Business intelligence as a prism: how three theories of decision-making are realized as a spectrum of practices

Business intelligence (BI), as a professional discipline, prescribes organizational and technological interventions aimed at providing timely, accurate information to the right people, at the right times, to enable and improve business decision-making (Watson, 2009). It falls firmly within the domain of decision support systems (DSS), and scarcely an academic paper or vendor sales-pitch fails to mention BI's power to impact decision-making (Russell et al., 2010; Imhoff & White, 2008). Despite the volume of ink spilled on BI, though, we still have little understanding of how its tools and strategies affect decisions, because most of the literature has a data-centric or technology-centric focus which takes decision-making for granted. Given this blind spot, it is no wonder that management and information systems researchers are asking why we aren't making better decisions already (Davenport, 2009).

One of the challenges in taking a decision-centric view of business intelligence is the embattled state of decision-making theories themselves. While the perfectly-informed, rational decision maker (or "economic man") is still held up as a normative ideal in some economics theorizing, organization theorists have for decades chipped away at it, charting the limits of information and rationality, and ultimately questioning the whether intentional decision-making even occurs. In the "garbage can" model of organizational choice, for example, decisions "happen" as collisions between problems looking for solutions, solutions looking for problems, and people looking for occasions to make choices (Cohen, March, & Olsen, 1972). Behavioral science has won a pyrrhic victory over economics, piling up reasons why rational choice is impossible, but providing no useful guidance to fill the gap. Cabantous and Gond (2010) argue that we can thread the needle between the assumption of rationality, and the despair of its impossibility, by studying how managers *informed by normative theory* effortfully engineer rational *decision-making prescriptions* in organizations.

We believe it would be fruitful to view business intelligence as a practice of *decision engineering*. This view bridges the chasm between *normative* theories of how decisions *should* be made, and the actionable, prescriptive theories-in-use applied by practitioners, by conceiving of BI managers and developers as "design theorists" who turn theory into sociotechnical reality. We hold that good theories of decision-making are performative, meaning that they have an impact on the world of practice , and we highlight the role that academic researchers play as producers of the normative theories that inform decision engineers. The subjects of our study are, therefore, the ways in which BI practitioners *think* about decision-making, and how their thinking *shapes* the kinds of interventions they make in the name of business intelligence.

Cabantous and Gond (2010) consider only one theory, the theory of rational choice, as a performative norm that managers attempt to engineer into organizational reality, but we believe this is too great of a simplification. We know that competing theories abound in the field of information systems, often contradicting one another, and as Sarker and Lee (2002) found in the practice of business process redesign, they are transformed by practitioners into diverse prescriptions of varying merit. We assume that there is no single "best" theory of decision making, but that different theories may apply according to the goals and constraints of the decision engineering context. In this we take inspiration from Graham Allison's famous study of the Cuban missile crisis (Allison & Zelikow, 1999), which showed that the "rationality" of a decision depends upon the lens through which one views it. A decision that doesn't seem rational in terms of short-term goals, for example, may reflect a rational organizing principle, or a rational political calculus; moreover, all three norms may be at work in the same decision. As a prism divides white light into its hidden component colors, the Allison study teased apart the narrative of the crisis to reveal three equally-valid theories of decision-making that *together* explained the events.

Similarly, the objective of our research is to separate and make visible the multiple and complementary theories of decision-making that, through the decision engineering process, become transformed into the prescriptive theories-in-use of business intelligence. As Information Systems is a science of the artificial, concerned with IT artifacts and business systems (Gregor, 2006, 2009), our goal and approach differs somewhat from Allison's. Whereas his study explained how multiple logics contributed to an array of decisions, choices, and actions taken, we wish to understand how multiple decision logics contribute to a wide set of technological and organizational *prescriptions* for improving decision-making outcomes. To do so, we focus on the decision engineering process that connects them. Our conceptual model is illustrated by Figure 1.

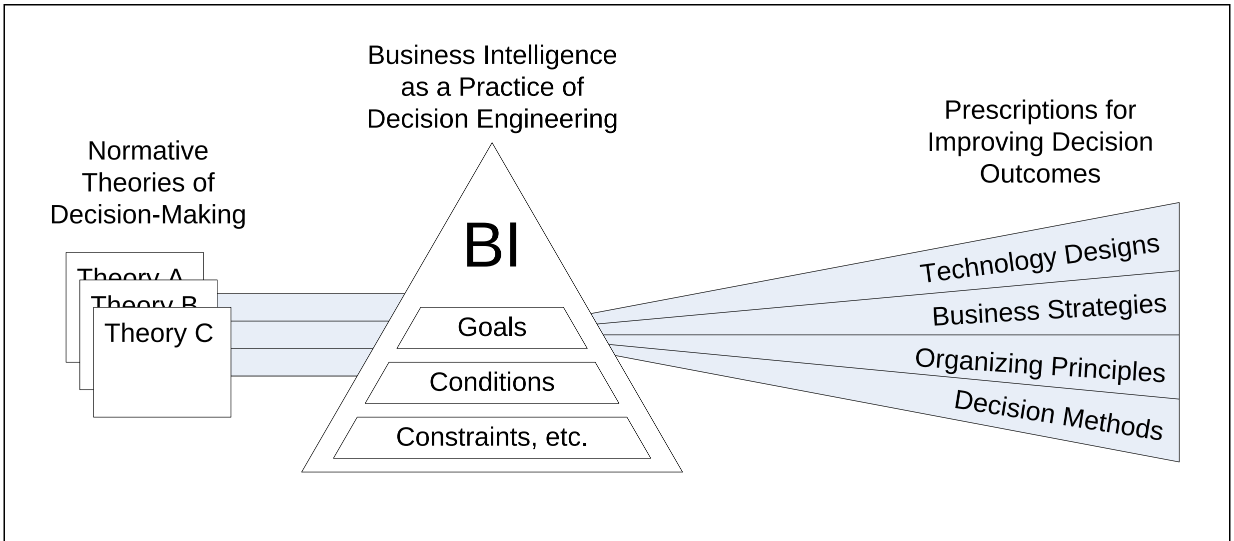

Figure 1. Business intelligence as a "prism" linking decision theory to BI prescriptions.

From normative and descriptive theory to prescriptive theory-in-use

This challenge requires us to consider the interaction between normative, descriptive, and prescriptive theories of decision-making. As Bell, Raiffa, and Tversky (1988) observed, these types of theory are generally separated by disciplinary walls: mathematical economists are concerned with normative theories of how people *should* make decisions according to fundamental principles, psychologists and behavioral scientists study how people *do* make decisions, rational or otherwise.

> *"But there is a third interest group, the methodologists, the consultants, call them what you will. Some of us are concerned with the bottom line: how do you improve the quality of decisions in practice? It is one thing to talk of axioms and proofs and paradoxes and cognitive limitations – but how can you really help?"* (Bell et al 1988, *p. ix)*

The Information Systems discipline, along with operations research and other sciences of the artificial (Simon 1996), is concerned with organizational interventions (such as information technologies) that help our subjects strive for the normative "should" while coping with the descriptive "do". According to our field's philosophy of science (Gregor, 2006), we approach the prescriptive in a couple of different ways: on the one hand, we generate (design) prescriptions ("recipes") for building specific IT solutions to specific business problems, evaluating them in the normative terms of utility and originality; on the other, we create descriptive theories *about* prescriptions, in what is called "behavioral" or "economics" IS research of the phenomena that occur around purposeful IT artifacts. These two modes of knowledge creation are fundamental to our field and seen as mutually informing (Gregor, 2009; Hevner et al., 2004).

This paradigm unfortunately skips over the question of how descriptive and normative theory *leads into* prescriptions, treating design as more of an art than a science. To answer it, we need a conceptual model that incorporates descriptive or normative theories as inputs, and prescriptions (theoretical or otherwise) as outputs. The lens of *design theory* is one such model, explaining how IS researchers draw on *kernel theories* from natural and social sciences, using a logic of procedural rationality, to transform them into formalized meta-designs for classes of information systems (Gregor, 2007; Walls, Widmeyer, & El Sawy, 1992, 2004). We also see research that views practitioners as design theorists of a sort – such as the already-mentioned Sarker and Lee (2002) case study in which three theories of business process redesign were translated into different types of prescriptions. These "theories-in-use" (Argyris & Schön, 1978) differ from design theory in that they are informal, possibly implicit, and their roots in formal theory may be unknown to their users. They make claims about effective action, not about objective truth.

The main weakness shared by the design theory and theory-in-use perspectives is that they take the normative/descriptive kernel theories as givens, which prevents us from hypothesizing about how practice might evolve differently if different theories predominated – as might be the case when comparing practices across cultures. Following Gregor (2006, 2007), we need to put kernel theories into their proper perspective, as abstract, man-made artifacts belonging to Popper's "World 3". It is important to view these theories as contingencies upon which the prescriptive process could hinge. Another problem with the two perspectives is that while the design theory view attributes all the work of prescriptive theorizing to academic researchers, the theories-in-use perspective assumes practitioners proceed by trial-and-error in a kind of theory vacuum.

Neither perspective has room for an entire discipline, full of back-and-forth between research and practice, as is the case in DSS and business intelligence.

We believe that the Cabantous and Gond (2010) framework is an excellent blueprint for a theory that overcomes these weaknesses. They discuss how the normative theory of rational choice, created by academics, known to practitioners, becomes social reality as it is embedded and engineered ("performed") into the activities, conventions, and tools of decision-making practice. The focus of their perspective is neither on the kernel theory nor on the prescriptions, but rather on the activities that connect them. Their model is reproduced in Figure 2.

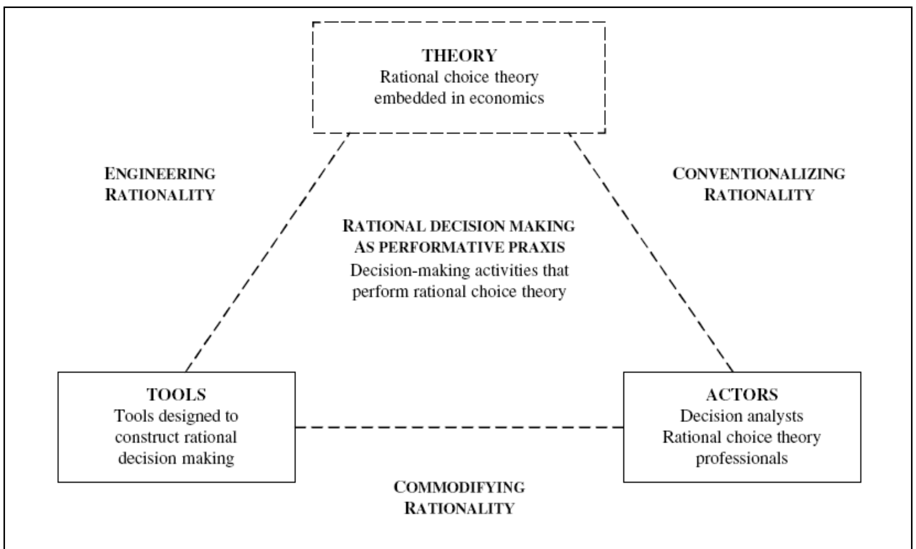

Figure 2. Conceptual model of Cabantous and Gond (2010)

Three archetypes of decision engineering

Our main extension of the Cabantous and Gond framework is to suggest that there may be more than one "wavelength" within the "white light" of normative decision theory, and that actors with different goals, assumptions, and constraints, may "perform" different types of rationality.

Allison and Zelikow (1999) demonstrate that three distinct theories of decision-making, erstwhile competitors, are best seen as co-present and complementary in explaining a narrative of decisions made and taken in a real-world scenario. We, too, see three kernel theories as complementary influences on business intelligence practice – each theory speaking to different parts of the academic and practitioner audience, and contributing to different performative practices.

The kernel theories we consider are: bounded rational choice, in which decision-makers select from among alternative choices to the best of their abilities given limited information, time, and attention; distributed cognition, in which decisions result from the interdependent behavior of individuals within a group context characterized by culture, processes, and cognitive tools like IT; and organized anarchies, in which decisions emerge as people and groups with markedly different interests, abilities, and political options interact, bargain, and compete. This is not an exhaustive list of all *possible* theories of decision-making rationality, rather, these are three theories that are readily evident in academic and practical business intelligence literature today. They are different from, but can trace their lineage to, the three models that Allison identified in 1971. If other understandings of decision-making rationality held sway – as they have in the past, and will in the future – we would expect different prescriptions to drive practice.

As each theory is characterized by different assumptions, and focuses on different kinds of actors and behaviors, we believe that each is received by a different audience with its own goals, constraints, and levers of control, and attracts the interest of a different set of academics as well. Thus a few theories that are all "rational" lead to a broad spectrum of practical prescriptions and research questions after passing through the decision engineering prism that is the business intelligence discipline. We now consider the three performative practices as archetypes of decision engineering: decision development, capability crafting, and issue illumination.

Decision development

Kernel theory

The main approach to understanding how people make decisions in business, under uncertainty, has been to identify their limitations – finite time, incomplete information, limited attention and computational abilities – and to research the techniques that allow decision-makers to cope with these limitations. Since at least the 1950s this research has been advanced under the rubric of *bounded rationality* because it was initially proffered as a replacement for the untenable assumptions of perfect information and perfect rationality that underpinned classical and neoclassical economics (Simon, 1979). It includes descriptive theories in fields as diverse as economics, psychology, and organization theory, and prescriptive traditions in operations research and decision support systems. The common thread is a shared interest in *how* decision-makers choose between alternatives, as opposed to the classical economist's focus on *what* they choose.

Central to a theory of bounded rationality is the concept of *search*. Decision-makers may not know all their alternative choices, and even if they do, they have only incomplete information about each alternative and its consequences, so they engage in active search for alternatives and for information. This search is guided by *heuristics*, a set of rules about how to search, when to stop searching ("stopping rules"), and how to choose from among the alternatives based on the information. The term "heuristics" has been associated with "rules of thumb" and with "biases" (Tversky & Kahneman, 1974), suggesting pessimistically that bounded rationality is merely a handicapped version of some economic principle like utility maximization, but it can be seen as a normative theory in terms of *procedural rationality* (Simon, 1979), the *logic of appropriateness* (Cabantous & Gond, 2010), or *ecological rationality* (Todd & Gigerenzer, 2000).

Todd and Gigerenzer (2000) argue that a generalized decision criterion like utility maximization is not normative, not only because it is impossible, but because our attempts to achieve it are destructive of the ends decision-makers care about: making the right choices under uncertainty with limited time and resources. Heuristics constitute an “adaptive tool kit” of computational strategies that allow decision-makers to exploit the structures of their information environments to make fast and accurate decisions with limited time and effort. Each decision domain may have a different information environment, so multiple heuristic strategies are needed. For example, a domain characterized by perishable alternatives that arrive one at a time, such as potential mates or potential investments, calls for a “satisficing” heuristic in which the decision maker sets an aspiration level of some criterion and continues the search for alternatives until one that exceeds the aspiration level is found, then chooses that one.

Because of their environmental specificity, heuristics are adaptive rather than optimizing. Because of their simplicity, however, they are robust to environmental change. Attempts to achieve the classical economic norm of rationality, by applying massive amounts of data and computational power to a problem, could easily become *less* optimal by overfitting its statistical models to unimportant variables in historical data. The normative theory of bounded rationality, therefore, requires that we understand the information structure of the problem environment(s), and utilize heuristic rules (for guiding search, stopping search, and choosing an alternative) that optimally balance accuracy, frugality, and speed of decision making. Different rules may apply to different decisions.

Assumptions and focus of kernel theory

Theories of bounded rationality share three main assumptions: we are given a specific decision-maker, faced with a specific decision (or type of decision), for which a finite set of alternatives exist, though they may not be known to the decision-maker. Although these assumptions may not seem too bold, scholars familiar with decision-making literature know that in many real-life situations it isn't certain what is to be decided, when, or by whom, from what set of options. Nevertheless, the assumptions of bounded rationality apply to many realistic decision-making situations, such as the casting of a vote, the purchase of a car, the selection of an employee, and so on.

The focus of the bounded rationality lens is not the who, the what, or the when, but the *how* of decision-making. What rules and heuristic strategies best guide decision-makers as they search for alternatives, make inferences about the relative merits and consequences of each choice, and ultimately stop searching and make a choice?

Community of practice

As bounded rationality takes decisions and decision-makers as givens, its perspective most likely resonates with DSS developers and BI consultants who are tasked with supporting a decision-making process that is given to them. As mentioned above, there are many real-world situations in which these things are fixed – for example, a banker judging a loan application, a project manager selecting a contractor for a job, or a doctor diagnosing a patient's illness – and the BI practitioner's job is to create and improve the tools and methods for gathering information, evaluating alternatives, and making the decision. By contrast, this community of practice probably excludes high-level managers like CIOs who are more concerned with organizational matters than with the details of specific decision support applications. For that reason, we feel that the label *decision development* best captures this first archetype of decision engineering.

DSS research questions

Academic researchers participate in the decision development practice by improving our understanding of decision types (what Todd & Gigerenzer, 2000 call "the information structure of the environments"), and by developing, prescribing, and testing heuristic strategies for different types of environments. They combine what we know about the human mind, and human-computer interaction, with our practical knowledge of technology, to make and test prescriptive theories about IT tools that enable new heuristics, or enhance their performance for particular decision types. A great deal of DSS research since the early days of the field has joined this conversation. Here are just a few examples of relevant research questions:

- How does the technology for data visualization affect decision performance, or not affect it, depending on the task environment and decision type? (Dickson, DeSanctis, & McBride, 1986)
- How do decision support systems reduce the cognitive effort associated with particular decision-making strategies, to enable decision-makers to employ those strategies instead of others? (Todd & Benbasat, 1994)
- What are the best "matches" between the characteristics of decisions and the characteristics of DSS technology? (Goodhue & Thompson, 1995)
- How does the provision of long-term historical data in a marketing data warehouse affect decision-makers' performance on a market prioritization decision? (Park, 2006)

- How can DSS be intentionally designed to reduce logical and mechanical errors in DAMCDA (discrete alternative multicriteria decision analysis) decision-making? (Williams et al., 2007)

Example prescriptions

Description development provides a number of prescriptions; some highly technical, such as new decision-making algorithms from statistics and computer science, and others more accessible to laymen, such as the claim that a balanced scorecard (BSC) is an effective tool for balancing the organization's priorities in strategic decision making (Kaplan & Norton, 1996). At any DSS industry exposition one encounters prescriptions for new types of data visualizations, new types of analytical tools, real-time systems, and dashboards, each intended to support a specific type of decision-making task better than its competitors. Any prescription that deals with search for information, analysis of alternatives, or methods for choosing, is part of decision development.

Capability crafting

Kernel theory

Bounded rationality has good explanatory and normative value when decisions and decision-makers are well-defined, but in research and practice we are often challenged with more complicated phenomena. Real-world enterprises are social systems composed of interdependent parts, but greater than the sums of their parts, and instead of discrete decisions they face "messes" made up of multiple, interrelated problems (Ackoff, 1994). To cope with such messes we need to be able to take a broader view of the context within which simple decisions are made. The theory of *distributed cognition* provides a lens for such inquiry. Initially developed from ethnographic studies of airplane cockpits and ship crews, this theory views teams and organizations as *cognitive systems* that produce cognitive outputs such as decisions (Hutchins, 1995a, 1995b; Weick & Roberts, 1993). As with individual cognition, group cognition involves the acquisition and interpretation of information, the storage and retrieval of memory, the formation of judgments and choices, and the motivation of action. The difference is that in distributed cognition, these activities take place in sociotechnical systems of people, information technology, rules and routines (Hollan, Hutchins, & Kirsh, 2000). As a result, the cognitive processes are easier to observe, and possibly subject to a different set of conditions and influences than individual cognition.

Some of the core concepts of a theory of distributed cognition are the technologies (*work materials*) that serve functions analogous to memory or analysis, the human actors who interact with the technologies and with each other, and the organization of the cognitive system: "Cognitive processes involve trajectories of information (transmission and transformation), so the patterns of these information trajectories, if stable, reflect some underlying *cognitive architecture*." (Hollan et al., 2000, p. 177). Another major feature of these theories is attention to the *context* within which distributed cognition takes place. Cognitive systems are part of larger systems, such as the enterprise, and they draw on organizational *culture*, which can be viewed as a language for interpreting information and a tool kit of decision frameworks. As such, culture both facilitates and constrains group cognition.

A related theoretical idea is that of capability development. (Ethiraj et al., 2005). Capabilities are routines or methods that a group develops and maintains at some cost of time or money, and that it can reproduce when necessary to solve a particular type of problem. When we're talking about cognitive tasks like decision making, some relevant capabilities are knowledge capabilities (Freeze & Kulkarni, 2007), sensing and responding capabilities (Overby, Bharadwaj, &

Sambamurthy, 2006), absorptive capacity (Malhotra, Gosain, & El Sawy, 2005), and transactive memory (Wegner, 1986). Like organizational culture, a system's capability set enables it to carry out some types of decision-making very efficiently, but at the expense of some flexibility as new types of information or decisions may not fit perfectly with established cognitive routines. A cognitive system may also develop dynamic capabilities and improvisational capabilities that help it modify or work around its existing capability set as the environment changes (Pavlou & El Sawy, 2010). The mix of capabilities clearly depends on the demands of the task environment. We begin to move into normative territory.

In a normative theory of distributed cognition, the manager or architect of a sociotechnical system is expected to engineer (Upton et al., 2008) or *orchestrate* (Lovallo & Sibony, 2010) the whole system. He or she assigns people to roles, decides what will be decided, and by whom, determines the resources and technologies (work materials) that they will use, and architects the decision-making routines that make up the system's capability set, all while managing the interdependencies of system parts. The normative ideal is one in which the cognitive system is designed in an optimal way for the kinds of tasks it faces. One does not "satisfice" or even "optimize" the system to deal with a particular problem, one *redesigns* the system so that the problem no longer exists, or is transformed into something more tractable (Ackoff, 1994).

Assumptions and focus of kernel theory

Unlike bounded rationality, the distributed cognition approach doesn't take the identity of decision-makers or even the specific type of decision as a given, instead, it assumes that cognitive architectures are contingent and malleable. What they do generally assume is that some sort of a manager can perceive the system, has some lever(s) of control over it, and cares about its performance according to an objective standard. If a distributed cognition study features neither a manager nor a goal, then it is a descriptive ethnography and not a normative or prescriptive theory.

A second major assumption, which underpins the concepts of cognitive architecture and routine capabilities, is that of stability, or *gradual* change. It serves no purpose to study the architecture of a cognitive group if it is a unique one-time collaboration, and social-cognitive capabilities make little sense if problem characteristics, information sources, and objective goals are changing radically and unpredictably.

Actors, technologies, and routines feature in the distributed cognition perspective, but there is a special focus on the *interactions* between them, and the decisions that result from these interactions. For example, Michel (2007) compared two organizations and showed that their different methods of assigning people to tasks (a work process) created different levels of uncertainty (a task characteristic) which caused workers to develop different information-sharing routines (a capability) resulting in different types of decision-making (a cognitive outcome).

Community of practice

The distributed cognition lens is mere sophistry to practitioners who have neither the opportunity to perceive an entire system nor the authority to change it, so it may not find a receptive audience among *decision developers*, as described in the previous section. Instead, distributed cognition is most relevant to managers, whether of small groups or large ones, to CIOs and CEOs and entrepreneurs, and perhaps to the data architects and consultants who also take a big-picture view of all the decision-making activity in a department, process, or enterprise. These practitioners are interested not in designing specific tools, but in orchestrating human resources and IT systems into efficient, effective, and adaptive groups. They may have the opportunity to redesign entire business processes from scratch (Sarker & Lee, 2002), but often they pull smaller levers, for example introducing a new technique to confront decision bias (Lovallo & Sibony, 2010), or auditing and intervening in one decision-making process at a time (Davenport, 2009). In this practice, the enterprise is a kind of cognitive machine, whose parts are capabilities, and practitioners are more often its mechanics than its designers, so we call the practice *capability crafting*.

DSS research questions

Information systems researchers have much to contribute to the practice of capability crafting. As the only major discipline that studies systems in organizations (Alter, 2003), we can contribute generally by developing frameworks and theories for mapping out and intervening in complex and interdependent systems. More specifically, we have developed and tested theories about methods of process management and process redesign, we have studied technologies that perform cognitive roles – such as organizational memory IS (Stein & Zwass, 1995) and knowledge management systems (Alavi & Leidner, 2001) – and we have studied how individual technologies enhance social-cognitive capabilities like coordination and communication. Some examples of research questions that may inform capability crafting are listed here:

- How should we design a “distributed cognition IS” that enables users to represent, share, discuss, and act on their different understandings of an issue? (Boland, Tenkasi, & Te’eni, 1994).
- What are the different kinds of knowledge capabilities, how is each one leveraged by the organization, and which technologies enable them? (Freeze & Kulkarni, 2007)

- What constitutes an effective methodology for mapping out the "distributed cognition configuration" of a complex sociotechnical enterprise, and modeling the effects of organizational change on that configuration? (Upton et al., 2008)
- How do knowledge management activities moderate the effect of business intelligence technologies on decision-making? Can "knowledge exchange protocols" help integrate BI into business processes? (Herschel & Yermish, 2009)
- Do business intelligence (BI) systems enhance organizational coordination mechanisms? (Ferrari & Rossignoli, 2010)

Example prescriptions

Prescriptions abound in practitioner literature, hailing the benefits of new analytics-based strategies (Davenport, 2006) or re-engineering all the decisions in an organization (Davenport, 2009). They prescribe changes to decision techniques, decision ownership, decision criteria, and decision-maker incentives. There is no shortage of software vendors or consultants who will advise starting over from scratch with entirely new organizational information systems to implement new and better cognitive strategies. If there is a weakness with these prescriptions, it is that they assume near-omnipotence and a great deal of money besides, and are consequently more relevant to CEOs than to ordinary managers, or harried CIOs, who are simply trying to work with the systems they're given. Perhaps we need more prescriptions for levers of change that can be implemented one at a time to "tune up" cognitive systems as business goes on.

Issue illumination

Kernel theory

The curiosity of decision theorists and the vexations of decision-makers have led to a third line of inquiry into those problems that are neither well-defined nor stable and controllable. These are the “wicked problems” (Rittel & Webber, 1973). Among their characteristics are: the involvement of many stakeholders with different goals and priorities; complex and tangled interdependencies; uniqueness that flummoxes routines and precedents; and the need for ongoing effort as the problem can never really be ended (Camillus, 2008). We use the term *organized anarchies* for the family of theories about how people grapple with these apparently-intractable problems. Cohen et al (1972) simulate organized anarchies with a “garbage can model” in which problems looking for solutions, solutions looking for problems, and people looking for decision-making opportunities, bounce around randomly until they come together by happenstance, resulting in a decision. However, their description of these anarchies as “organized” is an acknowledgement that we can at least try to “tract” these apparently-intractable problems.

Organized anarchies are characterized by a particular type of organizational, industry, or technological environment in which systems are so complex and dynamic that it is hard for any individual to perceive its cause-effect structure (Weick, 1990). These environments are *turbulent* in the sense they are not regulated by any equilibrium-seeking forces, so small actions and one-time decisions can compound unpredictably into large system shocks (Emery & Trist, 1965). Tight coupling between systems and dependencies that cross levels of analysis (Selsky, Goes, & Baburoglu, 2007) mean that theories of organized anarchies must connect phenomena at the individual and group levels with phenomena at the organization and industry levels in order to make sense of them. The theoretical term *sensemaking* (Weick, 1990) describes how we try to understand and solve these problems.

Normative theories for organized anarchies are less developed than those for bounded rationality and distributed cognition, but they do exist. Emery and Trist (1965) argued that the most important implication of *turbulent fields* was that the fortunes of all participants are positively correlated, so their participants (firms, in their view) should find ways to compete with one another that do not harm the overall health of the field. They recommended the collective establishment of *social values* to stabilize and transform the turbulence into order. This is similar to Camillus's (2008) proposition at a within-firm level of analysis: stakeholders should "tame" wicked problems by creating a shared understanding of them, and framing them in terms of a collective mission and values. One normative ideal is that individuals and groups should come together at a whole-field level to implement collective solutions , thereby transforming turbulence into order (Iansiti & Levien, 2004; Selsky et al., 2007). Another is that we should remain flexible in our sensemaking, ready to "drop our tools" and view problems from a different level of analysis at any moment (Meyer, Gaba, & Colwell, 2005). Needless to say, transforming these norms into effective prescriptions is a challenge worthy of great minds.

Assumptions and focus of kernel theory

Organized anarchies are characterized by problematic preferences as stakeholders have different goals or, at least, different priorities; by unclear technology, in other words, processes and cognitive systems are too complex for any stakeholder to perceive the whole picture of how they work; and by fluid participation, as different stakeholders participate unpredictably in different issues and decisions (Cohen et al., 1972). Like distributed cognition, a normative theory of organized anarchies assumes that a manager cares about the achievement of some sort of goal, but unlike distributed cognition, there is no assumption that the task or environment will be the same tomorrow as it is today. Thus, the assumed goal may be something more Darwinian, like organizational survival, rather than specific, such as optimizing an investment portfolio or maximizing widget sales. Moreover, the manager will be assumed to be in competition with *other* actors who have their own goals and levers of control.

The foci of organized anarchy theories are the *causal texture* of the environment, which may be of several types, and the *coping strategies* that people and groups employ to deal with them. There may also be a focus on the political process or bargaining that takes place between stakeholders (Eisenhardt & Zbaracki, 1992). These theories are also self-conscious of multiple levels of analysis, connecting phenomena at individual, group, organization, and industry levels.

Community of practice

If decisions and objectives are well-defined, one practices decision development, and if a stable socio-technical system can make a problem tractable, one practices capability crafting. By the process of elimination, then, the idea of organized anarchy is most applicable to those situations where the DSS professional has little say in who the participants of a social-cognitive process might be, which decision criteria will be important to each of them, or how they will make decisions. This may be the challenge for designers of technology, like GDSS, that are to be sold to other firms. It is also highly relevant to CIOs and project managers tasked with building information systems for groups who do not agree on their requirements or goals, as is often the case in strategic decision support at the executive level. El Sherif and El Sawy (1988) described the plight of an IT department tasked with serving a national Cabinet defined by conflicting goals, unstructured decision-making processes, and shifting political winds. They had no mandate to change the political process, so instead they made their job one of “structuring” strategic issues by integrating and cleaning data, defining a single version of the truth, and increasing its accessibility to Cabinet ministers. As this archetype of decision engineering practice is concerned not with resolving disagreements but with supporting and enlightening them, we call it *issue illumination*.

DSS research questions

There are number of questions about issue illumination that we need better answers to. The technological challenges alone of building systems for unknown users to make unstructured decisions based on unknown criteria strain the industry's best engineers, so design theories that can inform them are in serious demand. Beyond that, though, we need to understand organizational issues like how decision support systems affect debate, and how political tension in organizations impacts technology adoption and use. Commensurate with issue illumination's interest in multiple levels of analysis, we are also interested in how IT changes the causal texture of the problem environment, for example, how it makes firms more or less dependent on their ecosystem of partners and competitors. Research that falls under this heading is not yet tightly integrated, but a few exemplary research questions include:

- How does an "issue based" methodology for DSS design and deployment differ from a "decision based" methodology, and what are its benefits? (El Sherif & El Sawy, 1988)
- What are the characteristics of new technologies that affect sensemaking, and what new concepts (e.g. structuration) are needed to understand their effects on organizations? (Weick, 1990)
- What is an effective design (meta-design) for a class of information systems to support emergent knowledge practices (EKPs) in which process, participants, and information requirements are unpredictable? (Markus, Majchrzak, & Gasser, 2002)

- How can firms take advantage of “creation nets” to plug into open innovation at the multi-firm (ecosystem) level, and what changes does that require in the way leaders think about competition? (Hagel & Brown, 2008)

- How are tasks and information systems reciprocally linked? How can we identify the trajectory of their emergent evolution? (Berthon et al., 2007)

Example prescriptions

The most widespread product of the issue illumination practice is the mantra “single version of the truth” that is found in virtually every textbook or brochure about data warehousing and digital dashboards. The core of this idea is not that it ends disagreement, but quite the contrary: it allows us to make our arguments more emphatically with unimpeachable data to back us up. Other prescriptions that come from this practice are those that advise the “innovation champion” in an organization as to how he can garner top management sponsorship for his project, or attract users to his DSS. Integrating the data and structuring the digital scorecards doesn’t constitute a successful BI project unless and until a number of business users with very different goals can agree to use the system to structure their interactions.

Contrasting archetypes

Table 1 differentiates the main archetypes of decision engineering according to their kernel theories, assumptions and focus, the DSS practitioners that they apply to, and important topics for academic research that inform them.

Table 1. Comparison of decision engineering archetypes

	Decision development	Capability crafting	Issue Illumination
Kernel theories	Bounded rationality; heuristics	Distributed cognition; organizational capabilities	Organized anarchies; turbulent environments; wicked problems
Theory assumptions and focus	Decision and decision-maker are specified; alternatives are finite but may not be known	A socio-technical cognitive system with a manager who can perceive it, a clear goal, and some levers of control	Multiple stakeholders with different priorities; intractably complex interrelations; cross-level reverberations
Community of practice	DSS developers	Business process owners, entrepreneurial managers, and CIOs	Harried CIOs and developers, supporting multiple user groups who don't know what they want
DSS research topics	Decision types; IT-enabled heuristics for search and analysis; IT's effects on cognitive limitations	Business process design/management; organizational memory, knowledge management; IT-enabled social-cognitive capabilities	Issue structuring; GDSS for unstructured decisions; digital business ecosystems; social/political effects of DSS

A theory of decision engineering

We conceptualize decision engineering as a complex practice in which multiple normative theories of decision-making are processed and transformed into prescriptive theories-in-use through a rich dialogue involving academics, practitioners, and others: consultants, publishers, students, vendors, professional organizations. This dialogue constitutes the *discipline* of business intelligence. Our view helps bring some order to what might seem like a cacophony of diverse and conflicting technological, organizational, and strategic prescriptions, by teasing it apart into three distinct practices that are each rationally sound.

But how, we may ask, do practitioners access academic theories, and furthermore, how do they apply the right theory for the right problems? Cabantous and Gond (2010) explain that theories of rationality are performative, that they *become* social reality by being performed in practice. Three activities contribute to this process: we engineer rationality, by building information systems like DSS that have rational rules and work processes embedded in them; we conventionalize rationality, by spreading our normative theories to students in business schools and through professional organizations; and we commoditize rationality by identifying what works in one firm and packaging it in books and consulting toolkits to be copied and sold to everyone else. We believe that different technologies, different classes, and different books speak to different parts of the audience. It is the match between the assumptions and focus of a theory, and the constraints and goals of a practitioner, that guides each decision engineer to the theory-based practice that is most relevant to his needs.

An important part of our theory is that the three archetypes of decision engineering are complementary. Organizations do not need to choose one practice, as if they were competing methodologies, but instead researchers need to understand how all three fit together. One way to do this might be to look at the three practices in terms of strategic, tactical, and operational business intelligence, a common distinction made by practitioners. Decision development is the engineering of IT-enhanced strategies for supporting specific decisions at the operational level. Capability crafting can be seen as the engineering of man-machine systems composed of many decisions in order to achieve predetermined goals, a tactical endeavor. Issue illumination can be seen as IT support for the unstructured decision-making and uncertainty management that often characterizes the strategy-making process. The three archetypes may be used by different people at different levels of the same company.

Alternatively, we may show that multiple archetypes are used by the same people. For example, Simons (1990, 1991) discovered that top managers (at the strategic level) had two distinct modes of using management information systems: for most issues, they would build an organization, give it a performance target, and let their subordinates run things; for one or two issues that represented great strategic uncertainty, though, they would make an information system their central sensemaking tool, using its numbers as part of every major meeting and frequently asking for new measures and targets. It seems like they were applying both capability crafting and issue illumination to strategic business intelligence. We can also see at the operational level that capability crafting and decision development are frequently done together: while a decision engineer is thinking about what kinds of data and analyses are necessary to make a particular decision well (decision development), he is also thinking about the processes that will collect the data, how it should be communicated, and how the decision is to be disseminated and carried out in conjunction with other decisions (capability crafting). We may conceive of all three practices as mutually informing.

Directions for future DSS theorizing

We have presented a descriptive theory of decision engineering which shows that business intelligence practice is composed of three parallel and interdependent practices, each based on an identifiable normative theory of decision-making. We believe that this approach provides a better avenue to understanding BI phenomena than other approaches that focus on technology or data, and creates several opportunities for future research. An important next question is, what prescriptions can we produce based on *this* theory? Would it serve our students better, for example, if we divided the teaching curriculum into three units on decisions, cognitive systems, and unstructured issues, respectively? Or can the three archetypes of practice be combined into a systematic methodology for business intelligence development?

Perhaps knowledge of decision engineering's components will help us to build theory around the business intelligence "maturity models" that are ubiquitous in the literature. Perhaps the developmental stages of a business intelligence project naturally progress from decision development to capability crafting and issue structuring as the BI team discovers new needs, or perhaps it is the case that business intelligence efforts evolve over time as academic theories of decision-making fall in and out of favor. Russell et al (2010) present a promising BI maturity model that allows for both a typical progression of implementation skills and the evolution of what business intelligence is thought to mean.

Since we see kernel theories as important inputs to the BI discipline, it is natural to wonder how the practice might be different if theories were different. There are two good ways to explore this question. First, we can go back in time, and see how prescriptions differed in bygone eras, for example back before bounded rationality had supplanted perfect rationality as a model of choice between alternatives. Second, we can look to other cultures, where people and institutions may think about decision-making differently. Would the findings of El Sherif and El Sawy (1988), based on research in Egypt's government, translate well or poorly to similar executive Cabinets in other countries?

Conclusion

Our central thesis has been that the way people think about decision-making, as a result of the dialogue between research and practice, is important to understanding the interventions that they make in the name of business intelligence. We have shown that three different kernel theories from decision science explain three different sets of prescriptions, and three different categories of research questions, that are actually evident in the discipline. Moving forward, we intend to design a research project to examine how these three practices complement and interact with one another in a successful enterprise BI project.

References

Ackoff, R. L. (1994). Systems thinking and thinking systems. *System Dynamics Review, 10*(2-3), 175-188.

Alavi, M., & Leidner, D. E. (2001). Review: Knowledge Management and Knowledge Management Systems: Conceptual Foundations and Research Issues. *MIS Quarterly*, *25*(1), 107-136.

Allison, G., & Zelikow, P. (1999). *Essence of Decision: Explaining the Cuban Missile Crisis* (2nd ed.): Longman.

Alter, S. (2003). Sidestepping the IT Artifact, Scrapping the IS Silo, and Laying Claim to" Systems in Organizations". *Communications of the AIS, 12*(1), 30.

Argyris, C., & Schön, D. A. (1978). *Organizational learning: a theory of action perspective*. Reading, Mass.: Addison-Wesley Pub. Co.

Bell, D. E., Raiffa, H., & Tversky, A. (1988). *Decision making: Descriptive, normative, and prescriptive interactions*. Cambridge: Cambridge University Press.

Berthon, P., Pitt, L. F., Watson, R. T., & Campbell, C. (2007). *"Made" to "Morph": Technology as Transmutation*. Paper presented at the ICIS 2007: JSIS Ancillary Workshop.

Boland, R. J., Jr., Tenkasi, R. V., & Te'eni, D. (1994). Designing Information Technology to Support Distributed Cognition. *Organization Science, 5*(3), 456-475.

Cabantous, L., & Gond, J.-P. (2010). Rational Decision Making as Performative Praxis: Explaining Rationality's *Éternel Retour*. *Organization science*.

Camillus, J. C. (2008). Strategy as a Wicked Problem. *Harvard Business Review, 86*(5), 98-106.

Cohen, M. D., March, J. G., & Olsen, J. P. (1972). A Garbage Can Model of Organizational Choice. *Administrative Science Quarterly*, *17*(1), 1-25.

Davenport, T. H. (2006). Competing on Analytics. *Harvard Business Review, 84*(4), 98-107.

Davenport, T. H. (2009). Make Better Decisions. *Harvard Business Review, 87*(11), 117-123.

Dickson, G. W., DeSanctis, G., & McBride, D. J. (1986). Understanding the effectiveness of computer graphics for decision support: a cumulative experimental approach. *Commun. ACM, 29*(1), 40-47.

Eisenhardt, K. M., & Zbaracki, M. J. (1992). Strategic decision making. *Strategic Management Journal, 13*(S2), 17-37.

El Sherif, H., & El Sawy, O. A. (1988). Issue-Based Decision Support Systems for the Egyptian Cabinet. *MIS Quarterly, 12*(4), 551-569.

Emery, F. E., & Trist, E. L. (1965). The Causal Texture of Organizational Environments. *Human Relations, 18*(1), 21-32.

Ethiraj, S. K., Kale, P., Krishnan, M. S., & Singh, J. V. (2005). Where do capabilities come from and how do they matter? A study in the software services industry. *Strategic Management Journal, 26*(1), 25-45.

Ferrari, A., & Rossignoli, C. (2010). Do Business Intelligence Systems Enforce Organizational Coordination Mechanisms? In A. D'Atri & D. Saccà (Eds.), *Information Systems: People, Organizations, Institutions, and Technologies* (pp. 433-440): Physica-Verlag HD.

Freeze, R. D., & Kulkarni, U. (2007). Knowledge management capability: defining knowledge assets. *Journal of Knowledge Management, 11*(6), 94-109.

Goodhue, D. L., & Thompson, R. L. (1995). Task-Technology Fit and Individual Performance. *MIS Quarterly, 19*(2), 213-236.

Gregor, S. (2006). The Nature of Theory in Information Systems. *MIS Quarterly, 30*(3), 611-642.

Gregor, S., & Jones, D. (2007). The Anatomy of a Design Theory. *Journal of the AIS, 8*(5), 312-335.

Gregor, S. (2009). *Building Theory in the Sciences of the Artificial*. Paper presented at the Proceedings of the 4th International Conference on Design Science Research in Information Systems and Technology.

Hagel, J., & Brown, J. S. (2008). Creation Nets: Harnessing the Potential of Open Innovation. *Journal of Service Science, 1*(2), 27-40.

Herschel, R., & Yermish, I. (2009). Knowledge Management in Business Intelligence. In W. R. King (Ed.), *Knowledge Management and Organizational Learning* (Vol. 4, pp. 131-143): Springer US.

Hevner, A. R., March, S. T., Park, J., & Ram, S. (2004). Design Science in Information Systems Research. *MIS Quarterly, 28*(1), 75-106.

Hollan, J., Hutchins, E., & Kirsh, D. (2000). Distributed cognition: toward a new foundation for human-computer interaction research. *ACM Transactions on Computer-Human Interaction, 7*(2), 174-196.

Hutchins, E. (1995a). *Cognition in the wild*. Cambridge, Mass.: MIT Press.

Hutchins, E. (1995b). How a cockpit remembers its speeds. *Cognitive Science, 19*(3), 265-288.

Iansiti, M., & Levien, R. (2004). Strategy as Ecology. *Harvard Business Review, 82*(3), 68-78.

Kaplan, R. S., & Norton, D. P. (1996). Linking the Balanced Scorecard to Strategy. *California Management Review, 39*(1), 53-79.

Imhoff, C., & White, C. (2008). Full Circle: Decision Intelligence (DSS 2.0). *BeyeNetwork*. Retrieved from http://www.b-eye-network.com/channels/1138/view/8385

Lovallo, D., & Sibony, O. (2010). The case for behavioral strategy. *McKinsey Quarterly*.

Malhotra, A., Gosain, S., & El Sawy, O. A. (2005). Absorptive Capacity Configurations in Supply Chains: Gearing for Partner-Enabled Market Knowledge Creation. *MIS Quarterly, 29*(1), 145-187.

Markus, M. L., Majchrzak, A., & Gasser, L. (2002). A Design Theory for Systems That Support Emergent Knowledge Processes. *MIS Quarterly, 26*(3), 179-212.

Meyer, A. D., Gaba, V., & Colwell, K. A. (2005). Organizing Far from Equilibrium: Nonlinear Change in Organizational Fields. *Organization Science, 16*(5), 456-473.

Michel, A. A. (2007). A Distributed Cognition Perspective on Newcomers' Change Processes: The Management of Cognitive Uncertainty in Two Investment Banks. *Administrative Science Quarterly, 52*(4), 507-557.

Overby, E., Bharadwaj, A., & Sambamurthy, V. (2006). Enterprise agility and the enabling role of information technology. *European Journal of Information Systems, 15*(2), 120-131.

Park, Y.-T. (2006). An empirical investigation of the effects of data warehousing on decision performance. *Information & Management, 43*(1), 51-61.

Pavlou, P. A., & Sawy, O. A. E. (2010). The "Third Hand": IT-Enabled Competitive Advantage in Turbulence Through Improvisational Capabilities. *Information Systems Research*(forthcoming).

Rittel, H. W. J., & Webber, M. M. (1973). Dilemmas in a general theory of planning. *Policy Sciences, 4*(2), 155-169.

Russell, S., Haddad, M., Bruni, M., & Granger, M. (2010). *Organic Evolution and the Capability Maturity of Business Intelligence*. Paper presented at the Americas Conference on Information Systems (AMCIS).

Sarker, S., & Lee, A. S. (2002). Using a Positivist Case Research Methodology to Test Three Competing Theories-in-Use of Business Process Redesign. *Journal of the AIS, 2*(1, Article 7).

Selsky, J. W., Goes, J., & Baburoglu, O. N. (2007). Contrasting Perspectives of Strategy Making: Applications in 'Hyper' Environments. *Organization Studies, 28*(1), 71-94.

Simon, H. A. (1979). Rational Decision Making in Business Organizations. *The American Economic Review, 69*(4), 493-513.

Simon, H. A. (1996). *The Sciences of the Artificial* (3rd ed.). Cambridge, Massachusetts: The MIT Press.

Simons, R. (1990). The role of management control systems in creating competitive advantage: new perspectives. *Accounting, Organizations and Society, 15*(1/2), 127-143.

Simons, R. (1991). Strategic Orientation and Top Management Attention to Control Systems. *Strategic Management Journal, 12*(1), 49-62.

Stein, E., & Zwass, V. (1995). Actualizing organizational memory with information systems. *Information Systems Research, 6*(2), 85.

Todd, P. A., & Benbasat, I. (1994). The influence of decision aids on choice strategies under conditions of high cognitive load. *Systems, Man and Cybernetics, IEEE Transactions on, 24*(4), 537-547.

Todd, P. M., & Gigerenzer, G. (2000). Précis of Simple heuristics that make us smart. *Behavioral and Brain Sciences, 23*(05), 727-741.

Tversky, A., & Kahneman, D. (1974). Judgment under Uncertainty: Heuristics and Biases. *Science, 185*(4157), 1124-1131.

Upton, C., Doherty, G., Gleeson, F., & Sheridan, C. (2008). Designing decision support in an evolving sociotechnical enterprise. *Cognition, Technology & Work, 12*(1), 13-30.

Walls, J. G., Widmeyer, G. R., & El Sawy, O. A. (1992). Building an Information System Design Theory for Vigilant EIS. *Information Systems Research, 3*(1), 36-59.

Walls, J. G., Widmeyer, G. R., & El Sawy, O. A. (2004). Assessing information system design theory in perspective: How useful was our 1992 initial rendition? *Journal of Information Technology Theory and Application, 6*(2), 43-58.

Watson, H. J. (2009). Tutorial: Business Intelligence - Past, Present, and Future. *Communications of the AIS, 25*, 487-510.

Wegner, D. M. (1986). Transactive memory: A contemporary analysis of the group mind. In B. Mullen & G. R. Goethals (Eds.), *Theories of Group Behavior* (pp. 185-208). New York: Springer-Verlag.

Weick, K. E. (1990). Technology as equivoque: Sensemaking in new technologies. In P. S. Goodman & L. S. Sproull (Eds.), *Technology and organizations. The Jossey-Bass management series.* (pp. 1-44). San Francisco, CA: Jossey-Bass.

Weick, K. E., & Roberts, K. H. (1993). Collective Mind in Organizations: Heedful Interrelating on Flight Decks. *Administrative Science Quarterly, 38*(3), 357-381.

Williams, M. L., Dennis, A. R., Stam, A., & Aronson, J. E. (2007). The impact of DSS use and information load on errors and decision quality. *European Journal of Operational Research, 176*(1), 468-481.