their use of time by being exposed deliberately to activities for different durations, and receiving feedback regarding the value achieved by doing so. They can practice acting quickly and acting late to realize the value (or lack of value) of timely action, and they can practice different deliberation methods for different periods of time to learn about delayed benefits (in the case of weak methods) and declining returns (in the case of strong methods).

[20] The proposed mechanism has numerous experimentally testable implications, which we can only suggest here (Cohen, 2011). The general principle is that highly experienced decision makers will divide their time between uncertainty handling and more immediate action, and among different uncertainty handling methods, only if there are time allocations for which the losses and efficiencies are equal. Let us assume decision makers have been exposed to different uncertainty handling methods and different action latencies as described in note 19, and efficiency curves have been measured for different methods as described in note 18. With respect to *allocation of time among methods*, simulation with efficiency curve values in Figure 6A and B and assuming a fixed moderate amount of time for deliberation (slightly less than necessary for any method to reach its maximum), predicts the following: If both adaptive intuitive and adaptive deliberative methods are afforded, time will probably be divided between them; for example, use of the intuitive method will be followed by a brief period of deliberation, unless redundancy has reduced the latter's value to one-seventh or less. In Figure 6A, if either strong method is used, nether weak method will be, even if afforded. By contrast, in Figure 6B, if the top down weak method is afforded along with either strong method, the strong method may be used alone; but the top-down weak method (e.g., taking a look at theory) may be used first followed by transition to the strong method (e.g., well-learned practices) (melioration in this case approaches a mixed equilibrium from only one direction). Finally, in both Figure 6A and B, in unfamiliar situations where strong methods are not available, only one of the weak methods will be used; whichever is strongest initially will push out the other even if there is no redundancy. With respect to *delay in*

One additional control tactic should be mentioned, which exploits the interaction between decision modes and uncertainty. As we have noted, different decision making modes pose questions about different *issues*; thus, the amount of uncertainty may be affected when the same decision maker views the same decision problem from the perspective of different modes. A situation that offers high affordance for matching (e.g., by means of well-known rules for traditional role-based behavior) may be difficult to parse in terms of outcomes and their relative values (i.e., choice). Conversely, in other situations it may be easy to identify the most advantageous option (choice) but difficult to decide if it is ethical (matching). Finding a successful action by experimentation (i.e., reassessment) may be easier than choice or matching if the situation is novel and actions are reversible; but choice or matching will be easier when actions are costly and the domain is well understood. These affordances may well influence which mode is used. For the same reasons, decision makers may find themselves changing modes if uncertainty becomes excessive; for example, a decision maker may decide to follow the usual practice rather than trust an assumption-laden analysis of outcomes. Whether this happens depends not only on uncertainty but on dispositions to favor different modes, which may vary in strength with individual experience, personality, social context, task, domain, and cultural background.[21]

All three modes play a role in determination of the *decision making process*: Matching primes secondary questions and a path to commitment, and reassessment fills in when the process includes random search rather than guidance by questions; the balance of the two determines action uncertainty. Choice helps manage time, the selection of methods, and to some degree the selection of modes – i.e., the amount of effort that is actually expended to overcome action uncertainty. Simultaneously, the first-order *decision making mode* that is supported by these

---

commitment, simulation with cost of delay curves taken from Figure 6A and uncertainty handling efficiency curves from Figure 6B, predicts the following*:* When delay costs and uncertainty handling efficiency are both negatively accelerated, the prediction is straightforward: uncertainty handling will occur before commitment if and only if its benefits rise faster than its costs. However, if either delay cost or method efficiency has an S-shaped curve, results can be problematic. For example, if shorter periods of deliberation are tried first, they may be decremented until they eventually disappear in favor of immediate commitment, despite potential benefits from longer deliberation; if long periods of deliberation are tried first, they may be incremented while action is indefinitely delayed, despite benefits from earlier action. These findings underscore the importance of deliberate exposure to different durations under these conditions (note 19).

[21] The hypothesis that decision makers will tend to avoid the use of modes where uncertainty is high relative to other modes can be experimentally tested, by presenting scenarios with characteristics like those described in the text, and assessing the impact on decision mode by criteria described in note 9. Mode-specific uncertainty may interact with cultural and personal values in their effects on decision mode, as described in note **Error! Bookmark not defined.**.

processes may be matching, choice, or reassessment, as shown in Table 2. It is not first-order decision making mode that determines action uncertainty and cognitive effort; it is decision making *process*, and the second-order role of the three modes at that level.

CONCLUSION

Summary

We have argued for a conceptualization of decision making not as choice, but more simply and generally as *change in commitment*. This led to a Trimodal Decision Making theory, based on recognition of three ways commitments can change, corresponding to modes of decision making that both complement, and sometimes conflict with, one another:

- *Matching* generates an action that fits the agent's role in the current situation. Reasons refer to features of the current setting and decision maker role that are linked to an appropriate action via *practices*, such as social customs, organizational routines, expert strategies, personal traits, or moral principles.
- *Reassessment* monitors, evaluates, and adjusts existing practice to deal with unexpected problems, resolve disagreement, or take advantage of opportunities. Reasons for action refer to its having *survived or not survived* testing thus far.
- *Choice* selects one among given alternatives by comparing their ability to achieve desirable outcomes. Reasons for choice are based on the superior efficiency of the selected option in achieving the decision maker's ends.

Different modes are associated with different criteria for successful performance: i.e., appropriate action (matching), desirable outcomes (choice), and meeting challenges (reassessment). Each mode applies to commitments created by the others and creates commitments to which the others can be applied. Yet they represent competing visions of rationality, which in the most difficult cases can directly conflict.

TDM includes an ecological conception of *action uncertainty*, as *distance in cognitive effort from actionable commitment*. High or low action uncertainty, and methods that are high or low in cognitive effort, may occur in all three modes of decision making, highlighting similarities in the difficulties they encounter and the strategies they use to manage them, despite differences in the content of reasoning and the normative definition of success. Relationships among existing decision making research paradigms are clarified by locating them in a matrix of decision modes by uncertainty handling methods. A process called *melioration* can result in reasonable allocations

of time among competing methods or strategies, including immediate action versus continued decision making. A variety of predictions follow when we classify cognitive strategies by the efficiency with which they achieve actionable commitment. The three decision modes operate at the second-order level of cognitive strategies to determine action uncertainty and effort expended.

TDM reframes the notions of decision making and uncertainty in a way that is suited for naturalistic research, decision aiding, and training. On the research side, it offers categories and definitions to guide naturalistic observations, implies empirically testable hypotheses, and brings out similarities and differences among various decision paradigms. On the application side, it calls for an expanded normative basis and enlarges the function of cognitive engineering beyond the supporting role assigned to prescription in the standard story.

Prescriptive Implications

The TDM framework provides a basis for decision support that is more comprehensive than current variants of decision theory, naturalistic decision making, or critical thinking. We will mention a handful of high level design guidelines that are suggested by the framework, which we hope to illustrate and elaborate in the future.[22]

Compatibility. Our most basic proposition is that decision aids should support modes of decision making that are appropriate for the situation and the decision maker role, according to matching. This implies that a traditional decision model for comparing multiple alternatives should almost never be the default decision aiding mode. Likewise, support for reassessment should direct attention to aspects of the current course of action that are permitted to change under ethical, organizational, or other constraints. This does not mean that an aid must always impose an initial phase of support for matching (although this may often be the case). It means that cognitive engineers must be aware of the operation of matching and the constraints it imposes on users.

Focus. Displays should direct attention to factors that are relevant to the problem as viewed through the current mode: i.e., (a) situation awareness, roles, values, and course of action creation for matching, (b) comparison of outcomes and actions for choice, and (c) course of action monitoring, challenging, and modification for reassessment. For matching, the aid should compare the actual situation with the conditions of relevant practices. For choice, the aid should compare projected action outcomes with desired features across choice alternatives. For reassessment, the aid should compare actual events and actions with those expected.

[22] Some on-going decision aiding work (Cohen et al., 2011; Cohen et al., 2008) works with these principles. A prototype system has achieved significant acceptance from US Army end users in the field (Weltman et al., 2011).

Control. Aids should help users adapt to constraints of time and stakes. This includes reducing risks by increasing the robustness of courses of action (*Forestalling*), reducing the cost of delaying irreversible commitment (Buying time), and acting immediately on the current best solution when benefits of further risk reduction are outweighed by the costs of delay (the *Suppress* tactic).

Transitions. Displays for different modes should support seamless transitions to other modes when appropriate. For example, displays for creating options and goals (matching) should relate transparently to displays for comparing options in terms of goals (when choice is appropriate), and both should relate transparently to displays for vetting courses of action before implementation and for monitoring goal achievement during implementation (reassessment).

Uncertainty. Aids should focus on the appropriate intersection of decision mode and action uncertainty. Issues relevant to the decision mode should be highlighted as discussed under Focus, e.g., obstacles in determining rule fit in matching, option efficiency in choice, and course of action reliability in reassessment. Variations in levels of action uncertainty across problems and users should be accommodated, e.g., pattern recognition tactics should be reserved for the lowest levels of action uncertainty; weak methods and lengthy reasoning processes should be reserved for the highest levels of action uncertainty; strategies for adaptive deliberation should be supported for moderately difficult problems in familiar domains. Displays should seamlessly track uncertainty levels as they change in response to uncertainty handling tactics; for example, affording users the opportunity to evaluate and modify peripheral assumptions in cases of low action uncertainty, or more central ones if overall implausibility increases.

The standard story, which identifies decision making with choice, encourages people to make retrospective sense of their actions to themselves and others as deliberate and optimal (March, 1979) – and conversely, to fault others' decisions for not achieving the same perfection. This story is an ideal congealed into a cultural myth. It has also proven to be a trap for designers, who are still capable of being surprised by the low level of user acceptance even for well-engineered decision support, when it focuses exclusively on deliberative comparison of options. Trimodal Decision Making insists on the normative validity of three paths to commitment, describes different levels of affordance for those paths, and prescribes principles of decision aiding. It fleshes out and corrects the standard story and other overly narrow approaches to decision making, and, we hope, will be a useful source of implications and challenges for cognitive engineering practice and research in the future.

REFERENCES

Alexander, E.R. (1979). The design of alternatives in organizational contexts: A pilot study. *Administrative Science Quarterly, 24*, 382-404.

Alexander, L. & Moore, M., (2008). Deontological ethics. In E.N. Zalta (ed.), *The Stanford Encyclopedia of Philosophy*. <http://plato.stanford.edu/archives/fall2008/entries/ethics-deontological.

Appleby, J. (2010). *The Relentless Revolution: A History of* Capitalism. New York: Norton.

Argyris, C., & Schön, D.A. (1974). *Theory in practice: Increasing professional effectiveness*. San Francisco: Jossey Bass.

Argyris, C., & Schön, D. A. (1996). *Organizational learning II: Theory, Method, and Practice.* reading, MA: Addison-Wesley.

Atran, S. (2002). *In Gods We Trust: The Evolutionary Landscape of Religion*. Oxford: Oxford University.

Axelrod, R. 1984, *The Evolution of Cooperation*, Basic Books.

Bacharach, M. (2006). *Beyond Individual Choice: Teams and Frames in Game Theory*. Princeton NJ: Princeton University Press.

Bak, P. (1996). *How Nature Works: The Science of Self-Organized Criticality*. New York: Springer-Verlag.

Barber, B., M. Heath, C., O'dean T. (2003). Good reasons sell: Reason–based choice among group and individual investors in the Stock Market. *Management Science, 49,* 1636-1652.

Baron, J. (1993). *Morality and rational choice*. Dordrecht: Kluwer.

Baron, J., & Leshner, S. (2000). How serious are expressions of protected values? *Journal of Experimental Psychology: Applied*, *6*(3), 183-194.

Beach, L. R. (2010). *The Psychology of Narrative Thought*. Xlibris Corporation.

Beach, L. 1990, *Image theory: decision making in personal and organizational contexts*, John Wiley & Sons, New York.

Beach, L., & Connolly, T. (2005). *The Psychology of Decision Making*. Thousand Oaks: Sage.

Bennett, K. B., & Walters, B. (2001). Configural display design techniques considered at multiple levels of evaluation. *Human Factors*, *43*(3), 413-434.

Bennett, K. B., Posey, S. M., & Shattuck, L. G. (2008). Ecological interface design for military command and control. *Journal of Cognitive Engineering and Decision Making, 2*(4), 349-385.

Berry, D. C., & Dienes, Z. (1993). *Implicit learning*. Mahwah NJ: Lawrence Erlbaum Associated, Inc.

Betsch, T. & Haberstroh, S. (eds.) 2004, *The Routines of Decision Making*, Psychology Press, East Sussex UK.

Bicchieri, C. 2005, *The Grammar of Society: The Nature and Dynamics of Social Norms*, Cambridge University Press, Cambridge, UK.

Bicchieri, C., Jeffrey, R. & Skyrms, B. (eds.) 2009, *The Dynamics of Norms*, Cambridge University Press, Cambridge, UK.

Binmore, K. (1994). *Game Theory and the Social Contract Volume I: Playing Fair*. Cambridge, MA: The MIT Press.

Binmore, K. (2009). *Rational Decisions*. Princeton NJ: Princeton University Press.

Bratman, M. E. (1987). *Intention, plans, and practical reason*. Cambridge, MA: Harvard University Press.

Braybrooke, D. & Lindblom, C. 1963, *A Strategy of Decision: Policy Evaluation as a Social Process*, The Free Press, New York.

Brown, R. V., & Paschoud, P. (2005). *Rational choice and judgment: Decision analysis for the decider.* Hoboken, NJ: Wiley

Camerer, C. (2003). *Behavioral game theory: Experiments in strategic interaction*: Princeton University Press Princeton, NJ.

Chase, W., & Simon, H.A. (1973). Perception in chess. *Cognitive Psychology*, *4*, 55-81.

Chemero, A. (2009). *Radical Embodied Cognitive Science*. Cambridge MA: M.I.T. Press.

Cohen, M. S. (1993). The naturalistic basis of decision biases. In G. Klein, J. Orasanu, R. Calderwood & C. Zsambok (Eds.), *Decision Making in Action: Models and Methods* (pp. 51-99). Norwood, NJ: Ablex.

Cohen, M. S. (2004). Leadership as the orchestration and improvisation of dialogue: Cognitive and communicative skills in conversations among leaders and subordinates. In D. V. Day, S. J. Zaccaro & S. M. Halpin (Eds.), *Leader development for transforming organizations: Growing leaders for tomorrow* (pp. 177-208). Mahwah, NJ: Erlbaum.

Cohen, M. S. (2011). Knowns, Known Unknowns, and Unknown Unknowns: Time and Uncertainty in Naturalistic Decision Making. In K. L. Mosier & U. M. Fischer (Eds.), *Informed by Knowledge: Expert performance in complex situation* (pp. 371-96). New York: Taylor & Francis.

Cohen, M. S., & Ratwani, R. (2011). *Counterinsurgency operational planning tool to train and facilitate complex decision making*. Paper presented at the Human Systems Integration Symposium 2011, https://www.navalengineers.org/publications/symposiaproceedings/Pages/HSIS2011.aspx.

Cohen, M., Schum, D.A., Freeling, A.N.S., & Chinnis, J.O., Jr. (1984). *On the art and science of hedging a*

conclusion: alternative theories of uncertainty in intelligence analysis. Falls Church, VA: Decision Science Consortium, Inc.

Cohen, M., Laskey, K.B., Chinnis, J.O., Jr., & Ulvila, J.W. (1986). *Report on uncertainty models: identification and critical analysis of models of uncertainty with potential applicability to intelligence data bases.* Falls Church, VA: Decision Science Consortium, Inc.

Cohen, M., Freeman, J.T., & Wolf, S. (1996). Meta-recognition in time stressed decision making: Recognizing, critiquing, and correcting. *Human Factors*, *38*(2), 206-219.

Cohen, M.S., Freeman, J.T., & Thompson, B.B. (1998). Critical thinking skills in tactical decision making: A model and a training method. In J. Cannon-Bowers & E. Salas (Eds.), *Decision-Making Under Stress: Implications for Training & Simulation* (pp. 155-89). Washington, DC: American Psychological Association.

Cohen, M. S., & Thompson, B. B. (2001). Training teams to take initiative: Critical thinking in novel situations. In E. Salas (Ed.), *Advances in Human Performance and Cognitive Engineering Research* (Vol. 1, pp. 251-91). Amsterdam: JAI.

Cohen, M.S., Adelman, L.A., Bresnick, T., Freeman, F.M., Salas, E., & Riedel, S. (2006). Dialogue as a medium (and message) for training critical thinking. In R. Hoffman (Ed.), *Expertise Out of Context*, pp. 219-261. Mahwah, NJ: Erlbaum.

Cohen, M. S., de Visser, E., Freedy, E., Freedy, A., LeGoullon, M., & Weltman, G. (2008). *Intuitive and analytic support for decision making.* Paper presented at the Conference on Human Centered Processes, Delft, The Netherlands.

Connolly, T., & Wagner, W. (1988). Decision cycles. In R. Cardy, S. Puffer & M. Newman (Eds.), *Advances in Information Processing in Organizations* (Vol. 3, pp. 183-205). Greenwich, CT: JAI Press.

Cross, N. (2006). *Designerly ways of knowing*. New York: Springer Verlag.

Cross, N., (2007). Forty years of design research, *Design Studies 28*, 1-4.

Dawes, R. M. (1979). The robust beauty of improper linear models in decision making. *American Psychologist, 34*, 571-582.

Dawkins, R. (2006). *The Selfish Gene*. Oxford: Oxford University.

de Finetti, B. (1964). Foresight: Its logical laws, its subjective sources. In J. Kyburg, H.E. & H. E. Smokler (Eds.), *Studies in Subjective Probability*. New York: Wiley.

Dennett, D. C. (1978). Why the Law of Effect will not go away. In *Brainstorms: philosophical essays on mind and psychology*. Montgomery VT: Bradford Books.

Dennett, D. C. (1996). *Darwin's Dangerous Idea: Evolution and the Meanings of Life*. New York: Touchstone Books.

Dosi, G., Marengo, L., Bassanini, A., & Valente, m. (1998). Norms as emergent properties of adaptive learning: The case of economic routines. In P. A. Danielson (Ed.), *Modeling Rationality, Morality, and Evolution* (pp. 442-63). Oxford: Oxford University Press.

Dubra, J., Maccheroni, F., & Ok, E. A. (2004). Expected utility theory without the completeness axiom. *Journal of Economic Theory*(115), 118-133.

Einhorn, H. J., & Hogarth, R. M. (1985). Ambiguity and uncertainty in probabilistic inference. *Psychological Review*, *92*(4), 433-461.

Ellsberg, D. (1988). Risk, ambiguity, and the Savage axioms. In P. Gardenfors & N.-E. Sahlin (Eds.), *Decision, probability, and utility: Selected reasoning* (pp. 245-69). Cambridge, UK: Cambridge University Press.

Elster, J. (1974). Ulysses and the sirens: Studies in rationality and irrationality. Cambridge, Ma: Cambridge University Press.

Elster, J. (2000). *Ulysses unbound*. Cambridge, MA: Cambridge University Press.

Endsley, M. (1995). Toward a theory of situation awareness in dynamic systems. *Human Factors*, *37*(1), 32-64.

Endsley, M. R. (2000). Theoretical underpinnings of situation awareness: A critical review. In M. R. Endsley & D. J. Garland (Eds.), *Situation awareness analysis and measurement* (pp. 3-32). Mahwah, NJ: Erlbaum.

Ericsson, K. A., Krampe, R. T., & Tesch-Romer, C. (1993). The role of deliberate practice in the acquisition of expert performance. *Psychological Review, 100*(3), 363-406.

Ericsson, K. (2005). Superior decision making as an integral quality of expert performance: Insights into the mediating mechanisms and their acquisition through deliberate practice. In H. Montgomery, R. Lipshitz & B. Brehmer (Eds.), *How Professionals Make Decisions* (pp. 135-67). Mahwah, NJ: Lawrence Erlbaum Associates.

Etzioni, A. (2000). Social norms: Internalization, persuasion, and history. *Law & Society Review, 34*(1), 157-178.

Etzioni, A. (1988). *The Moral Dimension: Toward a new economics*. New York: The Free Press.

Evans, J. S. T. (2007). *Hypothetical Thinking: Dual Processes in* Reasoning and Judgment. New York: Psychology Press.

Fiedler, K., & Juslin, P. (Eds.) (2006). *Information Sampling and Adaptive Cognition*. Cambridge, UK: Cambridge

University Press.

Fiske, A. P., & Tetlock, P. E. (1997). Taboo tradeoffs: Reactions to transactions that transgress the spheres of justice. *Political Psychology*, *18*(2), 255-297.

Garud, R. & Karnoe, P. 2001, *Path Dependence and Creation*, Erlbaum, Mahwah NJ.

Gentner, D., Holyoak, K. J., & Kokinov, B. N. (Eds.). (2001). *The analogical mind: Perspectives from cognitive science*. Cambridge MA: M.I.T. Press.

Gibbard, A. (1990). *Wise choices, apt feelings: A theory of normative judgment*. Cambridge, MA: Harvard University Press.

Gibson, J. J. (1979). *The Ecological Approach to Visual Perception*. Boston: Houghton Mifflin.

Gilovich, T., Griffin, D., & Kahneman, D. (Eds.). (2002). *Heuristics and biases: The psychology of intuitive judgment*. Cambridge, UK: Cambridge University Press.

Gigerenzer, G., Hoffrage, U., & Kleinbolting, H. (1991). Probabilistic mental models: A Brunswickian theory of confidence. *Psychological Review*, *98*(4), 506-528.

Gigerenzer, G., & Selten, R. (2001). *Bounded rationality: The adaptive toolbox*. Cambridge MA: MIT Press.

Gigerenzer, G. (2000). *Adaptive thinking: Reasoning in the real world*. Oxford: Oxford University Press.

Gigerenzer, G., & Goldstein, D. G. (1996). Reasoning the fast and frugal way: Models of bounded rationality. *Psychological Review, 103*(4), 650-669.

Gigerenzer, G., Todd, P. M., & Group, A. R. (1999). *Simple Heuristics that Make Us Smart*. Oxford: Oxford University.

Gilbert, M. (1996). *Living Together: Rationality, Sociality, and Obligation*. Lanham MA: Rowman & Littlefield.

Ginges, J., Atran, S., Medin, D. L., & Shikaki, K. (2007). *Sacred bounds on rational resolution of violent political conflict*, from www.pnas.org/cgi/doi/10.1073/pnas.0701768104.

Gillies, D. (2000). *Philosophical Theories of Probability*. New York: Routledge.

Gilovich, T., Griffin, D., & Kahneman, D. (Eds.). (2002). *Heuristics and biases: The psychology of intuitive judgment*. Cambridge, UK: Cambridge University Press.

Goldstein, D. G. (2001). Group Report: Why and when do simple heuristics work?. In G. Gigerenzer & R. Selten (Eds.), *Bounded rationality: The adaptive toolbox* (pp. 173-90). Cambridge MA: MIT Press.

Goodman, N. (1965). *Fact, fiction, and forecast.* Indianapolis: The Bobbs-Merrill Company.

Goschke, T., & Kuhl, J. (1996). Remembering what to do: Explicit and implicit memory for intentions. In M.

Brandimonte, M. A. McDaniel & G. O. Einstein (Eds.), *Prospective Memory: Theory and Applications*. Mahwah NJ: Erlbaum.

Gould, S. J. (1977). *Ever Since Darwin: Reflections in Natural History*. New York: Norton.

Grice, P. (1989). *Studies in the way of words*. Cambridge MA: Harvard University Press.

Gunther, K. (1993). *The Sense of Appropriateness: Application Discourses in morality and Law*. Albany: State University of New York.

Guth, W., Schmittberger, R., & Schwartze, B. (1982). An experimental analysis of ultimatum game bargaining. *Journal of economic Behavior and Organization*, *3*, 367-388.

Hammond, P. J. (1988). Consequentialist Foundations for Expected Utility. *Theory and Decision*(25), 25-78.

Hammond, K. R., & Stewart, T. R. (Eds.). (2001). *The Essential Brunswick: Beginnings, Explications, Applications*. Oxford, UK: Oxford University Press.

Hanselmann, M., & Tanner, C. (2008). Taboos and conflicts in decision making: Sacred values, decision difficulty, and emotions. *Judgment and Decision Making*, *3*(1), 51-63.

Herrnstein, R. J. (1982). Melioration as behavioral dynamism. In M. L. Commons, A. Kacelnik & S. J. Shettleworth (Eds.), *Quantitative Analysis of Behavior: Matching and Maximizing Accounts* (Vol. II, pp. 433-58). Cambridge, MA: Ballinger

Herrnstein, R. J. (1997). *The Matching Law: Papers in Psychology and Economics*. New York: Russell Sage Foundation.

Herrnstein, R. J., & Prelec, D. (1992). A theory of addiction. In G. Loewenstein & J. Elster (Eds.), *Choice over Time* (pp. 331-60). New York: Russell Sage Foundation.

Hoffman, R. R. (Ed.). (2007). *Expertise Out of Context*. Mahwah, NJ: Erlbaum.

Holyoak, K. J., & Thagard, P. (1995). *Mental leaps: Analogy in creative thought*. Cambridge MA: M.I.T. Press.

Hooker, B., & Streumer, B. (2004). Procedural and substantive practical rationality. In A. R. Mele & P. Rawling (Eds.), *The Oxford Handbook of Rationality* (pp. 57-74). Oxford: Oxford University.

Hutchins, E. (1995). *Cognitive in the wild*. Cambridge MA: MIT.

Kahneman, D., Slovic, & Tversky, A. (Ed.). (1982). *Judgment under uncertainty: heuristics and biases*. New York: Cambridge University Press.

Kliegel, M., McDaniel, M. A., & Einstein, G. O. (Eds.). (2009). *Prospective Memory: An Overview and Synthesis of an Emerging Field*. New York: Psychology Press.

Janis, I.L., & Mann, L. (1977). *Decision making: A psychological analysis of conflict, choice and commitment* New York: Free Press.

Jeffrey, R. C. (1965/1983). *The logic of decision*. Chicago, IL: University of Chicago Press.

Jensen, F. V. (2001). *Bayesian Networks and Decision Graphs*. New York: Springer.

Keeney, R. (1992). *Value focused thinking: A path to creative decision making*. Cambridge, MA: Harvard University Press.

Keeney, R. & Raiffa, H. 1976, *Decisions with multiple objectives: Preferences and value tradeoffs*, Wiley, New York.

Khoo, Y.-L., & Mosier, K. (2008). The impact of time pressure and experience on information search and decision-making processes. *Journal of Cognitive Engineering and Decision Making*, *2*(4), 275-294.

Klayman, J. (1995). Varieties of confirmation bias. In J. Busemeyer, R. Hastie & D. Medin (Eds.), *Decision making from the perspective of cognitive psychology* (pp. 385-418). New York: Academic Press.

Klein, G. (1993). *Naturalistic Decision Making: Implications for Design.* Wright Patterson Air Force Base, Ohio: Crew System Ergonomics Information Analysis Center.

Klein, G. (1998). *Sources of power: How people make decisions*. Cambridge, MA: MIT Press.

Klein, G. (2001). The fiction of optimization. In G. Gigerenzer & R. Selten (Eds.), *Bounded rationality: The adaptive toolbox* (pp. 103-22). Cambridge MA: MIT Press.

Klein, G., Orasanu, J., Calderwood, R., & Zsambok, C.E. (Ed.). (1993). *Decision making in action: models and methods*. Norwood, NJ: Ablex.

Klein, G., Phillips, J. K., Rall, E. L., & Peluso, D. A. (2007). A data-frame theory of sensemaking. In R. R. Hoffman (Ed.), *Expertise Out of Context* (pp. 113-55). Mahwah, NJ: Erlbaum.

Knight, F. H. (1921). *Risk, Uncertainty, and Profit*: Houghton Mifflin.

Kolodner, J. 1993, *Case-based reasoning*, Morgan Kaufmann, San Mateo CA.

Kuhn, T. (1996). *The structure of scientific revolutions*. Chicago IL: University of Chicago Press.

Lakoff, G. (1987). *Women, Fire, and Dangerous Things: What Categories Reveal about the Human Mind*. Chicago, IL: University of Chicago Press.

Langlois, R. N., & Cosgel, M. M. (1993). Frank Knight on Risk, Uncertainty and the Firm: A new Interpretation. *Economic Inquiry, XXXI*, 456-465.

LeDoux, J. (1996). *The Emotional Brain*. New York: Simon & Schuster.

Levi, I. (1986). *Hard choices: decision making under unresolved conflict*. New York: Cambridge University Press.

Levinthal, D. (2008). Understanding the role of novelty in organizational adaptation. In J. March (Ed.), *Explorations in Organizations* (pp. 95-103). Stanford, CA: Stanford University Press.

Levinthal, D. (1998). The slow pace of rapid technological change: Gradualism and punctuation in technological change. *Industrial and Corporate Change*, *7*(7), 217-247.

Levinthal, D. (1997). Adaptation on rugged landscapes. *Management Science*, *43*, 934-950.

Levitt, B., & March, J. G. (1988). Organizational Learning. *Annual Review of Sociology*, *14*, 319-340.

Lindblom, C. E. (1959). The science of muddling-through. *Public Administrative Review , 19*, 79-88.

Lindley, D. (1982). Scoring rules and the inevitability of probability. *International Statistical Review*(50), 1-26.

Lipshitz, R. (1993a). Decision making in realistic settings. In Klein, G.A., J. Orasanu, R. Calderwood, & C. Zsambok (Eds.), *Decision making in action: Models and methods*, 172-181). Norwood, NJ: Ablex Publishing.

Lipshitz, R. (1993b). Decision making as argument driven action. In Klein, G.A., J. Orasanu, R. Calderwood, & C. Zsambok (Eds.), *Decision making in action: Models and methods*, 172-181). Norwood, NJ: Ablex Publishing.

Lipshitz, R., & Cohen, M. (2005). Warrants for prescription: Analytically-based vs. empirically-based approaches to prescriptive decision making. *Human Factors,* 47: 102-20.

Lipshitz, R., & Strauss, O. (1997). Coping with uncertainty: A naturalistic decision making analysis. *Organizational Behavior and Human Decision Processes, 69*, 149-163.

Luce, R. D. (2003). Rationality in choice under certainty and uncertainty. In S. L. Schneider & J. Shanteau (Eds.), *Emerging Perspectives on Judgment and Decision Research* (pp. 64-83). Dordrecht, The Netherlands: Kluwer Academic Publishers.

Luce, R. D., & Raiffa, H. (1989). *Games and Decisions: Introduction and Critical Survey*: Dover.

Mandler, J. M. (1984). *Stories, Scripts, and Scenes: Aspects of Schema Theory.* Hillsdale, NJ: Erlbaum.

March, J.G. (1994). *A primer on Decision Making: How Decisions Happen*. NY: Free Press.

March, J. G. (2010). *The Ambiguities of Experience*. Ithaca, NY: Cornell University Press.

March, J. G., Schulz, M., & Zhou, X. (2000). *The dynamics of rules: Change in written organizational codes*. Stanford, CA: Stanford University Press.

Markman, K. D., Klein, W. M. P., & Suhr, J. A. (2009). *Handbook of Imagination and Mental Simulation*. New

York: Psychology Press.

McClennen, E. F. (1990). *Rationality and dynamic choice: Foundational explorations*. Cambridge, UK: Cambridge University Press.

McClennen, E. F. (2004). The rationality of being guided by rules. In A. R. Mele & P. Rawling (Eds.), *The Oxford Handbook of Rationality* (pp. 222-39). Oxford: Oxford University Press.

McDaniel, M. A., & Einstein, G. O. (2007). *Prospective Memory: An Overview and Synthesis of an Emerging Field*. Los Angeles, CA: Sage.

Minsky, M. (1985). *The Society of Mind*. New York: Simon & Schuster.

Mintzberg, H. (1994). *The Rise and Fall of Strategic Planning*. New York: MacMillan.

Mintzberg, H., Raisinghani, D., Théorêt, A., (1976). The Structure of 'Unstructured' Decision Processes. *Administrative Science Quarterly, 21*(2) 246 -275.

Montgomery, H. (1993). The search for a dominance structure in decision making: examining the evidence. In G. Klein, J. Orasanu, R. Calderwood & C. E. Zsambok (Eds.), *Decision Making in Action: Models and Methods* (pp. 182-7). Norwood, NJ: Ablex.

Mosier, K., & Fischer, U. (Eds.). (2011). *Knowledge management: Expert performance in complex situations*. New York: Taylor & Francis.

Neisser, U. (1976). *Cognition and reality: Principles and implications of cognitive psychology*. San Francisco: W.H. Freeman and Company.

Newell, A., & Simon, H.A. (1972).Nisbett, R., & Ross, L. (1980). *Human inference: strategies and shortcomings of social judgment*. Englewood Cliffs, NJ: Prentice-Hall, Inc.

Newell, A., & Simon, H.A. (1972). *Human problem solving*. Englewood Cliffs, NJ: Prentice-Hall.

Nisbett, R. E. (2003). *The geography of thought: How Asians and Westerners think differently...and why*. New York: The Free Press.

Norman, D. (1988). *The Design of Everyday Things*. New York: Doubleday.

Norman, D. (1999). Affordance, conventions, and design. *Interactions*, 38-43.

Nozick, R. (1993). *The nature of rationality*. Princeton, NJ: Princeton University Press.

Orasanu, J. (1993). Decision making in the cockpit. In E. Wiener, B. Kanki & Helmreich, R. L. (Eds.), *Cockpit Resource Management* (pp. 137-72). New York: Academic Press, Inc.

Orasanu, J. & Strauch, B. (1994). *Temporal factors in aviation decision making.* Paper presented at the Proceedings

of the Human Factors and Ergonomics Society 38th Annual Meeting, Santa Monica, CA.

Orasanu, J., & Fischer, U. (1997). Finding decisions in natural environments: The view from the cockpit. In C. Zsambok & G. Klein (Eds.).*Naturalistic Decision Making* (pp. 343-357). Hillsdale, NJ: Erlbaum.

Payne, J., Bettman, J.R., & Johnson, E.J. (1993). *The adaptive decision maker*. New York: Cambridge University Press.

Pearl, J. (2000). *Causality: models, reasoning, and inference*. Cambridge UK: Cambridge University Press.

Pearl, J. (1989). *Probabilistic reasoning in intelligent systems: Networks of plausible inference*. San Mateo CA: Morgan Kaufmann Publishers.

Pennington, N., & Hastie, R. (1992). Explaining the evidence: Tests of the story model for juror decision making. *Journal of Personality and Social Psychology*, *62*(2), 189-206.

Pirolli, P. (2007). *Information Foraging Theory: Adaptive Interaction with Information*. Oxford: Oxford University.

Poletiek, F. (2001). *Hypothesis-testing behavior*. Philadelphia: Psychology Press.

Popper, K. R. (1994). *The Myth of the Framework*. New York: Routledge.

Popper, K. (1995). *Conjectures and refutations*. New York: HarperCollins.

Prelec, D., & Bodner, R. (2003). Self-signaling and self-control. In G. Loewenstein, D. Read & R. Baumeister (Eds.), *Time and Decision: Economic and Psychological Perspectives on Intertemporal Choice* (pp. 277-98). New York: Russell Sage Foundation.

Quine, W. V. O., & Ullian, J. S. (1970). *The web of belief*. New York: Random House.

Raiffa, H. (1968). *Decision analysis: Introductory lectures on choices under uncertainty*. Reading, MA: Addison-Wesley Publishing.

Rasmussen, J., Pejtersen, A. M., & Goodstein, L. P. (1994). *Cognitive systems engineering.* New York: Wiley.

Raz, J. (1990). *Practical Reason and Norms*. Princeton NJ: Princeton University Press.

Rescorla, R. A., and Wagner, A. R. (1972). Theory of Pavlovian conditioning: Variations in the effectiveness of reinforcement and nonreinforcement. In J. B. Black & .. F. Prokasy (Eds.), *Classical Conditioning II* (pp. 64-99). New York: Appleton-Century-Crofts.

Rudolph, J. W. (2003). *Into the Big Muddy and Out Again: Error Persistence and Crisis Management in the Operating Room*: Boston College The Carroll Graduate School of Management Department of Organization Studies.

Sacks, H. (1995). *Lectures on conversation*. Malden, MA: Blackwell.

Savage, L. J. (1954/1972). *The foundations of statistics*. New York, NY: Dover Publications.

Shafir, E., Simonson, I., & Tversky, A. (1993). Reason-based choice. *Cognition*, *49*, 11-36.

Schank, R. C. (1990). *Tell me a story: Narrative and intelligence*. Evanston IL: Northwestern University Press.

Schank, R. C. 1999, *Dynamic memory revisited*, Cambridge University Press, Cambridge UK.

Schank, R., & Abelson, R.P. (1977). *Scripts, plans, goals, and understanding: an inquiry into human knowledge structures*. Hillsdale, NJ: Erlbaum.

Schank, R. C., & Abelson, R. P. (1995). Knowledge and Memory: The Real Story. In R. S. Wyer (Ed.), *Knowledge and Memory: The Real Story. Advances in Social Cognition Vol. VIII*. Hillsdale, NJ: Erlbaum.

Schauer, F. (2009). *Thinking Like a Lawyer: A New Introduction to Legal Reasoning*. Cambridge, MA: Harvard University Press.

Schön, D. A., (1983). *The reflective practitioner.* NY: Basic Books.

Schön, D. A. (1995). *Frame Reflection: Toward the Resolution of Intractable Policy Controversies.* New York: Basic Books.

Shachter, R. D. (2007). Model building with belief networks and influence diagrams. In W. Edwards, R. F. Miles, Jr. & D. von Winterfeldt (Eds.), *Advances in Decision Analysis: From Foundations to Applications* (pp. 177-201). Cambridge: Cambridge University Press.

Schoemaker, P. J. H. (2007). Forecasting and scenario planning: The challenges of uncertainty and complexity. In D. J. Koehler & N. Harvey (Eds.), *Blackwell Handbook of Judgment and Decision Making* (pp. 274-96). Malden MA: Blackwell.

Schwartz, S. (1992). Universals in the content and structure of values: Theoretical advances and empirical tests in 20 countries. In *Advances in experimental social psychology* (Vol. 25, pp. 1-65). San Diego: Academic Press.

Schwartz, S., & Bilsky, W. (1987a). Toward a universal psychological structure of human values. *Journal of Personality and Social Psychology*, *53*, 550-562.

Schwartz, S., & Bilsky, W. (1987b). Toward a theory of the universal content and structure of values: Extensions and cross-cultural replications. *Journal of Personality and Social Psychology*, *58*, 878-891.

Serfaty, D., MacMillan, J., Entin, E.E., & Entin, E.B. (1997). The decision making expertise of battle commanders. In &. G. K. C.E. Zsambok (Ed.), *Naturalistic Decision Making*. Mahwah, NJ: Erlbaum.

Simon, H. A. (1956). A comparison of game theory and learning theory. *Psychometrika*, *21*, 267-272. Reprinted in Simon, H. A. (1982). *Models of Bounded Rationality: Behavioral Economics and Business Organization*

(Vol. 2). Cambridge MA: MIT Press.

Simon, H. A. (1955). A behavioral model of rational choice. *Quarterly Journal of Economics*, *69*, 99-118.

Simon, H, A. (1960). *The New Science of Management Decision.* New York: Harper & Row.

Simon, H. A. (1972). Theories of bounded rationality. In C. B. Radner & R. Radner (Eds.), *Decision and Organization* (pp. 161-76). Amsterdam: North-Holland. Reprinted in Simon, H. A. (1982). *Models of Bounded Rationality: Behavioral Economics and Business Organization (Vol. 2)*. Cambridge MA: MIT Press.

Simon, H. A. (1976). From substantive to procedural rationality. In S. Latsis (Ed.), *Method and Appraisal in Economics* (pp. 129-48). Cambridge UK: Cambridge University Press. Reprinted in Simon, H. A. (1982). *Models of Bounded Rationality: Behavioral Economics and Business Organization (Vol. 2)*. Cambridge MA: MIT Press.

Simon, H. A., & Newell, A. (1958). Heuristic problem solving: The next advance in operations research. *Operations Research*, *6*(2), 1-10.

Simon, H. A. (1987). Satisficing. In U. Eatwell, M. Milgate & P. Newman (Eds.), *The New Palgrave: A Dictionary of Economics* (Vol. 4, pp. 243-5). New York: Macmillan.

Skinner, B. F. (1953). *Science and Human Behavior*. New York: The Free Press.

Smith, V. L. (2008). *Rationality in Economics*. Cambridge, UK: Cambridge University Press.

Smithson, M. (1989). *Ignorance and uncertainty: Emerging paradigms*. New York: Springer-Verlag.

Sperber, D., & Wilson, D. A. (1995). *Relevance: communication and cognition*. Oxford: Blackwell.

Spivey, M. (2007). *The Continuity of Mind*. Oxford: Oxford University.

Staw, B. M. (1976). Knee-deep in the big muddy: a study of escalating commitment to a chosen course of action. *Organizational Behavior and Human Performance*, *16*(1), 27-44.

Staw, B. M., & Ross, J. (1989). Understanding behavior in escalation situations. *Science*, *246*(216), 220.

Sutton, R. S., & Barto, A. G. (1998). *Reinforcement Learning*. Cambridge, MA: The M.I.T. Press.

Tetlock, P. E. (1996). An alternative metaphor in the study of judgment and choice: People as politicians. In R. M. Hogarth & W. Goldstein (Eds.), *Judgment and Decision Making: An Interdisciplinary Reader*. Cambridge: Cambridge University Press.

Tetlock, P. E., Kristel, O. V., Elson, S. B., Green, M. C., & Lerner, J. S. (2000). The psychology of the unthinkable: taboo tradeoffs, forbidden base rates, and heretical counterfactuals. *Journal of Personality and Social*

Psychology, *78*(5), 853-870.

Thagard, P. (2000). *Coherence in thought and action*. Cambridge MA: MIT Press.

Thompson, M., Ellis, R., & Wildavsky, A. (1990). *Cultural Theory*. Boulder: Westview Press.

Toulmin, S. E. (1958). *The uses of argument*. Cambridge: Cambridge University Press.

Toulmin, S., Rieke, R., & Janik, A. (1984). *An introduction to reasoning*. New York: MacMillan Publishing Co.

Tversky, A. (1972). Elimination by aspects. *Psychological Review*, *79*(4), 281-299.

von Neumann, J., & Morgenstern, O. (1944). *Theory of Games and Economic Behavior*. Princeton NJ: Princeton University Press.

von Winterfeldt, D. & Edwards, W. 1986, *Decision analysis and behavioral research*, Cambridge University Press, Cambridge, UK.

Wakker, P. P. (1999). *Justifying Bayesianism by dynamic decision principles*. Paper presented at the FUR IX, Marrakesh.

Wakker, P. P. (2010). *Prospect Theory for Risk and Ambiguity*. Cambridge, UK: Cambridge University Press.

Walton, D. N. (1996). *Argumentation schemes for presumptive reasoning*. Mahwah NJ: Lawrence Erlbaum Associated, Inc.

Wason, P. (1968). Reasoning about a rule. *Quarterly Journal of Experimental Psychology*, *20*, 273-281.

Watkins, E. T. (2000). *Improving the analyst and decision-maker's perspective through uncertainty visualization*. Wright Patterson Air Force Base, Ohio: Air Force Institute of Technology.

Weinreb, L. L. (2005). *Legal Reason: The Use of Analogy in Legal Arguments*. Cambridge, UK: Cambridge University Press.

Weltman, g., Freedy, A., Ratwani, R., Freedy, E., Chung, L., Todd, K., et al. (2011). *Decision Infrastructure for COIN Operational Planning (DICOP): Iraq Operational Evaluation Report*. Sharman Oaks, CA: Perceptronics Solutions.

Wilson, T., & Schooler, J. W. (1991). Thinking too much: Introspection can reduce the quality of preferences and decisions. *Journal of Personality and Social Psychology*(60), 181-192.

Witte, Eberhard (1972), "Field research on complex decision-making processes - the phase theorem", *International Studies of Management and Organization*, 156-182.

Yates, J. F., Veinott, E. S., & Patalano, A. L. (2003). Hard decisions, bad decisions: On decision quality and decision aiding. In S. L. Schneider & J. C. Shanteau (Eds.), *Emerging perspectives on judgment and*

decision research (pp. 13-63). New York: Cambridge University Press.

Table Captions

Table 1. Summary of three decision making modes considered as arguments for action.

Table 2. Examples from the research literature classified by decision mode and uncertainty handling methods. Bold items are illustrated by figures in this article.

	Matching	**Choice**	**Reassessment**
Starting commitment	No COA	Disjunction of mutually exclusive COAs	A single COA
Implicit framing question	What should a person in my role do in this kind of situation?	Which option is the best means to my ends?	Is my current course of action reliable?
Grounds (focus of attention)	Present situation and decision maker role	Future outcomes of available options	Past performance of current COA
Claim (commitment change if successful)	Add a COA	Select a COA	Retain or reject a COA
Warrant (why grounds are relevant)	COA is *appropriate* for someone in my role in this kind of situation.	COA is at least as *efficient* a means to my ends as any other available option.	Reject COA that *failed* challenges in similar situations; otherwise, retain.
Backing (normative basis of warrant)	*Duty*: Personal policies, group norms, universal ethical principles	*Instrumental rationality*: Consistency among outcome preferences, outcome forecasts, and action choices	*Trial and error*: Incremental variation and selective retention of successes as a basis for reliable innovation
Intrinsic rebuttals to claims (characteristic potential errors)	1. Situation or role has changed. 2. Certain duties conflict in this situation or in general.	Beliefs or preferences regarding outcomes were assessed incorrectly due to limited knowledge.	Testing process was too restricted in scope or non-representative of current situation.
Rebuttals to the solutions generated in *other* modes	Choice and reassessment endorse actions that conflict with expertise, ethics, cooperation, self-control, and inherited wisdom.	Matching and reassessment fail to optimize due to bias in favor of the status quo and failure to consider all possibilities.	Choice and matching yield unintended consequences in novel situations because they do not test hidden assumptions or adapt to their failure.

	Matching	Choice	Reassessment
Adaptive Intuitive (strong)	Pattern recognition in highly familiar situations, e.g., by experts (Chase and Simon, 1972), **recognition-primed decisions** by experienced decision makers (Figure 2, Klein, 1998). **Condition-action decisions** (Figure 3, Orasanu & Fischer, 1997).	Learned, but relatively automatic time allocation among simultaneously available activities (e.g., melioration as described by Herrnstein, 1997); recognition heuristic (selecting an answer recognized as more familiar than alternatives, Gigerenzer & Selten, 2001)	Experienced practitioners overcome dangers of automaticity by learning *habits* of monitoring. mindfulness, and quick recognition of cue patterns that signal trouble (Schön, 1983). **Go-no go decisions** (Figure 3, Orasanu & Fischer, 1997).
Adaptive Deliberative (strong)	Situation awareness (Endsley, 1995, 2000), **story construction** (Figure 4, Pennington and Hastie, 1992; Cohen et al., 1996, 1998, 2001), and expert's searching for a correct problem representation (Larkin, 1981; Larkin, McDermott, Simon, & Simon, 1980; Adelson, 1984; Voss, Wolf, Lawrence, & Engle, 1991). **Procedural management decisions** (Figure 3, Orasanu & Fischer, 1997).	Cues, options, and outcomes are familiar and can be leveraged by informal heuristics, such as one-reason decision making (looks for a single criterion that discriminates one option from all the rest, Gigerenzer & Selten, 2001); elimination by aspects (successively eliminates options based on criteria in order of their importance, Tversky, 1972); dominance structuring (edits criteria to show a favored option is at least as good as the others on all dimensions, Montgomery, 1993). Choice (Orasanu & Fischer, 1997).	**Mental simulation** (Figure 2, Klein, 1998), **critiquing and correcting stories** (Figure 4, Cohen et al., 1996, 1998, 2001;), expert verification of problem solutions (Patel and Groen, 1991; Larkin, et al., 1980; Larkin, 1981; Chi, Glaser, and Rees, 1982; Serfaty, 1993), and reflection in action (Schön, 1983).
Top-down Analysis (weak)	Decisions in unfamiliar situations that are guided by political ideology, economic theory, religious doctrine, a legal code, or a scientific model (Gunther, 1993; March, Schultz, & Zhou, 2000). Experts override schema-driven processing to "engage a deeper, more basic kind of reasoning from first principles when they need to, for instance, in particularly difficult cases" (Feltovich, 2006, p. 132).	**Rational Choice theory** (Figure 2, Savage, 1954). Decision analysis (a general method for instrumental decisions, which formulates questions about options, outcomes, uncertainties, and preferences in novel problems, von Winterfeldt & Edwards, 1986; Keeney & Raiffa, 1976; Raiffa, 1968). Strategic planning (once popular in large corporations, Mintzberg, 1994). **Scheduling** (Figure 3, Orasanu & Fischer, 1997).	Mixed scanning (Etzioni, 1988); critical thinking (Cohen et al., 2006); and scenario-based planning (Schoemaker, 2007).
Bottom-up Search (weak)	Associative learning, classical conditioning (Rescorla & Wagner, 1972), and sampling event frequencies to acquire probabilistic mental models (Gigerenzer, Hoffrage, & Kleinbolting, 1991), and judgment policies (Fiedler & Juslin, 2006). These involve lengthy experience in a domain, in path-dependent exposure to knowledge; takes a lot of sampling before a decision maker acquires expertise relevant to a broad range of specific problems.	Reinforcement learning or operant conditioning (Herrnstein & Prelec, 1992), and implicit learning (Berry & Dienes, 1993).	Trial and error learning; although quick on a trial by trial basis, it may take a long time (and many errors) to achieve reasonable accuracy. In organizational settings, incremental planning (Lindblom, ,1959; Braybrooke & Lindblom, 1963) and **hedge clipping** (Figure 2, Connolly & Wagner, 1988). Problem solving by searching in a large space of potential problem solutions (Newell & Simon, 1972).

Mixed Top-down, Bottom-up	Creative problem solving (Orasanu & Fischer, 1997). Bottom-up search guided by creative conjectures and extrapolations, as in case-based reasoning and analogical reasoning (Weinreb, 2005, Holyoke & Thagard, 1995; Gentner, Holyoak, & Kokinov, 2001). Expert scientists draw on advanced domain knowledge (a top down element) to search for potentially promising unsolved problems (a bottom- process) and set goals for themselves (Ericsson, 2005). According to research by Ericsson, Krampe, & Tesch-Romer (1993), expertise develops in a process of *deliberate practice*, by means of which experts-to-be seek out increasingly difficult tasks, elicit feedback from teachers and coaches, and analyze their failures to improve performance beyond what was "good enough."

Figure Captions

Figure 1. The TDM framework, showing successful transitions from the initial state associated with each decision mode to commitment to a singular course of action, which is followed by reassessment and/or implementation.

Figure 2. How various research paradigms are represented in the TDM framework, from left to right, top to bottom: the normative core of Rational Choice Theory, Klein's Recognition-Primed Decision Making (RPD) model, Connolly's hedge clipping strategy; Beach's Image Theory.

Figure 3. Representation in TDM framework of five decision types in Orasanu & Fischer's (1997) decision taxonomy. Top left: Condition-action and Go-No Go. Bottom left: Choice. Middle column: Procedural management. Right column: Scheduling. Broken lines represent ways the decision cycle can fail to deliver an actionable commitment.

Figure 4. Story building in the TDM framework (above) and evolution of corresponding story model (below). In story models, lightly outlined rectangles are expectations, planned actions, or hypothesized explanations. Heavily outlined rectangles are events or actions that have been observed. Arrows represent causal influence. X's mark violations of expectations, and dotted lines show how rebuttals are used to explain these failures. Broken lines represent ways the decision cycle can fail to deliver an actionable commitment.

Figure 5. TDM model of action uncertainty. Ovals represent types of uncertainty, and rounded rectangles represent methods. Broken lines represent ways the decision cycle can fail to deliver an actionable commitment.

Figure 6. Increase in degree of success as a function of time or cognitive effort invested, for different uncertainty handling methods.

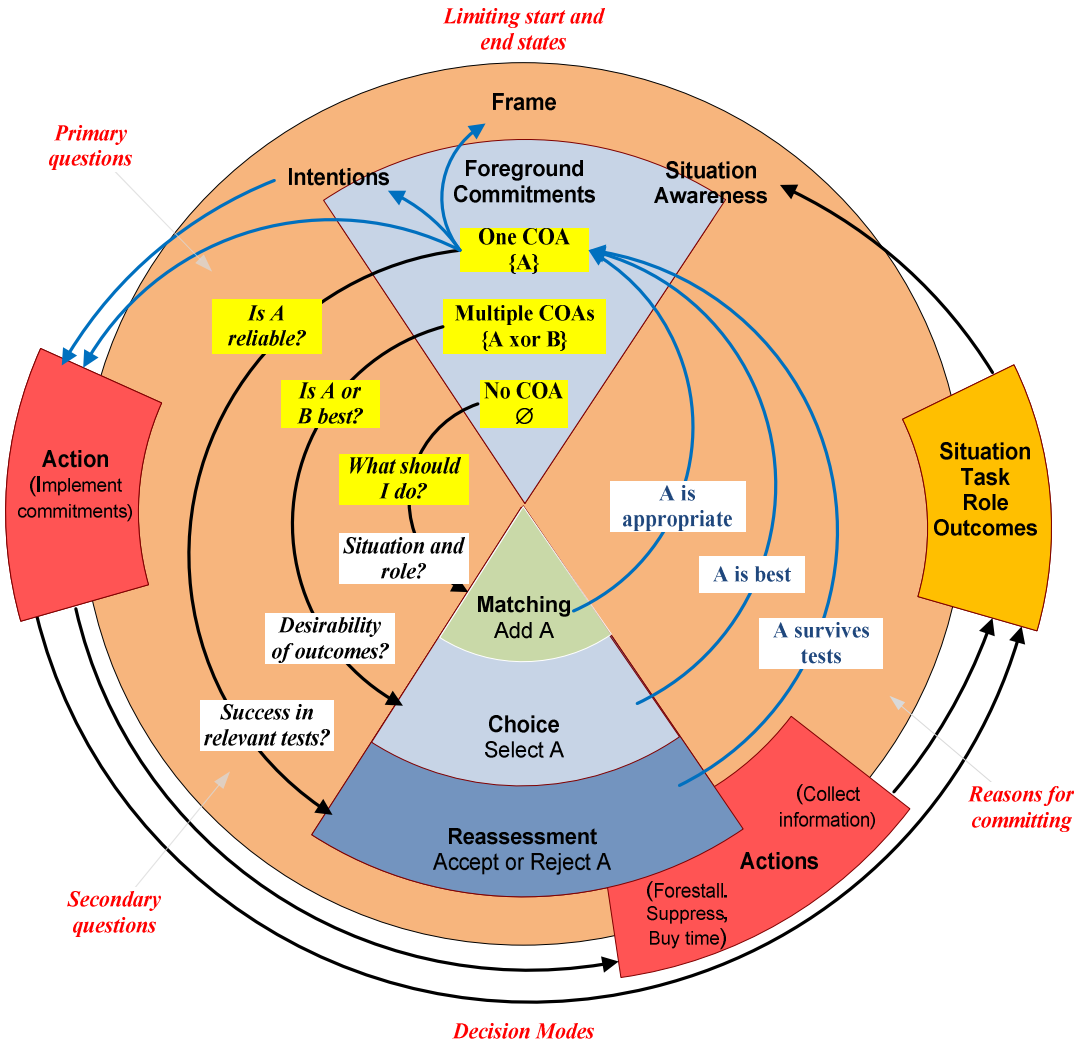
Limiting start and end states
Frame
Primary questions
Intentions
Foreground Commitments
Situation Awareness
One COA {A}
Multiple COAs {A xor B}
No COA ∅
Is A reliable?
Is A or B best?
What should I do?
Situation and role?
Desirability of outcomes?
Success in relevant tests?
Action (Implement commitments)
Situation Task Role Outcomes
A is appropriate
A is best
A survives tests
Matching Add A
Choice Select A
Reassessment Accept or Reject A
(Collect information)
Actions
(Forestall. Suppress, Buy time)
Reasons for committing
Secondary questions
Decision Modes

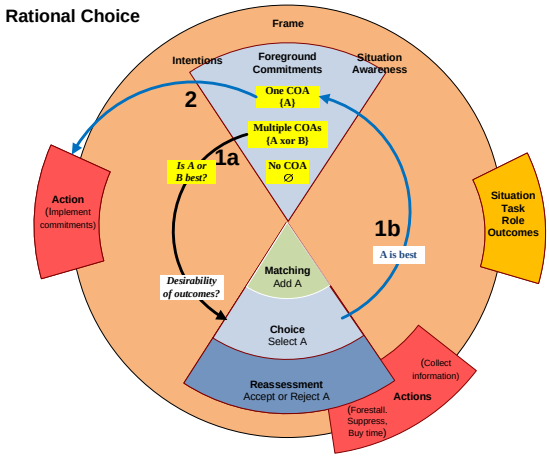
Rational Choice
Frame
Intentions
Foreground Commitments
Situation Awareness
2
One COA {A}
Multiple COAs {A xor B}
1a
Is A or B best?
No COA ∅
1b
A is best
Action (Implement commitments)
Situation Task Role Outcomes
Desirability of outcomes?
Matching Add A
Choice Select A
Reassessment Accept or Reject A
(Collect information)
Actions
(Forestall, Suppress, Buy time)

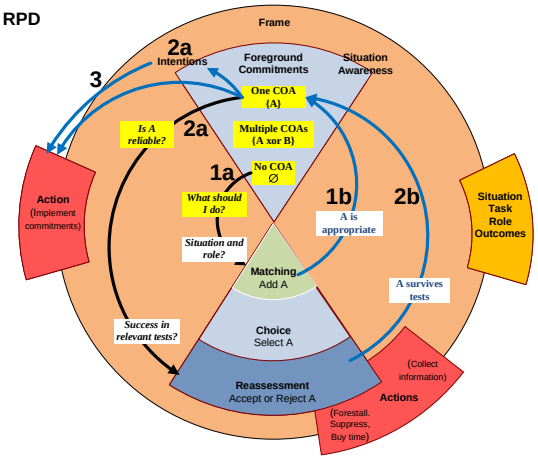
RPD
Frame
2a
Intentions
Foreground Commitments
Situation Awareness
3
One COA {A}
Is A reliable?
2a
Multiple COAs {A xor B}
1a
No COA ∅
What should I do?
1b
A is appropriate
2b
Situation and role?
Action (Implement commitments)
Situation Task Role Outcomes
Matching Add A
A survives tests
Success in relevant tests?
Choice Select A
Reassessment Accept or Reject A
(Collect information)
Actions
(Forestall, Suppress, Buy time)

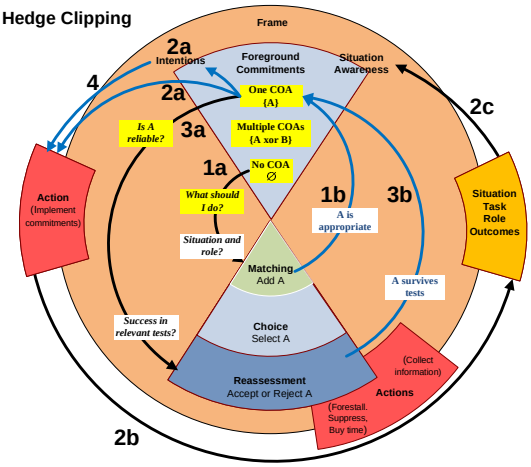
Hedge Clipping
Frame
2a
Intentions
Foreground Commitments
Situation Awareness
4
2a
One COA {A}
2c
Is A reliable?
3a
Multiple COAs {A xor B}
1a
No COA ∅
What should I do?
1b
A is appropriate
3b
Situation and role?
Action (Implement commitments)
Situation Task Role Outcomes
Matching Add A
A survives tests
Success in relevant tests?
Choice Select A
Reassessment Accept or Reject A
(Collect information)
Actions
(Forestall, Suppress, Buy time)
2b

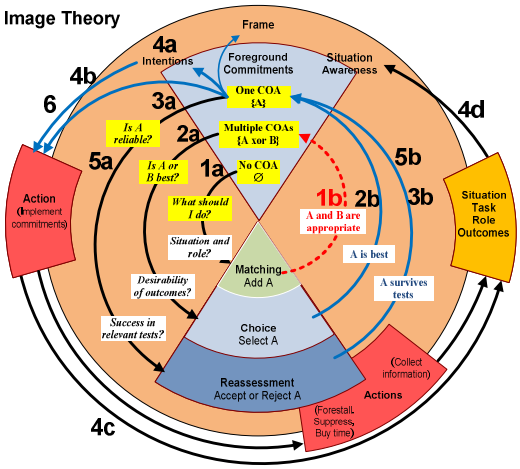
Image Theory
Frame
4a
Intentions
Foreground Commitments
Situation Awareness
4b
6
3a
One COA {A}
4d
Is A reliable?
2a
Multiple COAs {A xor B}
5a
5b
Is A or B best?
1a
No COA ∅
1b
A and B are appropriate
2b
3b
What should I do?
Situation and role?
A is best
Action (Implement commitments)
Situation Task Role Outcomes
Desirability of outcomes?
Matching Add A
A survives tests
Success in relevant tests?
Choice Select A
Reassessment Accept or Reject A
(Collect information)
Actions
(Forestall, Suppress, Buy time)
4c

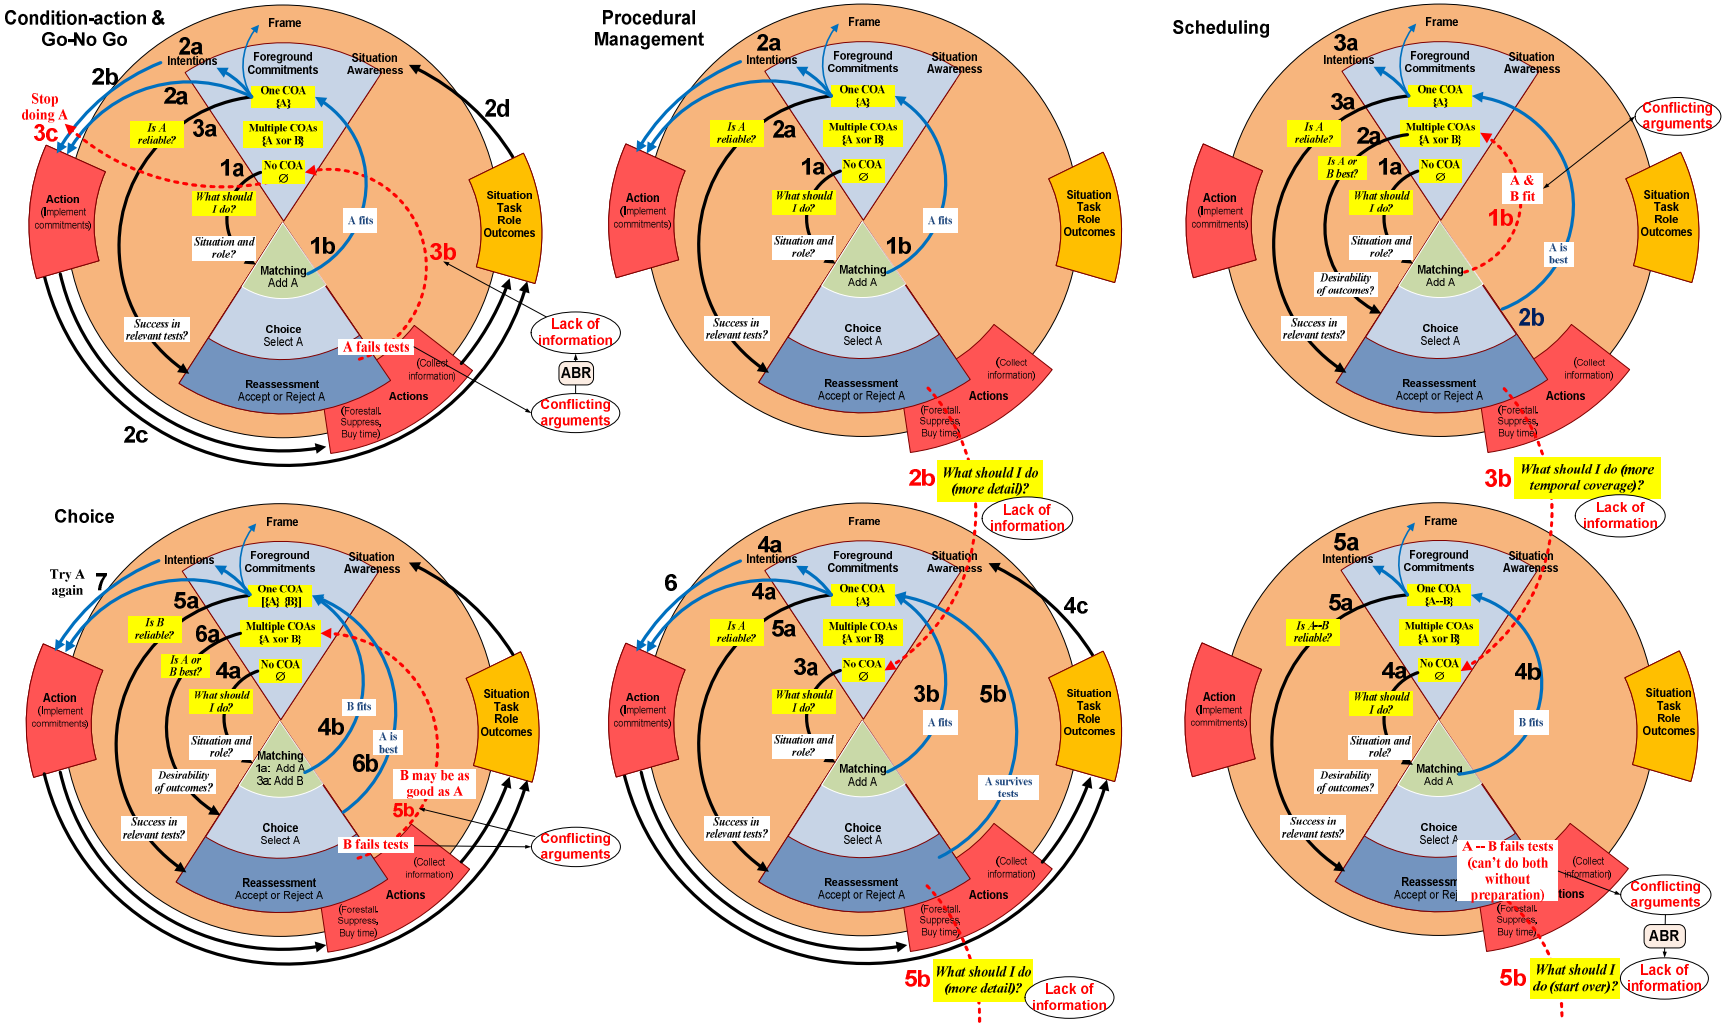
Condition-action & Go-No Go
Frame
Intentions
Foreground Commitments
Situation Awareness
One COA {A}
Multiple COAs {A xor B}
No COA ∅
Is A reliable?
What should I do?
Situation and role?
Matching Add A
Choice Select A
Reassessment Accept or Reject A
Actions
(Forestall, Suppress, Buy time)
(Collect information)
Action (Implement commitments)
Situation Task Role Outcomes
Success in relevant tests?
A fits
A fails tests
Stop doing A
Lack of information
ABR
Conflicting arguments
Procedural Management
2b What should I do (more detail)?
Lack of information
Scheduling
Conflicting arguments
A & B fit
A is best
3b What should I do (more temporal coverage)?
Lack of information
Choice
Try A again
Is B reliable?
Is A or B best?
Desirability of outcomes?
1a: Add A
3a: Add B
B fits
B may be as good as A
B fails tests
Conflicting arguments
Multiple COAs {[A} {B}]
6
4c
A survives tests
5b What should I do (more detail)?
Lack of information
A -- B fails tests (can't do both without preparation)
Is A--B reliable?
5b What should I do (start over)?
Lack of information

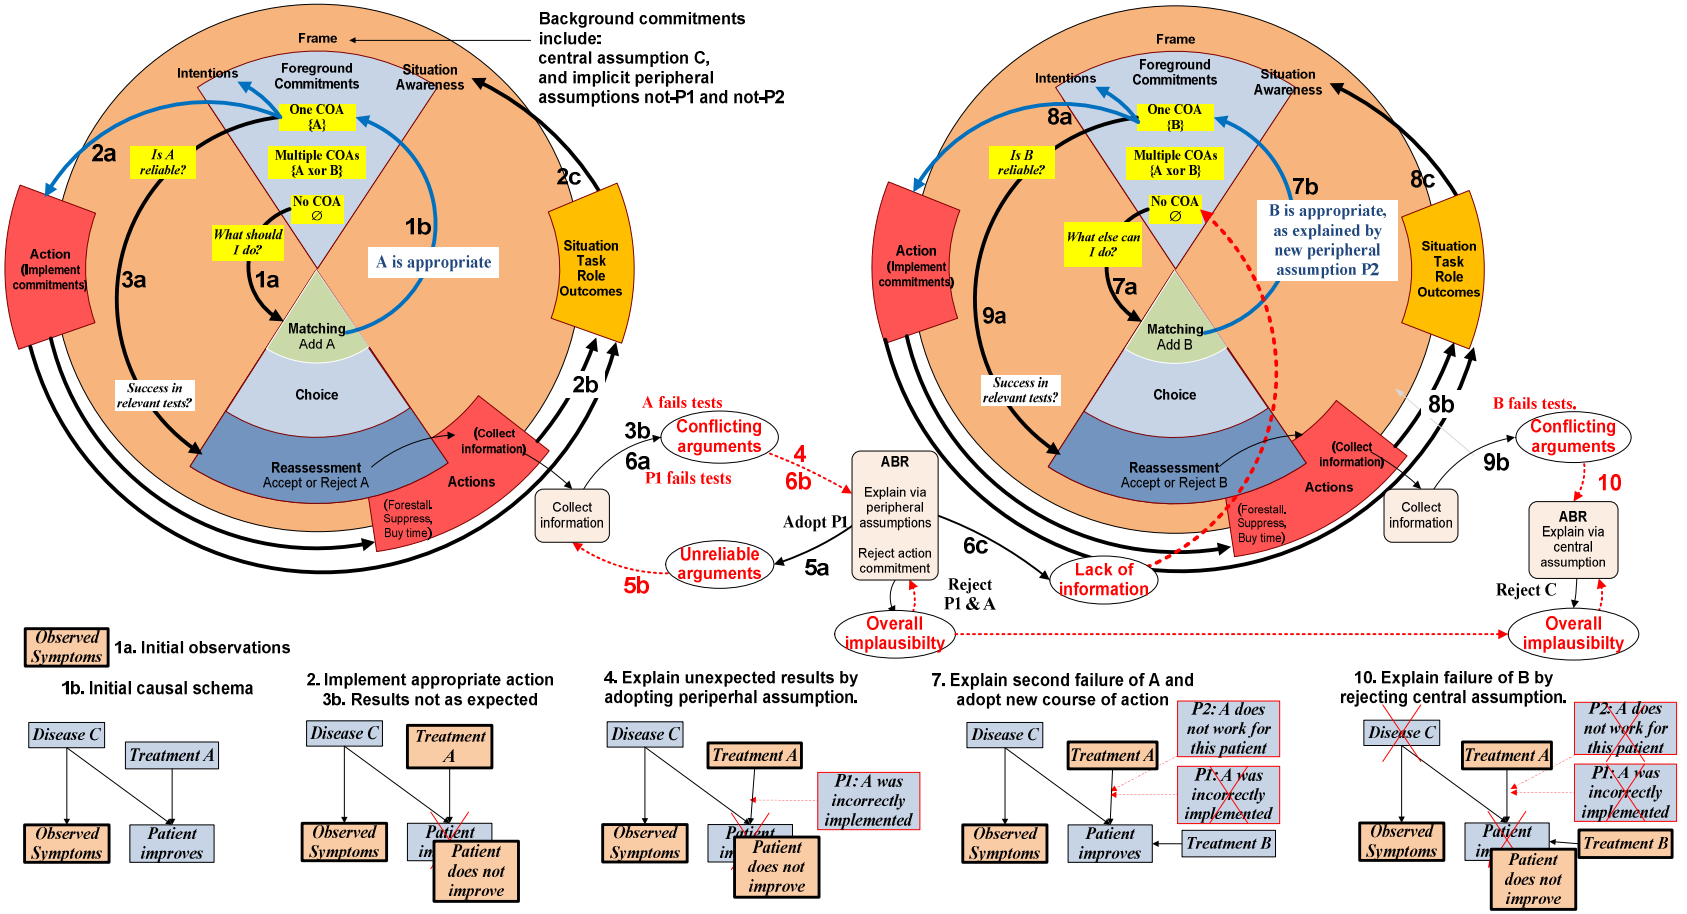
Background commitments include:
central assumption C,
and implicit peripheral assumptions not-P1 and not-P2
Frame
Intentions
Foreground Commitments
Situation Awareness
One COA {A}
Multiple COAs {A xor B}
No COA ∅
Is A reliable?
What should I do?
A is appropriate
Action (Implement commitments)
Situation Task Role Outcomes
Matching Add A
Choice
Reassessment Accept or Reject A
Actions
(Collect information)
(Forestall, Suppress, Buy time)
Success in relevant tests?
1a
1b
2a
2b
2c
3a
A fails tests
3b
Conflicting arguments
6a
P1 fails tests
4
6b
ABR
Explain via peripheral assumptions
Reject action commitment
Adopt P1
5a
Unreliable arguments
5b
Collect information
6c
Lack of information
Reject P1 & A
Overall implausibilty
Is B reliable?
What else can I do?
B is appropriate, as explained by new peripheral assumption P2
Matching Add B
Reassessment Accept or Reject B
7a
7b
8a
8b
8c
9a
B fails tests.
9b
10
ABR
Explain via central assumption
Reject C
Observed Symptoms
1a. Initial observations
1b. Initial causal schema
Disease C
Treatment A
Patient improves
2. Implement appropriate action
3b. Results not as expected
Patient does not improve
4. Explain unexpected results by adopting peripherhal assumption.
P1: A was incorrectly implemented
7. Explain second failure of A and adopt new course of action
P2: A does not work for this patient
Treatment B
10. Explain failure of B by rejecting central assumption.

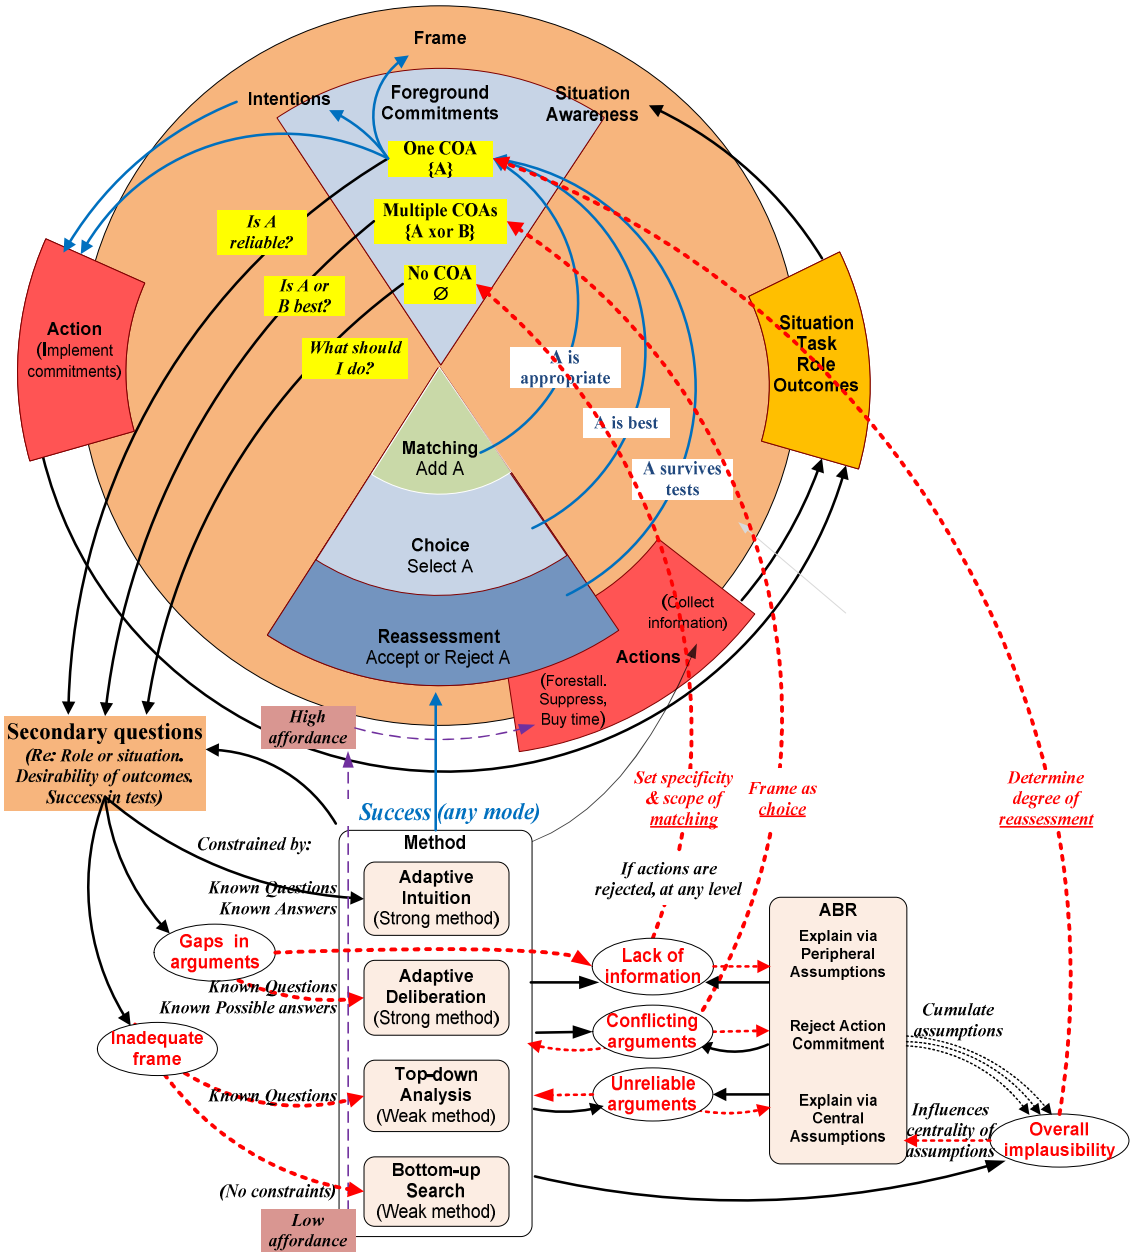
Frame
Intentions
Foreground Commitments
Situation Awareness
One COA {A}
Multiple COAs {A xor B}
No COA ∅
Is A reliable?
Is A or B best?
What should I do?
A is appropriate
A is best
A survives tests
Action (Implement commitments)
Situation Task Role Outcomes
Matching Add A
Choice Select A
Reassessment Accept or Reject A
(Collect information)
Actions
(Forestall. Suppress, Buy time)
Secondary questions (Re: Role or situation. Desirability of outcomes. Success in tests)
High affordance
Success (any mode)
Set specificity & scope of matching
Frame as choice
Determine degree of reassessment
Constrained by:
Known Questions
Known Answers
Known Questions
Known Possible answers
Known Questions
(No constraints)
Low affordance
Method
Adaptive Intuition (Strong method)
Adaptive Deliberation (Strong method)
Top-down Analysis (Weak method)
Bottom-up Search (Weak method)
Gaps in arguments
Inadequate frame
If actions are rejected, at any level
Lack of information
Conflicting arguments
Unreliable arguments
ABR
Explain via Peripheral Assumptions
Reject Action Commitment
Explain via Central Assumptions
Cumulate assumptions
Influences centrality of assumptions
Overall implausibility

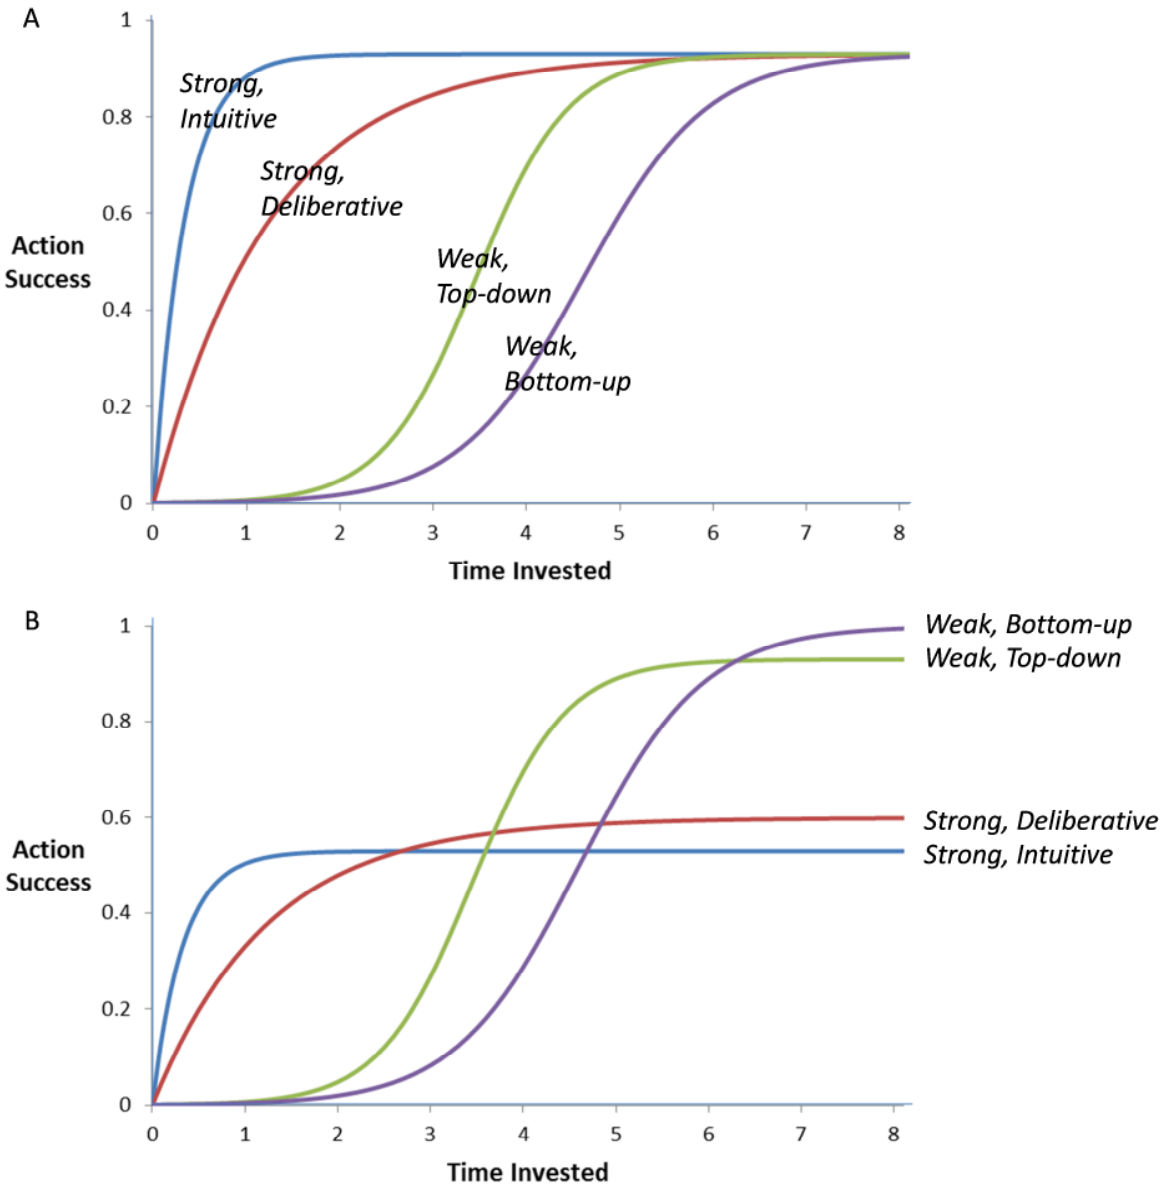
A
Action Success
Time Invested
Strong, Intuitive
Strong, Deliberative
Weak, Top-down
Weak, Bottom-up
B
Action Success
Time Invested
Weak, Bottom-up
Weak, Top-down
Strong, Deliberative
Strong, Intuitive