An Efficient Frontier in Organization Design: Organizational Structure As a Determinant of Exploration and Exploitation[1]

Felipe A. Csaszar

Stephen M. Ross School of Business

University of Michigan

701 Tappan Ave. R4336

Ann Arbor, MI 48109

Tel: (734) 615 4854 – Fax: (734) 936 6631

Email: fcsaszar@umich.edu

July 10, 2012

Forthcoming in Organization Science

[1]The author owes a special debt to Jitendra Singh and Raaj Sah for their early feedback and valuable insight. He thanks Dan Levinthal, Nicolaj Siggelkow, Sid Winter, and those in attendance at the Management Ph.D. Brown Bag at Wharton for offering many useful comments and suggestions. He also thanks Thorbjørn Knudsen and Ian Larkin for keen reviews and is grateful to Rich Burton and three anonymous referees for their helpful comments. Research funding from the INSEAD Alumni Fund and from the Ackoff Fund of the Wharton Risk Management and Decision Processes Center is gratefully acknowledged. Errors and omissions remain the author's.

Abstract

This paper develops a parsimonious process-level theory that connects organizational structure to exploration and exploitation. Toward this end, it develops a mathematical model of organizational decision making that combines an information processing approach in the spirit of Sah and Stiglitz (1986) with elements from signal detection theory (Green and Swets, 1966). The model is first used to explore a "design space" of organizations and identify trade-offs and dominance relationships among alternative organization designs. The paper then studies open questions in the organization design literature, such as the extent to which exploration and exploitation can be produced by one organization and what is the effect of organization size on exploration. More broadly, this research speaks to calls for introducing more process-level explanations in the organizations literature. The paper concludes with testable hypotheses and managerially relevant insights.

Keywords: organization design, organizational decision making, omission and commission errors, exploration and exploitation, ambidexterity

1 Introduction

There is anecdotal evidence suggesting that organizational structure may be connected to exploration and exploitation in ways not studied by research. One illustration is given by Lou Gerstner, CEO of IBM during its massive restructuring in the late 1990s. He mentions that one of his main contributions toward increasing IBM's innovativeness was to scrap the firm's "non-concur policy," which allowed any department to veto a project of any other department (Gerstner, 2003:189–199). He explains that this policy was sensible in the pre-PC era, when IBM competed with highly integrated, "monolithic" offerings. This policy ensured that IBM did not introduce products that would have adverse interactions with the rest of the firm's offerings. But gradually, as modular and open technologies started to emerge (eventually leading to the deconstruction of the IT industry value chain in the 1980s and '90s), the non-concur policy ceased to be useful. Gerstner's story suggests that, even though the non-concur policy had once been a sensible way of exploiting IBM's position, in the new scenario it was inhibiting IBM's exploration of new products.

Gerstner's story nicely illustrates the idea that there is a process (in his story, the non-concur policy) connecting organizational structure to both exploration and exploitation (March, 1991). The story is interesting because the organization design literature has not studied the possibility that exploration and exploitation (E&E) may be different outcomes of a single underlying process.

At least since the work of Burns and Stalker (1961) on organic and mechanistic organizations, studies have shown that some structures have a better fit with exploration, while others have a better fit with exploitation. Although these ideas were embraced in most of the subsequent work on organization design (Pettigrew and Fenton, 2000:6), there is little research describing the underlying process that connects structure to E&E. Most of the current explanations are "fit" type of arguments whereby a bundle of practices is shown to be self-reinforcing and congruent with achieving a given goal.

Although the field has achieved much clarity and detail about the contingencies that exhibit the best fit with each of the organizational types, the field remains far from having simple, process-level explanations that connect structure to E&E. In other words, the current explanations are not at the level of processes, but at higher levels of aggregation. Hence this paper deals first with providing an alternative and parsimonious logic for explaining the relationship between the structure of an organization and its ability to explore and exploit.

Another issue addressed by this paper is the extent to which a single organization can achieve exploration and exploitation simultaneously (or whether, as suggested by Gerstner's story, these goals are mutually exclusive). Because the challenge of achieving E&E simultaneously is akin to the challenge of pursuing ambidexterity (defined as achieving efficiency in the short term while remaining innovative in the long term; see Tushman and O'Reilly, 1996), the ambidexterity literature has pursued a similar question. The literature on *how* ambidexterity is achieved is scant (Raisch and Birkinshaw, 2008:380), but the few papers that have explored this suggest that ambidexterity is achieved by implementing two different processes (one pursuing exploration, the other exploitation) separated in either time or space (Raisch and Birkinshaw, 2008:389). Thus, the extent to which E&E can be achieved by a single organization at the same time and place is an open question.

In order to study in a rigorous manner the relationship between structure and E&E, this paper develops a mathematical model that builds on ideas first used by Sah and Stiglitz (1986) that conceptualize organizations as information filters made of fallible decision makers. This modeling methodology, with roots in reliability theory (Moore and Shannon, 1956/1993; von Neumann, 1956) and team theory (Marschak and Radner, 1972) and more recently used and extended in the management literature (Knudsen and Levinthal, 2007; Christensen and Knudsen, 2010; Csaszar, 2012), provides two main advantages over other approaches: (i) it is a process explanation, that is, it clearly postulates an underlying mechanism connecting structure to E&E; and (ii) it makes quantitative predictions—that is, it allows for empirical testing and falsification. By studying a broad range of organization designs in a quantitative way, this paper uncovers a "design space" of organizations and discusses the trade-offs and dominance relationships among alternative organization designs.

The aim of this paper is to contribute to a more detailed, process-level, and quantitative understanding of the connections among organizational structure, exploration, and exploitation. This is an important research subject because all organizations must choose their organizational structure and because E&E are important outcomes for most organizations.

In terms of contributions, this paper provides (a) an alternative and parsimonious logic that explains the relationship between structure and the achievement of E&E; (b) an answer to the question of to what extent it is possible to assemble "hybrid" or ambidextrous organizations, able

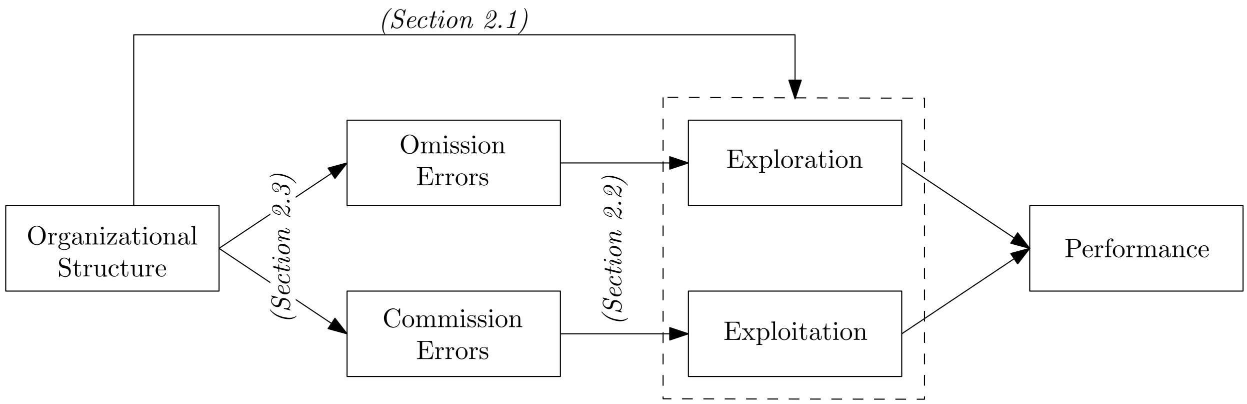

Figure 1: Causal relationships among the constructs explored in Section 2. Each arrow designates a causal relationship (as in, e.g., "Organizational Structure affects Omission Errors").

to pursue E&E simultaneously; (c) a methodology that allows to quantify the expected E&E of a broad range of structures; and (d) a set of "efficient" structures—that is, structures that are particularly good at achieving a given combination of E&E—leading to the concept of an "efficient frontier" in organization design.

2 Theoretical Motivation

This section develops theoretical connections between organizational structure and E&E. To do so, it proceeds as illustrated in Figure 1. First, it shows that little is known regarding how organizational structure affects E&E (Section 2.1). Second, it describes how E&E can be understood in terms of omission and commission errors (Section 2.2). Third, it shows that organizational structure is a determinant of omission and commission errors (Section 2.3) and hence that there is a causal connection between organizational structure and E&E. Additionally, this section highlights the research antecedents of the current work and describes the differences between the current and previous approaches (Section 2.4).

2.1 Achieving Exploration and Exploitation: Open Questions in Organization Design

Simon (1947/1997:18–19) defined an organization as "the pattern of communications and relations among a group of human beings, including the processes for making and implementing decisions." Accordingly, organization design has been defined in terms of choosing this pattern—which is sometimes called the firms' configuration, structure, or architecture (Burton and Obel, 2004)—and aims to produce coherence among the goals for which the organization exists, the patterns of division of labor and inter-unit coordination, and the people who will do the work (Galbraith, 1977). Organization design is concerned with the design of systems that ensure effective communication and integration of effort (Child, 1984), and it seeks to discover (i) who makes what decisions based on what information and (ii) what is the structure of communication (March and Simon, 1958/1993; Galbraith, 1973). Even though organization design is a critical decision of the executive (Burton and Obel, 2004), research on it is in many respects an emerging field (Daft and Lewin, 1993; Zenger and Hesterly, 1997; Foss, 2003).

An important question within the broad field of organization design is how to design organizations that achieve particular levels of E&E. March (1991:71) defined E&E as latent constructs underlying a broad set of organizational behaviors, stating that "exploration includes things captured by terms such as search, variation, risk taking, experimentation, play, flexibility, discovery, [and] innovation" whereas "exploitation includes such things such as refinement, choice, production, efficiency, selection, implementation, [and] execution." Exploration and exploitation are key dimensions of performance for many organizations. For example, R&D labs and high-tech firms are typically concerned with exploration; in contrast, organizations facing stable environments or harsh conditions are, in general, mostly concerned with exploitation. Moreover, organizations may have dual goals; for example, a pharmaceutical company may want to discover novel drugs (i.e., to explore) while at the same time avoiding FDA rejections and keeping operations efficient (i.e., to exploit).

The research on how to organize for E&E falls into at least three different literatures: the classic literature on organization design (e.g., Burns and Stalker, 1961; Lawrence and Lorsch, 1967); the literature on exploration, exploitation, and ambidexterity (e.g., March, 1991; Tushman and O'Reilly, 1996; Siggelkow and Levinthal, 2003); and the literature on behavioral approaches to reliability (e.g., Perrow, 1984; Roberts, 1990; Weick and Roberts, 1993). The following paragraphs summarize the main findings on how to organize for E&E.

Perhaps the first researchers to discuss what organization designs were most appropriate for pursuing either exploration or exploitation were Burns and Stalker (1961), who argued that organic structures foster exploration because they feature fluid job descriptions, high communication, and few rules—characteristics that free employees, giving them leeway to explore new ideas. Mechanistic structures, which use virtually opposite practices (e.g., clear line of command, centralization, and plenty of rules), support exploitation at the expense of being less explorative.

Other researchers have elaborated Burns and Stalker's (1961) ideas. For example, Lawrence and Lorsch (1967) discussed the challenges of differentiation and integration in firms that exhibit both organic and mechanistic divisions, and Miles and Snow (1978) created a two-dimensional typology of organizational forms in which organic and mechanistic traits are assigned separate and perpendicular axes, leading to four basic forms (Defenders, Analyzers, Reactors, and Prospectors). More recently, Burton, DeSanctis, and Obel (2006) proposed a framework that combines ideas from the classic literatures on information processing (Cyert and March, 1963; Galbraith, 1977) with Miles and Snow's (1978) typology and provides a unified platform from which to study the interrelationship between strategy and organization design. In general, all the theories that descended from Burns and Stalker (1961) have based their explanations on the concept of "fit"; that is, they have viewed firms as systems of highly interdependent elements whose overall performance is correlated with achieving consistency among those elements (Siggelkow, 2001).

The literature on E&E that has descended from March (1991) has primarily addressed the consequences of E&E, not their antecedents. Hence, structural concerns are less present in this literature. As noted by Siggelkow and Levinthal (2003:650), "little is known about how different organizational structures moderate this balance [between E&E]." Similarly, in the literature on ambidexterity, Raisch and Birkinshaw (2008:380) reported that "far less research has traditionally been devoted to how organizations achieve organizational ambidexterity." These remarks confirm the notion that structural concerns remain mostly unexplored in these literatures. In fact, even the effect of one of the coarsest measures of structure, organization size, remains unclear; for instance, Beckman et al. (2004) showed that exploration increases with firm size whereas Rothaermel and Deeds (2004) showed the opposite.

The few papers dealing with the organizational aspects of ambidexterity have proposed some complex integration mechanisms (Adler and Borys, 1996; Sheremata, 2000), but they have generally agreed that ambidexterity is more likely to be achieved by running two processes (one exploring and the other exploiting) that are separated either temporally (Brown and Eisenhardt, 1997; Nickerson and Zenger, 2002) or spatially (Tushman and O'Reilly, 1996; Christensen, 1997; Benner and

Tushman, 2003).

The other literature that can speak to the structural determinants of E&E is the literature on behavioral causes of reliability (Perrow, 1984; Roberts, 1990; Weick and Roberts, 1993), which has proposed several mechanisms to achieve reliability (i.e., a type of exploitation); these include setting reliability as an organizational goal, creating a culture of reliability, continually training, and mindfulness. However, few of these mechanisms relate to organizational structure. It is interesting that, the two main theories of organizational reliability—normal accidents theory (Perrow, 1984) and high-reliability organization theory (Roberts, 1990; Weick and Roberts, 1993)—do not agree on the effect of a basic organization design concept: redundancy. For the former, redundancy often causes accidents because it increases complexity and opaqueness and encourages risk taking; for the latter, redundancy enhances safety because duplication and overlap serve to detect errors and act as a backup system.

In sum, the different literatures that have studied how to organize for E&E propose few (and sometimes opposing) explanatory mechanisms that relate to organizational structure. Moreover, the arguments advanced in these literatures share at least two characteristics: (i) they are mostly arguments about "fit" or about high-level constructs, but not about underlying processes; and (ii) they do not suggest the existence of one single mechanism that can determine both exploration and exploitation at the same time. Although in a limited and simplified manner, this paper provides a logic that fulfills these two conditions, providing an alternative and parsimonious logic to explain the relationship between structure and E&E.

2.2 Omission and Commission Errors as Determinants of Exploration and Exploitation

To formally analyze the organizational determinants of E&E, this paper borrows from decision theory (Berger, 1985) the concept of omission and commission (O&C) errors.

Decision theory is concerned with devising optimal decision rules given imperfect knowledge about the state of nature, about the actions available to the decision maker, and about the consequences of those actions. Unlike economics, which is stated in terms of gains (utility), decision theory is stated in terms of losses—that is, deviations with respect to a theoretically optimal decision. These deviations are called errors, and can be of two fundamentally different types: omission errors (or Type I) and commission errors (or Type II). Most decision-making tasks produce these two types errors. For example, a venture capital firm may miss a great investment opportunity (an omission error) or may invest in a bad firm (a commission error). Similar examples can be devised for many managerially relevant contexts (e.g., decisions that concern hiring, mergers and acquisitions, expansion, research, and product development). More generally, given that decision making is a central task of organizations (Simon, 1947/1997), it seems natural to use decision theory as a lens through which to look at firm decisions.

Although economics and decision theory are stated in different terms (i.e., utilities versus losses), there is an equivalence between the two formulations in that losses can be viewed as negative utilities. Thus, a decision-theoretic view of a firm's performance may be written as

$$\text{Performance} = \beta_0 - \beta_1 \text{OmissionErrors} - \beta_2 \text{CommissionErrors}.$$

This equation assumes that the firm has a theoretically optimal performance (β_0), which decreases the more omission or commission errors the firm makes. Note that O&C errors carry different coefficients, since in most cases the errors entail different costs. For example, a highly leveraged retailer may succumb because of an expansion into a wrong market, yet might be unaffected by missing a market expansion opportunity (i.e., $\beta_2 > \beta_1$); conversely, a firm competing at the leading edge of technology may perish if it misses the next technological wave but could be relatively unaffected by trying one or more technological dead ends (i.e., $\beta_1 > \beta_2$).

Although rarely mentioned in organizational contexts, O&C errors are tightly connected to E&E. The following quotes, which represent some of the notable exceptions, serve to make concrete the connection between O&C errors and E&E. Garud et al. (1997:33) exemplify this connection as follows

> [D]uring the era of ferment ... decision makers need to encourage the exploration of all alternatives. In other words [omission[1]] errors need to be reduced so that consideration of potentially promising technologies is not prematurely foreclosed. ... During the

[1] The original text is stated in terms of Type I and II errors, with the labeling reversed with respect to Sah and Stiglitz's (1986) notation. See footnote 10 in Langlois (1997:91) for a discussion of the pros and cons of these alternative labeling schemes. To avoid confusion, in this paper the terminology of omission and commission errors is used.

> era of incremental change, decision makers need to encourage many incremental innovations. In other words, [commission] errors are to be reduced so that corporate resources are not misdirected towards trying to develop alternatives to the already dominant technology.

Similarly, Garicano and Posner (2005:157) establish the following link between O&C errors and E&E in the context of intelligence failures at the FBI.

> When the environment is unstable, the organization should be decentralized to maximize the likelihood that many fresh new ideas will be produced .... When the environment is stable, in the sense of changing slowly, so that adaptation is achievable by incremental adjustments, the organization should be designed with many filters so that the errors that are made are of the type "a few good ideas were not tried out" [i.e., minimizing omission errors] rather than "we went ahead with some terrible ideas and damaged our franchise" [i.e., minimizing commission errors]. 3M is an example of a business firm that has successfully implemented a "more ideas are better" strategy. ... In contrast, at Procter and Gamble, which operates in a more stable product environment, new product proposals go through 40 to 50 revisions until they reach the [CEO] ....

These quotes describe a definite relationship between O&C errors and E&E. However, it remains unclear to what other settings this relationship applies—in other words, under what conditions it is true. To explore the prevalence of this relationship, I first state it in unambiguous terms and then lay out a minimal set of sufficient conditions. Thus, the relationship between O&C errors and E&E may be stated as follows: *exploring firms pursue more good projects*, whereas *exploiting firms pursue fewer bad projects*. The former statement is true almost by definition: as exploring firms search more, they are more likely to discover good projects (March, 1991; Levinthal, 1997). The latter statement can be justified by showing that bad projects are costlier for exploiting firms than for exploring ones, which means that exploiting firms try harder to avoid them.

One possible explanation for exploiting firms' higher cost of bad projects involves modularity. Firms that explore are more likely to be organized in a modular way so that new projects can be started and terminated easily (Ethiraj and Levinthal, 2004). In contrast, firms that exploit will be more likely to have a monolithic, highly customized operation designed to fit their current projects (since these firms don't need to start new projects as frequently). It is therefore likely that accepting a project will require that the exploiting firm implement costly adaptations to its complex operations. If the project later fails, then the cost of adapting is wasted and, moreover, extra resources must be spent to undo the adaptations. In sum, because bad projects can be quite costly to an exploiting firm, it will be extra cautious when analyzing projects and hence will accept fewer bad ones.

Another factor that might explain the higher cost of bad projects to exploiting firms is financial slack. Pursuing exploration is riskier than pursuing exploitation, so many firms that pursue exploration create a buffer against bad luck by maintaining a financial cushion (O'Brien, 2003). Thus, a well-funded exploring firm can survive many bad projects, whereas an exploiting firm—which will typically operate on thin margins—is less likely to survive. In this view, then, bad projects have a higher cost for an exploiting firm because they are more likely to disrupt (or even bankrupt) exploiting than exploring firms.

The preceding logic implies, under some relatively common conditions, that decreasing omission errors will lead to increasing exploration and that decreasing commission errors will lead to increasing exploitation. That logic seems reasonable because it is parsimonious, fits the quoted stories, and is consistent with previous theories (e.g., the effects of modularity and financial slack).[2]

Using O&C errors to understand E&E is promising for at least two reasons. First, because of their decision-theoretic roots, O&C errors offer an unambiguous way to characterize and measure the sometimes elusive concepts of exploration, exploitation, and ambidexterity. Second, as shown next, O&C errors are affected by organizational structure and thus serve as a bridge to the underexplored topic of E&E's structural determinants.

[2]This logic does not imply that O&C errors are the only determinants of E&E or that O&C errors could not affect E&E differently under other assumptions. As pointed out by a referee, one may conceive of more complex models in which O&C errors could affect E&E in ways other than as illustrated by the quotations. For instance, with a model in which projects can be only of two types (exploration or exploitation), missing an exploitation project would lead to less exploitation overall. Also, a model that incorporates regret might predict that, after missing a good exploration project, the firm may actually increase the number of exploration projects it pursues. Hence the model presented in this paper should be interpreted not as an exhaustive explanation of E&E but instead as one parsimonious and plausible mechanism by which structure, mediated by O&C errors, affects E&E.

2.3 Organizational Structure as a Determinant of Omission and Commission Errors

The work of Sah and Stiglitz (1986, 1988) sheds light on the remaining set of arrows in Figure 1: how organizational structure affects O&C errors. The following paragraph summarizes the rationale behind Sah and Stiglitz's models.

Sah and Stiglitz (1986) analyzed the performance of two stylized organizational structures. Their basic model is in the tradition of team theory (Marschak and Radner, 1972; Radner, 1993; Van Zandt, 1999) and thus abstracts from incentive problems to focus solely on information processing considerations. The organizations in their model are the simplest possible (each composed of two individuals) and are in charge of "screening" projects—that is, deciding whether a project is good (and deserves to be approved) or not. Here is an example that reflects the intuition of their model: If two individuals (e.g., partners in a venture capital firm) must agree on the quality of a project before investing in it, then the probability that the project will be funded is lower than if either individual can approve it independently. Sah and Stiglitz (1986) referred to these two ways of organizing as a hierarchy and a polyarchy, respectively.[3] Because the polyarchy (the two independent individuals) approves a higher proportion of projects, it has a smaller chance of missing a good project than the hierarchy (the two dependent individuals) albeit at the expense of a greater chance of investing in a bad project. Sah and Stiglitz mathematically formalized this intuition and showed that the hierarchy minimizes commission errors whereas the polyarchy minimizes omission errors. Their work implies that the two different structures allow an organization to trade off one error against the other; hence, which structure is "better" depends on the context (i.e., on the relative cost of the two errors).

In a continuation of their 1986 work, Sah and Stiglitz (1988) extended their model to account for the O&C errors of committees. In a committee there are N decision makers, and the approval by C of them is required for a project to be approved (e.g., a seven-member committee using majority rule would be characterized by $N = 7$ and $C = 4$).

Although they abstract from reality, the models of Sah and Stiglitz (1986, 1988) capture some essential characteristics of organizing: that one important characteristic of organization is how

[3] Sah and Stiglitz's use of the word *hierarchy* is nonstandard, since neither of the two evaluators is a superior of the other. It should not be confused with the traditional meaning of the word in the literature on organizations.

information is processed (Cyert and March, 1963) and that O&C errors are important outcomes for many organizations. Because these models are agnostic with respect to what a "project" is, they can be used to illustrate a broad range of organizational decision-making phenomena. The only constraint on the nature of the projects is that it must be possible to make a binary decision regarding them (e.g., accept/reject, fund/do not fund, hire/do not hire, etc.). The work of Sah and Stiglitz is valuable for the purposes of this paper because it provides a transparent mechanism connecting organizational structure to O&C errors.

2.4 Differences Between the Current and Previous Approaches

Few papers in the management literature have used or extended Sah and Stiglitz's (1986; 1988) arguments. Among those that have used their arguments is work addressing mergers and acquisitions (Puranam et al., 2006), venture capital syndication (Lerner, 1994), and technological choices (Garud, Nayyar, and Shapira, 1997). The papers that are closest to the current one—because they use or extend the ideas of Sah and Stiglitz (1986, 1988) in a central way to their arguments—are Christensen and Knudsen (2002), Knudsen and Levinthal (2007), and Christensen and Knudsen (2010). The next few paragraphs briefly describe these papers and point out how they differ from the current work.

Christensen and Knudsen (2002) described one approach to extending the work of Sah and Stiglitz (1986, 1988) to deal with a broader set of structures. In doing so, they clarified several ambiguities present in previous project selection models and outlined a graph-traversing algorithm that allows one to compute O&C errors for this broader set of structures. Their paper focused on the representational and computational aspects of extending the ideas of Sah and Stiglitz and thus are an antecedent of the methodology (Section 3) of the current paper.

The method described in the current paper uses a representation similar to that employed by Christensen and Knudsen (2002), but provides a different approach to computing the O&C errors. The new methodology offers two advantages. First, it describes a key element of the previous models—the individuals' screening functions—not as a primitive of the model but rather as the outcome of noisy perception. This increases the model's behavioral plausibility and allows one to study the effect of varying the accuracy of individuals' perceptions in a manner concordant with the psychological literature in this respect.[4] Second, the new method yields closed-form expressions for computing O&C errors, which simplifies the task of computing them (this is especially useful given the large number of structures studied here).

Christensen and Knudsen (2010) gave a mathematical proof showing that it is possible to construct an arbitrarily reliable organization using unreliable individuals. The main goal of their paper was to provide an organizational equivalent to a famous result of reliability theory, the Moore–Shannon theorem (Moore and Shannon, 1956/1993), which states that it is possible to create a reliable system out of unreliable parts.

Knudsen and Levinthal (2007) explored the interactions between bounded rationality (in the form of local search) and imperfect evaluation of alternatives due to how information is aggregated by different organizational structures. They accomplish this goal by comparing how six different organizational structures perceive and search in an NK landscape (Levinthal, 1997). They extended the NK literature by showing how imperfect evaluation leads to a more robust search process (as it diminishes the relevance of the initial position of the firm) and show how some structures are better at exploring the landscape. The current paper shares with Knudsen and Levinthal (2007) an interest in innovation and in project selection models, but it pursues different questions using different methodologies.

3 The Model

The model in this paper allows for computing the probability that a given organization will make an omission error or a commission error. The inputs of the model are the noisiness or fallibility of the individual decision makers that belong to the organization, the organizational structure (conceptualized as a decision-making structure), and the probability distribution of projects received by the organization.

The model studies a particular class of organizations. Intuitively, any of these organizations

[4] Signal detection theory (Green and Swets, 1966; Macmillan and Creelman, 2004) has developed a rich set of models and experiments to measure the accuracy of individuals when confronted with decision tasks. The prototypical experimental setting of this theory is producing a yes/no answer depending on the quality of a stimulus. The model in the current paper attempts to create a seamless connection between the micro level explored by signal detection theory and the more macro level explored by project selection models. Connecting the models is straightforward; it consists in expressing the screening function in terms of the noisiness of the decision maker (see Section 3.1). Combining both theories is parsimonious and provides a sound foundation for project selection models.

can be viewed as a set of individuals, arranged in a particular way, whose task is to decide whether a given project should be accepted or rejected. The project enters into the organization through a preestablished agent and is passed from one agent to another depending each agent's evaluation (e.g., if agent 1 accepts a project then it is sent to agent 2; otherwise it is sent to agent 3). Agents' evaluation of the project depend on its intrinsic quality and on random noise due to the agent's fallibility as a decision maker. Some of the agents (those at the ends of the decision process) have the power to decide in the name of the organization—that is, to accept or reject the project outright.

The exposition of the model can be logically divided into three parts: (i) a characterization of the individuals that participate in the organization, particularly of their ability to discern good from bad projects; (ii) a notation to describe decision-making structures; and (iii) a method for calculating the organization-level performance, which is measured as the probability of making an omission or a commission error. The rest of this section describes in more detail the broad picture just outlined. For the sake of replicability, a mathematical derivation of all the model's elements is presented in Appendix A.

3.1 Describing One Individual

The individuals in the model are responsible for accepting or rejecting a project depending on how they perceive its quality. A literature that has extensively studied this exact phenomenon is signal detection theory in psychology (Green and Swets, 1966; Macmillan and Creelman, 2004). The current paper borrows from this literature the description of perception as a signal-plus-noise process.

Individuals' perception is modeled as follows. Suppose an individual has to evaluate a project whose real quality is q. Given her fallibility as a decision maker, we assume that she will perceive a noisy signal $q' = q + \tilde{n}$. Most of the examples in this paper assume that $\tilde{n} \sim N(0, 1)$, but the methodology can accept any arbitrary probability distribution.[5] The task of the individual is to recommend approving those projects that she perceives have a positive quality (i.e., $q' > 0$) and to

[5] The normal distribution is a convenient starting point but may not represent adequately real-life evaluators, who could be heterogeneous (e.g., some individuals may be more discriminating than others) or not independent (as in the case of herding). They may also have biased or skewed perceptions, as is known from prospect theory (Kahneman and Tversky, 1979). These considerations are discussed further in Section 4.1.

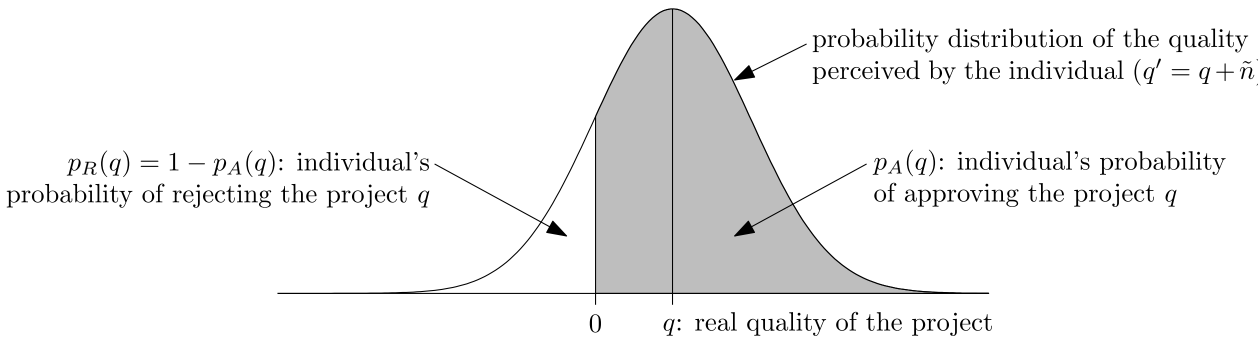

Figure 2: Graphical representation of the probability that an individual approves a project of quality q.

recommend rejecting otherwise. This formulation has the intuitive property that individuals are more likely to make errors when the noise is greater and when the project quality is closer to zero.

The shaded area in Figure 2 corresponds to the probability that the quality perceived by the individual will be positive or, in other words, that she will recommend approving a project of quality q. This probability, denoted $p_A(q)$, is called "the screening function of the individual." (See Appendix A for all the equations.) The probability of an individual recommending the rejection of a project of quality q is $p_R(q) = 1 - p_A(q)$.

3.2 Representing Decision-Making Structures

In line with previous models (Knudsen and Levinthal, 2007; Christensen and Knudsen, 2010), this paper uses a graph notation to represent the decision-making structure of organizations. Under the specific notation used in the current paper, the agents are arranged on a grid, and the structure of the organization is described by arrows connecting the agents. The representation used here works in the following way: (a) the agent in the top left is the first to receive the project; (b) once an agent receives a project, he takes one time period to evaluate it; (c) once an agent evaluates a project, he can *either* accept it and send it to the right *or* reject it and send it down; (d) if the project exits through the right side of the organization then it is approved, but if exits through the bottom side then it is rejected.

Figure 3 shows how this notation can be used to express two simple structures that correspond to Sah and Stiglitz's hierarchy (which approves a project only if both agents accept it) and polyarchy (which approves a project if either one of the two agents accepts it). These structures mirror the

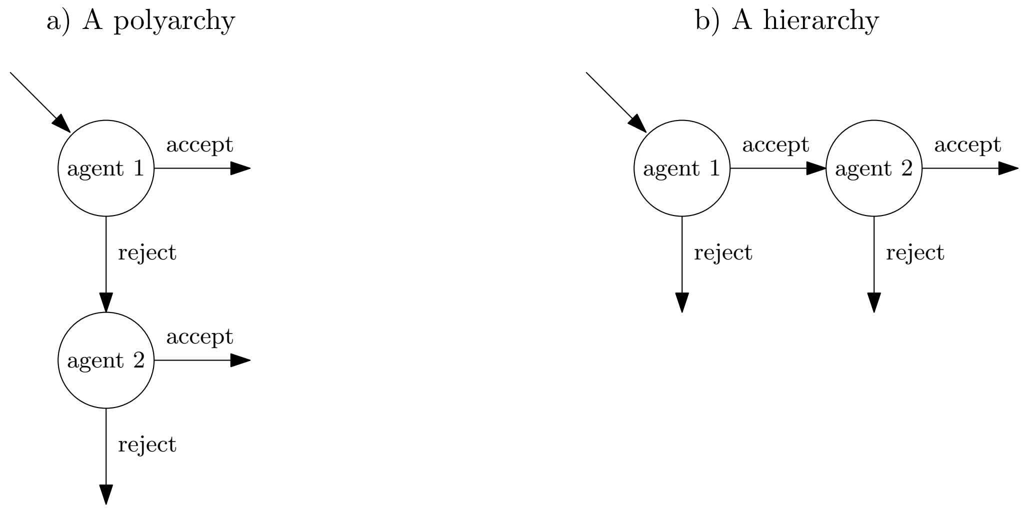

Figure 3: Sah and Stiglitz's hierarchy and polyarchy expressed using the graph notation.

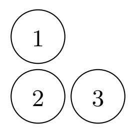

Figure 4: A simple example of an organization with three agents, which could represent the decision process used by three managers when determining whether to hire someone. Observe that the notation is unambiguous even if the accept and reject arrows have not been drawn.

OR and AND circuits in reliability theory (Rausand and Høyland, 2004): for project acceptance, a polyarchy requires the approval of agent 1 OR agent 2 whereas a hierarchy requires the approval of agent 1 AND agent 2.

The notation can be further simplified, since the arrows can be removed without loss of information (the arrows are not necessary because each agent has exactly one horizontal and one vertical arrow). Figure 4 shows how the simplified notation would describe the decision process used by three managers when determining whether or not to hire a technology consultant. In this figure, if agent 1 (e.g., the CEO) accepts, the consultant is hired immediately. If not, the hiring decision is pondered by agent 2 (e.g., the CFO), who may definitely reject the idea of hiring the consultant or pass the decision to agent 3 (e.g., the IT manager), who then makes the final hiring decision.

This notation can be used to describe a wide spectrum of decision-making structures. (See Appendix B for a formal definition of this set of organizations and a mechanism to exhaustively generate them; see Appendix C for a list of all the organizations with at most five agents that can be expressed using this notation.) However, this nomenclature is not capable of expressing the whole range of possible decision-making structures. For example, there is no way to represent a loop or an agent passing the decision on a project to two or more agents simultaneously. To distinguish the set of organizations that can be drawn on a grid from the set of all possible organizations, let us call the former "grid organizations."

The model is stated in sequential terms, but it can be used (with some care) to understand organizational processes that are not sequential. In fact, if agents are homogeneous then there is a one-to-one mapping between every organization that has a rectangular shape (i.e., , , , , , , , , , etc.) and the "committee structures" studied by Sah and Stiglitz (1988) because each of the rectangular organizations produce the same O&C errors as those committees.[6] For example, a 2×2 square structure ( ) will approve (reject) a project only if it was approved (rejected) by at least two agents. This is the exact same decision rule of a committee with three decision makers using the majority rule. More generally, any square grid organization corresponds to a majority-rule committee. This equivalence relationship is further explained near the end of Appendix B.

3.3 Computing Organization-Level Omission and Commission Errors

Before calculating the O&C errors of an arbitrary grid organization, it is first necessary to compute the screening function of the organization (i.e., the probability that a given organization will approve a project of a given quality q). Derivation of the organization-level screening function is facilitated by the theory of Markov chains, which is a method for computing any future state of a system that depends only on the system's current state and state transition probabilities (for a classical introduction to the subject see Feller, 1968). In the following analysis, a "state" is a variable that represents the location of the project being evaluated. Thus, in an organization with N agents there are $N + 2$ states, because at any time a project can be at one of the N agents, accepted, or rejected.

To compute the organization-level screening function via a Markov chain, the graphical notation

[6] Sah and Stiglitz's committees are an abstraction of real committees. Like in the current paper, Sah and Stiglitz model committees in a team-theoretic way: as voting by fallible decision makers. This leaves outside the model's scope such phenomena as groupthink (Janis, 1972), herding (Bikhchandani et al., 1992), and group polarization (Tindale et al., 2003:384). Further work could explore how to incorporate these complex phenomena into a model like the one presented here.

developed in Section 3.2 must be transformed into a transition matrix. Each cell (i, j) in this matrix represents the probability that a project will pass from state i to state j. For example, the organization shown in Figure 4 has the following transitions:

	1	2	3	4(A)	5(R)
1		$p_R(q)$		$p_A(q)$	
2			$p_A(q)$		$p_R(q)$
3				$p_A(q)$	$p_R(q)$
4(A)				1	
5(R)					1

In this example, columns 4 and 5 are used to encode transitions to the acceptance and rejection states, respectively. (More generally, for an organization with N individuals, this paper will always consider state $N + 1$ to be the organization's acceptance state and $N + 2$ to be the rejection state.) Hence, for example, the first row of the table can be read as "the project will pass from agent 1 to agent 2 with probability $p_R(q)$ (which happens when agent 1 rejects the project) or from agent 1 to the acceptance state with probability $p_A(q)$." The last two rows of the table (containing a 1 each) reflect that, once a project is accepted or rejected by the organization, it remains in that state (in Markov chain terms, these are "absorbing states"). Observe that every row adds up to 1, since each row represents a probability distribution.

Once a grid organization is represented as a transition matrix, Markov chain theory allows one to compute the organization's screening function directly (see equation (2) in Appendix A). The screening function of the organization with transition matrix $\mathbf{T}$ is denoted by $P_A(\mathbf{T}, q)$. The reverse of the screening function is the probability of rejecting a given project, and it is denoted by $P_R(\mathbf{T}, q) = 1 - P_A(\mathbf{T}, q)$.[7]

Given the organization-level probability of accepting a project of a given quality (i.e., $P_A(\mathbf{T}, q)$) and the probability distribution of the quality of the projects the organization can receive, the probabilities of omission and commission errors by a given organization $\mathbf{T}$ correspond to the shaded areas in Figure 5. The right area corresponds to the probability of rejecting good projects (omission errors) and the left area to the probability of accepting bad projects (commission errors). For

[7]Individual-level probabilities are denoted by lowercase p's and organization-level probabilities by uppercase P's.

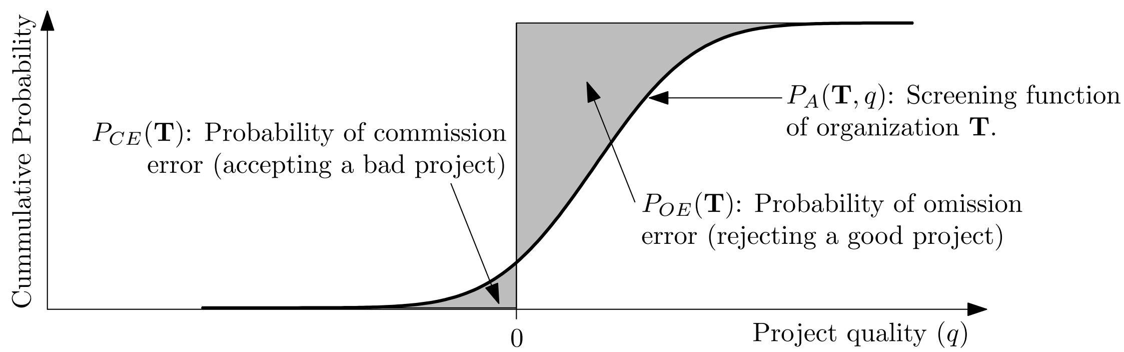

Figure 5: Screening function with omission and commission error areas shaded.

simplicity, the distribution of the q-values is assumed to be a uniform between $\underline{q}$ and $\overline{q}$. The formulas for the organization-level O&C errors of a given organization $\mathbf{T}$ are denoted $P_{OE}(\mathbf{T})$ and $P_{CE}(\mathbf{T})$ and derived in Appendix A (see equation (3)).

4 Results

This section examines results that stem directly from analyzing the model quantitatively. The main goal of this section is to develop an intuitive understanding of what the "design space" of the grid organizations looks like. This understanding is later invoked when explaining (in Section 5) the effects of organizational structure on exploration and exploitation.

Here the model is analyzed mostly by inspecting charts that plot all the grid organizations of a given size. Exploring the whole design space of the grid organizations, rather than some handpicked cases, allows for the discovery of relationships that would otherwise be difficult or impossible to determine; examples include estimating the proportion of organizations that are efficient and identifying structures that are particularly good or bad at achieving given levels of O&C errors. The algorithm that created all the grid organizations of a given size N is described in Appendix B, and its results for $N \leq 5$ are listed in Appendix C. The noise distribution used to create the figures in this section was $\tilde{n} \sim N(0, 1)$, and the distribution of project qualities used was $q \sim U[-3, 3]$. Robustness checks are presented in Section 4.1.

Figure 6 plots all the grid organizations with at most five agents according to their omission and commission errors. Note that the actual shape of each organization appears in the figure (e.g., is in one corner, in the other, near the middle, etc.). The x and y axes of this figure correspond (respectively) to the probability of an organization making an omission error and a commission error—that is, P_{OE} and P_{CE}. Several insights can be extracted from this chart.

First, and corroborating prior research (Sah and Stiglitz, 1986, 1988), the pure hierarchy and the pure polyarchy are at the extremes. For example, the five-agent hierarchy () is at the bottom right (this is the lowest commission error but the highest omission error). The opposite happens with the five-agent polyarchy ( ), which is at the top left corner. This figure also shows that "hybrid structures"—forms which are neither hierarchies nor polyarchies—populate the intermediate positions in the omission–commission plane. An example of a hybrid is (around 0.10, 0.02), which provides O&C errors that are in between those of the hierarchy and polyarchy of the same size (and , respectively). Structure could be optimal for an organization interested in achieving a low commission error but without incurring in the high omission error associated with a hierarchy. More generally, in the model we can describe the task of designing an optimal structure as a minimization problem: given the cost β_1 and β_2 of (respectively) omission and commission errors for a given organization, find the structure that minimizes the expected overall cost of the errors (i.e., $\arg\min_{\mathbf{T}}\{\beta_1 P_{OE}(\mathbf{T}) + \beta_2 P_{CE}(\mathbf{T})\}$). Graphically, the solution to this problem corresponds to the first structure in Figure 6 that would intersect a line with slope $-\frac{\beta_I}{\beta_{II}}$ approaching from the left.[8]

Second, the hybrid structures of a given size do not lie in a straight line connecting the corresponding hierarchy and polyarchy; instead, they form a convex space. This implies that, by arranging the same number of individuals in a different way, it possible not only to tradeoff one error for the other, but to decrease the sum of both errors. For example, the overall error (i.e., the sum of both errors, $P_{OE} + P_{CE}$) of is 0.1957 whereas the overall error of is 0.1109 (see Appendix C).

Third, not all organizations are efficient in the sense that, for a fixed level of one of the errors, a rational organization designer should always prefer the organization with the smallest value on the other error. For example, form is efficient because it is the only form that allows achieving such a small omission error; in contrast, form (around 0.05, 0.13) is inefficient because form

[8]The minimization problem could be more generally stated by considering expected losses that are not simply proportional to the error probabilities (i.e., E[cost of omission errors] + E[cost of commission errors], rather $\beta_1 P_{OE}(\mathbf{T}) + \beta_2 P_{CE}(\mathbf{T})$) and by including other costs such as the cost of evaluating the projects (e.g., setup costs, wages of the decision makers).

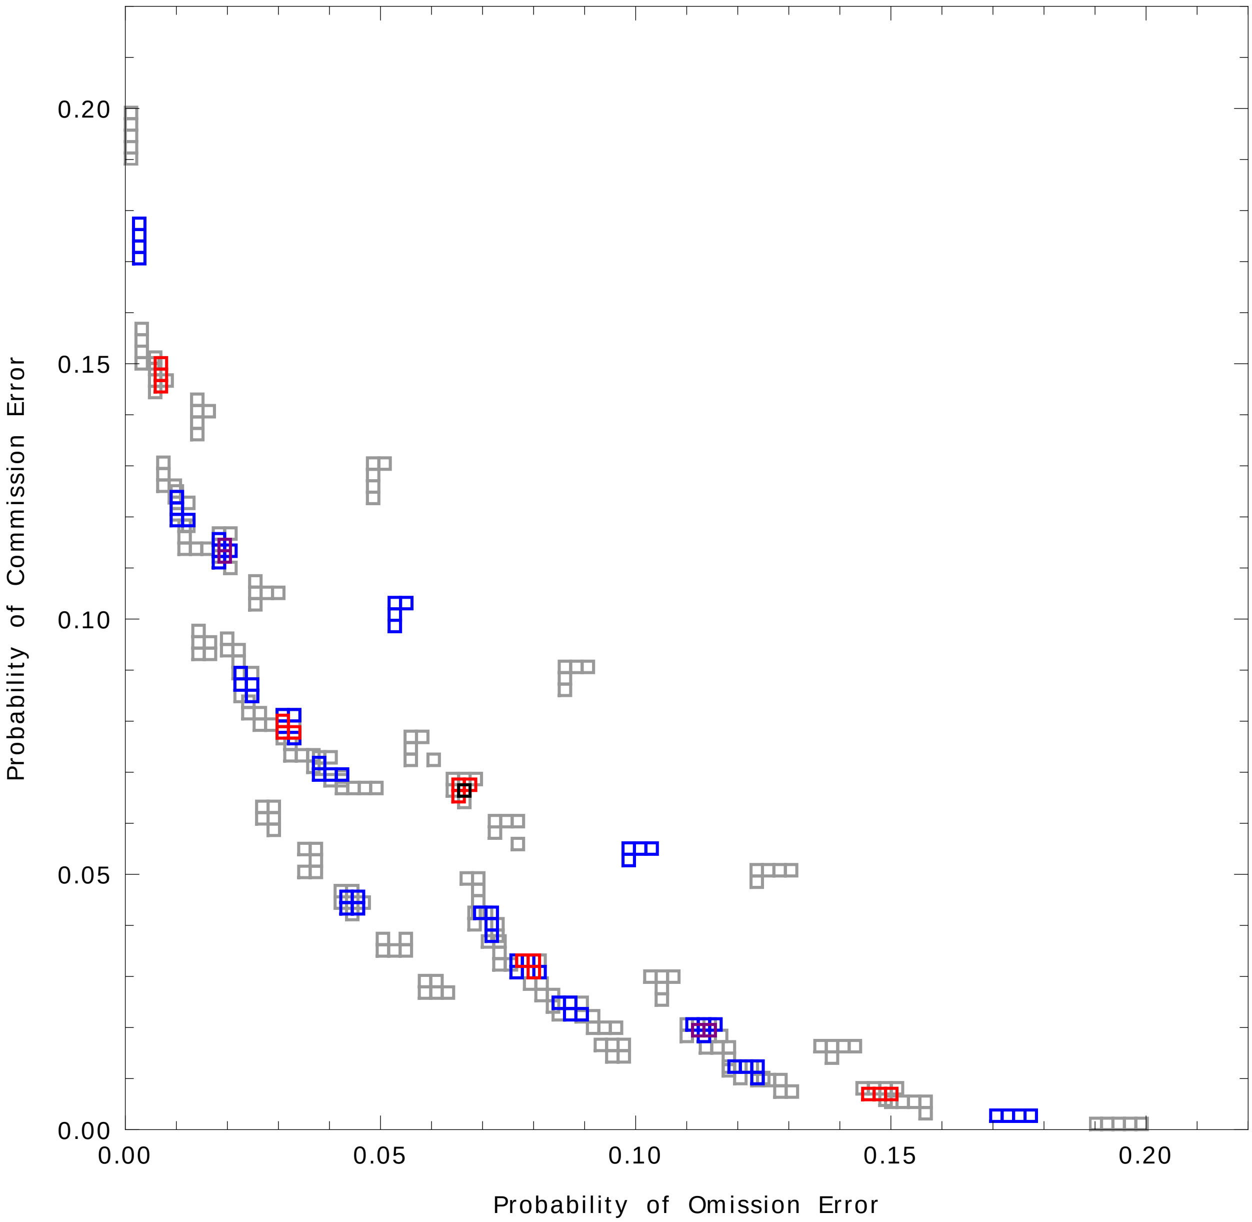

Figure 6: All grid organizations (with five or fewer agents) plotted in terms of their omission and commission errors; organizations with the same number of agents are drawn in the same color (shade of gray).

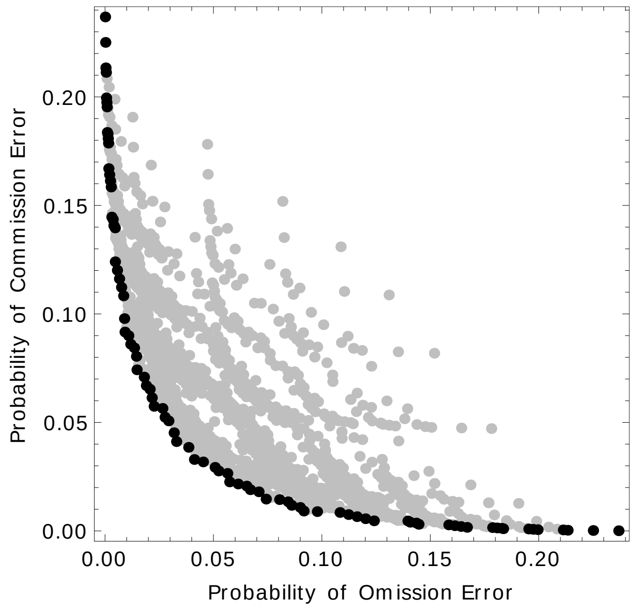

Figure 7: The efficient frontier for all grid organizations with eight or fewer agents; in this case, only 82 of the 1,982 organizations are on the frontier.

(around 0.01, 0.13) has less omission error at the same level of commission error. The set of efficient organizations, hereafter called the "efficient frontier," can be characterized as the set of all organizations that are not dominated or quasi-dominated by others.[9] The efficient frontier delineates an inescapable trade-off for the grid organizations of a given maximum size: past the frontier, it is impossible to decrease one error without increasing the other.

Fourth, the proportion of efficient organizations is small. Figure 7 plots all the organizations that can be formed with up to eight agents and highlights in black the organizations at the efficient frontier. Here the number of efficient organizations is much smaller than the total number of organizations (82 out of 1,982). One implication of this low proportion is that randomly selected organizations will seldom be efficient. Organizational forms should therefore be selected with care, because otherwise ending up with a suboptimal form is likely.

Fifth, the effect of adding more agents is larger on the middle part of the frontier than at the extremes. From Figure 8 it is clear that—although the successive frontiers advance a considerable amount toward the bottom left corner—this is not true for the top left and bottom right sections of the frontiers, where frontiers with as few as three agents perform similarly to frontiers with more

[9] Organization A (located at a_I, a_{II}) is dominated by organization B (located at b_I, b_{II}) if $b_I < a_I$ and $b_{II} < a_{II}$. Organization A is quasi-dominated by organization B if $b_I \leq a_I$ and $b_{II} < a_{II}$, or if $b_I < a_I$ and $b_{II} \leq a_{II}$.

agents. More precisely, the marginal effect of adding more agents is greater in the midsection of the frontier than at its extremes. Looking at the actual structures that lie in the midsection of successive frontiers reveals interesting details. In particular, the structures that lie at the midpoint of the frontiers (i.e., where each frontier would intersect a 45-degree line from the origin) include the square grid organizations (i.e., , , , etc.). As explained in Section 3.2, these structures correspond to majority-rule committees. Perhaps the widespread use of majority voting is due in part to the uniqueness of its position on the frontier: lying exactly where the frontier is closest to the origin and thereby jointly minimizing O&C errors.[10]

Finally, the only way to "move" the frontier into new territory is by adding more agents. Figure 8 shows how the frontiers created by the organizations with more agents reach smaller O&C errors. Note that more agents can add value only when they are arranged in "intelligent" ways. For example, from Figure 6 it is clear that organization (around 0.07, 0.07) carries the same errors as organization, so the two extra agents of are completely ineffectual. In this case, both configurations are equivalent because the first agent is redundant: irrespective of what she does, the project will be evaluated by another agent who has absolute power to accept or reject the project.[11]

4.1 Robustness

This section studies the robustness of the findings with respect to changes in the main parameters of the model: (i) the threshold above which individuals decide to approve projects; (ii) the noisiness of the decision makers (parameterized in terms of the mean and standard deviation of the individuals'

[10] Although they are beyond the scope of this paper, it is interesting to note that there exist other decision rules capable of jointly minimizing O&C errors. One such rule is averaging individuals' evaluations and then accepting the project if that average is positive. When we compare averaging with voting while using the same assumptions about noise and project distribution, averaging produces fewer O&C errors than majority voting using the same number of agents. But this comparison may not be the most apt when one considers that the performance of averaging is highly sensitive to factors that have little effect on voting. For instance, a noisy or malicious individual can completely derail averaging (by reporting an evaluation so negative that the average, too, is negative and thus causing the project to be rejected); in contrast, voting is more robust to such derangement (i.e., the effect of an individual like this is capped at one vote). For this reason, heterogeneity and incentives play a central role in the analysis of averaging. Further research could explore the advantages and disadvantages of other families of structures with respect to the grid organizations described in this paper.

[11] For more complex structures it is not usually possible, based on visual inspection alone, to determine whether or not two structures are equivalent; doing so requires the mathematical method developed in Section 3 (see, e.g., how different the six forms associated to organization 47 appear in Appendix C). Similarly, determining whether a given structure will lie on the frontier seems possible only via the mathematical method (except for rectangular grid organizations, which are efficient).

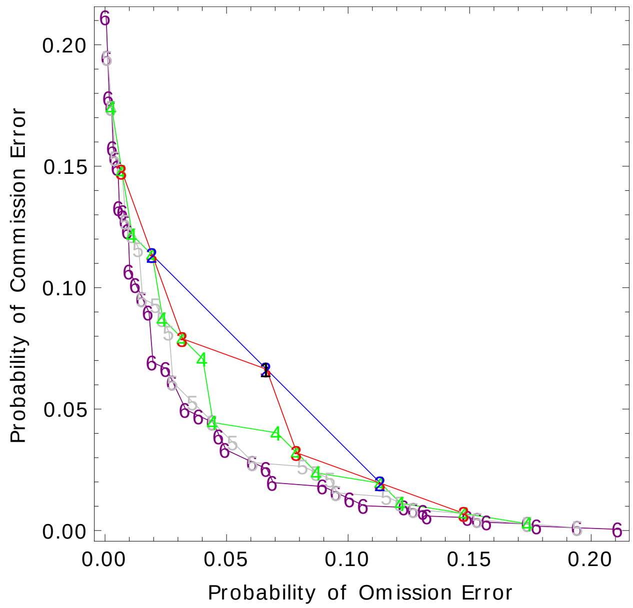

Figure 8: The six efficient frontiers for organizations with one to six agents.

noise, i.e., $\tilde{n} \sim N(\mu, \sigma)$); and (iii) the quality distribution of the incoming projects (parameterized in terms of the range $[\underline{q}, \overline{q}]$ of incoming project qualities, i.e., $q \sim U[\underline{q}, \overline{q}]$). In addition, this section addresses the effects of combining heterogeneous decision makers in the same organization and of introducing a correlation in the perceptions of the agents.

Most of the robustness checks in this section are illustrated by figures consisting of three panels: (a) a baseline scenario (with the same parameters as those used to generate Figure 6; i.e., agent's approval threshold set at 0, $\mu = 0$, $\sigma = 1$, and $q \sim U[-3, 3]$); (b) a scenario that diminishes the parameter being studied; and (c) a scenario that increases the same parameter.

An important result of these tests is that, although the exact locations of the structures in the omission–commission plane vary as the value of the parameters change, the ordering relationships among the structures remain constant. In other words, these tests show that the efficient frontier, the set of dominated structures, and the overall shape of the design space of organizations remain invariant under a broad range of conditions.

Robustness with respect to the approval threshold. Recall that, in the model presented so far, an agent recommends approval if she perceives its value to be above zero ($q' > 0$). In these robustness tests, the zero is replaced by an approval threshold (i.e., now the individual recommends approval if $q' >$ threshold). Comparing the first two panels in Figure 9 shows that, when the approval threshold

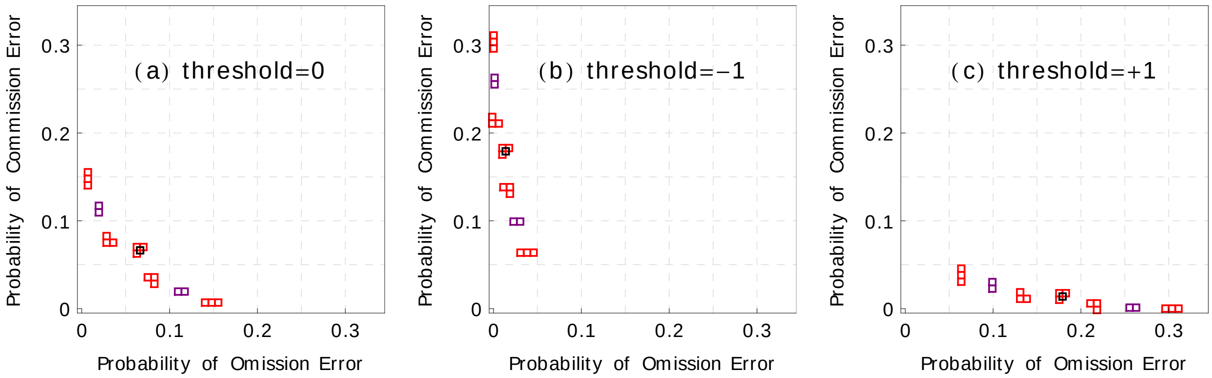

Figure 9: Effects of varying the project approval threshold.

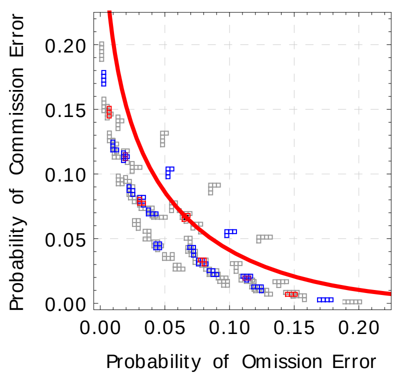

Figure 10: The curve shows all the combinations of O&C errors that are attainable by varying the threshold of the individual decision maker.

is decreased, commission errors increase and omissions errors decrease. This happens because, as the threshold decreases, it becomes easier to approve bad projects, and also, since more projects are approved, it becomes harder to miss good opportunities. Panel (c) in which the threshold is increased with respect to panel (a), tells the reverse story.

Figure 10 shows that some—but not all—of the effects of structure can be emulated by varying the approval threshold of one decision maker. The curve in this figure marks all the combinations of O&C errors that can be attained by varying the approval threshold of the structure with one individual (structure ▪). The mathematical explanation for this phenomenon is that threshold changes can cause a screening function to be shifted horizontally (i.e., swapping omission for commission errors, or vice versa) yet structural changes affect the *slope* of the screening function—allowing one not only to swap one error for the other but also to decrease both errors simultaneously.

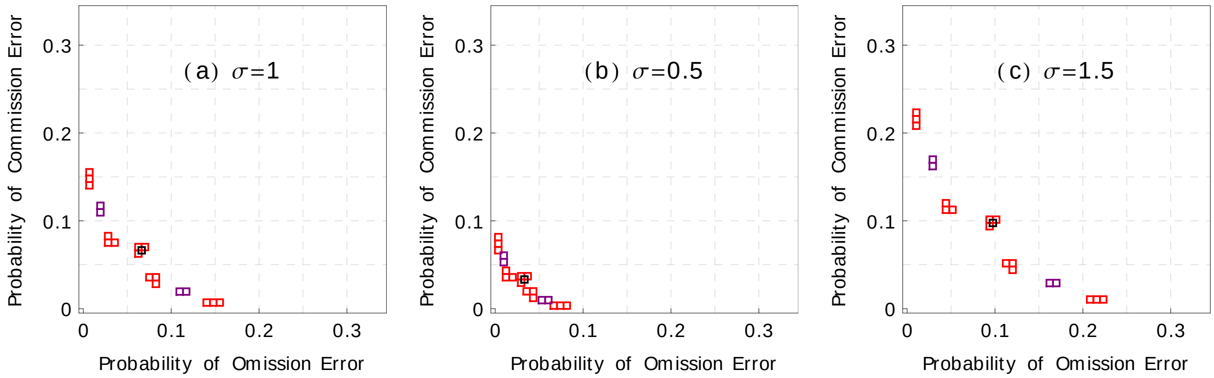

Figure 11: Effects of varying the magnitude of the agents' noise (σ).

Because noise is additive in the model, shifting the approval threshold can also be interpreted as shifting the noise distribution in the opposite direction. For instance, the effect of moving the approval threshold from 0 to 1 is equivalent to shifting the center of the agents' noise distribution (μ) from 0 to -1. One application of this equivalence relationship is that an organization composed of overoptimistic managers can be "de-biased" by increasing the approval threshold.

Robustness with respect to the magnitude of the agents' noise. Not surprisingly, the panels in Figure 11 show that, as the magnitude of the agent's noise (σ) increases (compare panels (a) and (c)), the expected O&C errors also increase.

Robustness with respect to the quality of incoming projects. Figure 12 shows that, as the environment becomes more munificent (compare panel (a), where the incoming projects take values from -3 to 3, to panel (c), where the distribution ranges from -1 to 5), it becomes harder to make commission errors. On the other hand, if the environment is less munificent (as in panel (b)), it becomes harder to make omission errors.

The effect of varying the functional form of the incoming projects' probability distribution was explored by replacing the uniform distribution with a Normal distribution and then testing a broad range of means and standard deviations. The results were qualitatively the same as those described in Figure 12 (the ordering relationship among the structures remained unchanged, and the effects of shifting the center of the distribution were in the same direction).

Effect of heterogeneous agents. The effect of heterogeneous agents was first explored by introducing a parameter that controlled the similarity of perceptions of the decision makers in the

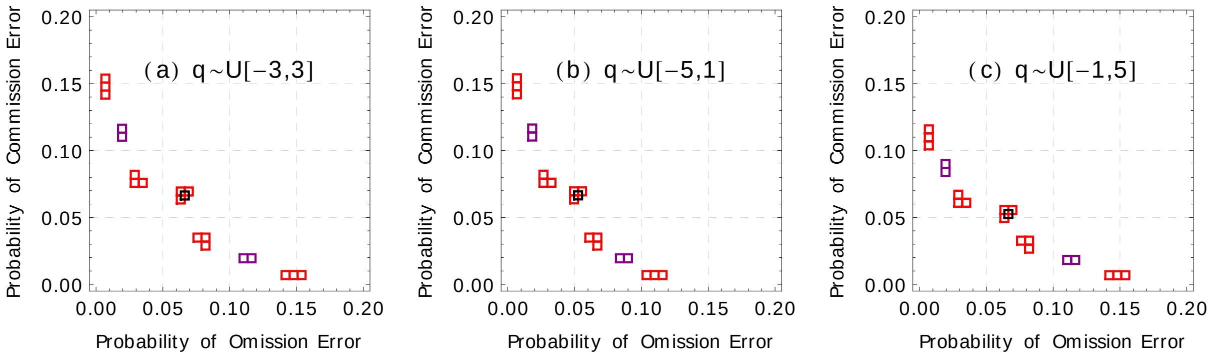

Figure 12: Effects of varying the quality of the incoming projects.

organization (i.e., as done by Clemen and Winkler, 1985). Technically, this parameter controls the correlation between the noise functions of the agents: when the correlation is 0 the perceptions are completely independent; when the correlation is 1, all agents perceive the same. The effect of gradually varying the correlation from 0 to 1 is to gradually collapse all the organizations in the omission–commission plane into the position of the structure with a single decision maker. In other words, increasing the correlation has an effect akin to reducing the number of decision makers.

Heterogeneity was studied also by having each agent draw noise from a different distribution (i.e., instead of all noises coming from a $N(\mu, \sigma)$, now each individual draws from a $N(\mu_i, \sigma_i)$). The effect of increasing heterogeneity in the noise distribution (e.g., by increasing z in $\mu_i \sim U[-z, z]$ or in $\sigma_i \sim U[0, z]$) is to increase both O&C errors. The reason is that adding heterogeneity to the noise distribution increases the overall level of noise of the agents, which in turn makes it more likely that their individual evaluations will be more extreme (either positive and negative). Thus, for instance, a hierarchy (in which everybody has veto power) will be more likely to reject projects owing to the influence of its most negative member. The converse holds in a polyarchy: with more noise, its most positive member will be more likely to accept. In other words, increasing the noise distribution's heterogeneity has an effect akin to increasing the overall noise level.

5 Discussion

This section uses the results just developed to illuminate the questions motivating the paper. The discussion is structured in three parts, ordered from the most practical to the most theoretical implications. First, it uses the theory developed in the paper to analyze a mini-case on Google versus Microsoft and discusses managerial implications. Second, it presents avenues for theoretical and empirical work and provides testable hypotheses. Finally, it connects the paper's findings to open questions in organization design. In the ensuing discussion, recall that all the relationships uncovered so far are probabilistic—that it, they are not bound to occur every time, but should have average effects observable given a suitably large number of observations.

5.1 Google versus Microsoft from the Model's Perspective

The issue of what organizational form is most appropriate to foster E&E has been a topic of much practitioner interest, especially in the high-tech industry (e.g., Noyce, 1978; Gerstner, 2003; Garvin and Levesque, 2005) and in hazardous industries (e.g., Perrow, 1984; Weick and Roberts, 1993; Vaughan, 1996). This paper began with IBM's CEO pondering the effects of one decision that involved structure and innovation. Now, armed with a more robust theoretical framework, I analyze the competition between Google and Microsoft in order to illustrate how the ideas developed in this paper can shed light on significant questions of structure, exploration, and exploitation.

Recently, the business press has hailed Google for producing a large number of innovations while Microsoft, which has a larger R&D budget, has not introduced as many innovations. A possible explanation for this difference is that the damage a bad project may cause (the cost of a commission error) is greater for Microsoft than for Google. In the first place, Microsoft has a broader product line and so the possibility of cannibalizing other projects is higher. Second, Microsoft has a larger and older code base, so new projects are more likely to have adverse interactions with existing code. And unlike Google's web-based platform, most Microsoft software requires installation and is difficult to update; thus, the cost of replacing malfunctioning software is therefore orders of magnitude higher for Microsoft. In sum, commission errors (which could disrupt current operations) are probably a much greater concern for Microsoft than for Google.

One could well hypothesize that Microsoft has created a hierarchical structure so that fewer (potentially disruptive) commission errors slip through. Yet the main side effect of this structure is that fewer good projects are approved. On the other hand, Google faces a smaller cost for commission errors and so its relevant cost is missing good opportunities (omission errors); hence the logical decision is to use a structure more closely resembling a polyarchy. In fact, Google's policy of allowing its engineers to spend 20% of their time on personal R&D projects, which do not require further approval, is a visible cue of Google's more polyarchy-like structure.

From the viewpoint of this paper, two approaches could enable Microsoft to increase its level of exploration. The first is to reduce the cost of the commission errors—for example, by moving to a more modular code base or by changing the distribution medium; the second is to use a hybrid structure that could keep both omission and commission errors low.

5.2 Testable Hypotheses and Further Work

The model suggests some relationships that merit further empirical investigation. A first testable relationship is that as the number of decision makers is increased, the efficient frontier achieves new terrain. That is, adding new individuals allows to decrease both omission and commission errors. This effect is most visible in the middle section of the frontier (see Figure 8). Because the majority-rule committees lie in the midsection of the frontier, the model predicts that organizations that make decisions using the majority-rule with a large number of individuals, would achieve a better performance than organizations using the majority-rule with fewer individuals. A particular case is that, keeping everything else constant, a majority-rule committee should make fewer errors (both omission and commission) than a single decision maker. In hypothesis form:

> **H1**: A majority-rule committee will make both fewer omission errors and fewer commission errors than a single decision maker.

Perhaps the most suggestive characteristic of the efficient frontier (Figure 6) is that it is convex. Thus, for example, structure (near 0.044, 0.044) achieves smaller errors than any linear combination of structures of the same size. For instance, a 50-50 combination between (near 0.002, 0.174) and (near 0.174, 0.002) would achieve position 0.088, 0.088 ($= \frac{0.174+0.002}{2}$), which is twice the errors of structure . In practice, forms can be combined by splitting the stream of projects into two halves and sending each half to each of the two structures—a process consistent with the recommendation of the ambidexterity literature of alternating between exploration and exploitation either temporally or spatially (Raisch and Birkinshaw, 2008:389). Interestingly, the convexity of the efficient frontier implies that hybrid organizational forms (those in between the extremes of hierarchy and polyarchy) have a unique power to achieve error levels that are unattainable by alternating between the hierarchy and the polyarchy. If this is so, hybrid organizations deserve further attention, as they could provide a particularly effective way to achieve ambidexterity. A hypothesis aimed at testing if the efficient frontier is indeed convex is the following:

> **H2**: The omission (commission) error of the majority-rule committee of size N will be smaller than the averaging of the omission (commission) errors of the hierarchy and the polyarchy of size N.

Further empirical work could test the predictions of the model presented in this paper (see Csaszar (2012) for a first step in this direction) and contrast the effects of decision-making structure with the effects of other mechanisms known to affect E&E. Because models in the spirit of Sah and Stiglitz (1986; 1988) have remained mostly unexplored in the organizations literature, there are many possible opportunities for further theoretical work. These include: (a) analyzing an even broader family of structures, such as pyramidal structures (Radner, 1993) or circular decision making (Romme, 2004); (b) analyzing a broader set of tasks, such as organizations that face a budgetary constraint and can only pick a limited number of projects, or organizations whose agents can communicate in a richer, non binary way (e.g., continuous signals or three-level signals with a "pass" alternative when the agent is unsure); and (c) studying the existence of non trivial interactions with other phenomena amenable to modeling, such as strategic behavior or learning.

5.3 Conclusions

From a theoretical standpoint, this research speaks to several calls for introducing more process-level explanations in the organizations literature. For example, Maritan and Schendel (1997:259) noted that "there has been surprisingly little work that has explicitly examined the link between the processes by which strategic decisions are made and their influence on strategy," and Pettigrew (1992:171) remarked that most of the current research on boards make "great inferential leaps" when connecting input to output variables. An example demonstrating the usefulness of this paper's focus on processes concerns the debate between Beckman et al. (2004) and Rothaermel and Deeds (2004), who reported contradictory findings on the effect of organization size on exploration. The current paper suggests that both results are possible depending on the structure used (i.e., larger polyarchies would explore more whereas larger hierarchies would explore less).

From a practical standpoint, this research sheds light on how organizations can compensate for the shortcomings of individuals and also allows to address several managerial concerns, such as: Is it true that hierarchy hampers innovation? What organization is needed to support exploitation? What organizational structures are conducive to more exploration? With respect to this last question, an important application area is how established organizations can be enabled to exhibit traits that are usually associated with entrepreneurial ventures. The 9/11 Commission Report contains an eloquent call for this sort of transformation: "Imagination is not a gift usually associated with bureaucracies. [...] [I]t is therefore crucial to find a way of routinizing, even bureaucratizing, the exercise of imagination." (National Commission on Terrorist Attacks upon the United States, 2004:344).

One implication of special interest to academics concerns the review process used by most journals.[12] Arguably, the review process looks more like a hierarchy than a polyarchy (e.g., in many cases the editor, associate editor, and reviewers must all agree about a paper before it will be published), which implies that the review process will make more omission than commission errors. If the goal of academia is to explore new ideas and produce innovative research, then this screening mechanism seems counterproductive. Yet moving to the other extreme (i.e., using a polyarchy to minimize omission errors) is not a good solution because then many bad papers would be published, producing bad science and wasting readers' time. Therefore, this paper calls for journals to make a conscious assessment of their desired level of omission and commission errors and then to design a review process that best accomplishes that goal.

The current paper has contributed to the literature in several respects. First, it has provided an alternative and parsimonious logic that serves to explain the relationship between the structure of an organization and its ability to achieve E&E. Second, it has reported on the existence of a wide range of "hybrids"—structures that allow not only trading off one error against the other but also for achieving a smaller error overall. Hybrids can offer a mechanism for achieving a high

[12] I thank an anonymous referee for suggesting this topic.

degree of *both* exploration and exploitation (i.e., ambidexterity) at the same time and place. Third, this paper has provided a methodology that allows one to quantify the expected O&C errors of a broad range of decision-making structures. This methodology can be used as a tool for examining the relative advantages and disadvantages of different organization designs. Fourth, by developing the concept of an "efficient frontier," this paper has shown that some organization designs actually dominate others and that, with even relatively small numbers of agents, the proportion of efficient structures is small.

Given that organizations are so central to human life, even a slight improvement in their performance may have important payoffs. This paper has aimed to achieve this by advancing a small step toward an "organizational engineering," which could be used similarly to how architects use structural engineering when they design buildings.

References

Adler, P. S., B. Borys. 1996. Two types of bureaucracy: Enabling and coercive. *Administrative Science Quarterly* **41**(1) 61–89.

Beckman, C. M., P. R. Haunschild, D. J. Phillips. 2004. Friends or strangers? Firm-specific uncertainty, market uncertainty, and network partner selection. *Organization Science* **15**(3) 259–275.

Benner, M. J., M. L. Tushman. 2003. Exploitation, exploration, and process management: The productivity dilemma revisited. *Academy of Management Review* **28**(2) 238–256.

Berger, J. O. 1985. *Statistical Decision Theory and Bayesian Analysis*. 2nd ed. Springer-Verlag, New York.

Bikhchandani, S., D. Hirshleifer, I. Welch. 1992. A theory of fads, fashion, custom, and cultural-change as informational cascades. *Journal of Political Economy* **100**(5) 992–1026.

Brown, S. L., K. M. Eisenhardt. 1997. The art of continuous change: Linking complexity theory and time-paced evolution in relentlessly shifting organizations. *Administrative Science Quarterly* **42**(1) 1–34.

Burns, T., G. M. Stalker. 1961. *The Management of Innovation*. Tavistock Publications, London, UK.

Burton, R. M., G. DeSanctis, B. Obel. 2006. *Organizational Design: A Step-by-Step Approach*. Cambridge University Press, Cambridge, UK.

Burton, R. M., B. Obel. 2004. *Strategic Organizational Diagnosis and Design: The Dynamics of Fit*. 3rd ed. Kluwer, Boston, MA.

Child, J. 1984. *Organization: A Guide to Problems and Practice*. Harper & Row, London, UK.

Christensen, C. M. 1997. *Innovator's Dilemma: When New Technologies Cause Great Firms to Fail*. Harvard Business School Press, Boston, MA.

Christensen, M., T. Knudsen. 2002. The architecture of economic organization: Toward a general framework. LINK Working Paper.

Christensen, M., T. Knudsen. 2010. Design of decision making-organizations. *Management Science* **56**(1) 71–89.

Clemen, R. T., R. L. Winkler. 1985. Limits for the precision and value of information from dependent sources. *Operations Research* **33**(2) 427–442.

Csaszar, F. A. 2012. Organizational structure as a determinant of performance: Evidence from mutual funds. *Strategic Management Journal* **33**(6) 611–632.

Cyert, R. M., J. G. March. 1963. *A Behavioral Theory of the Firm*. Prentice-Hall, Englewood Cliffs, NJ.

Daft, R. L., A. Y. Lewin. 1993. Where are the theories for the new organizational forms: An editorial essay. *Organization Science* **4**(4) i–vi.

Ethiraj, S. K., D. A. Levinthal. 2004. Modularity and innovation in complex systems. *Management Science* **50** 159–173.

Feller, W. 1968. *An Introduction to Probability Theory and its Applications*. Wiley, New York.

Foss, N. J. 2003. Selective intervention and internal hybrids: Interpreting and learning from the rise and decline of the Oticon spaghetti organization. *Organization Science* **14**(3) 331–349.

Galbraith, J. R. 1973. *Designing Complex Organizations*. Addison-Wesley, Reading, MA.

Galbraith, J. R. 1977. *Organization Design*. Addison-Wesley, Reading, MA.

Garicano, L., R. A. Posner. 2005. Intelligence failures: An organizational economics perspective. *Journal of Economic Perspectives* **19**(4) 151–170.

Garud, R., P. R. Nayyar, Z. Shapira. 1997. Technological choices and the inevitability of errors. R. Garud, P. R. Nayyar, Z. Shapira, eds., *Technological Innovation: Oversights and Foresights*, chap. 3. Cambridge University Press, New York, 20–40.

Garvin, D. A., L. C. Levesque. 2005. *Emerging Business Opportunities at IBM (A)*. Harvard Business School Case, Cambridge, MA.

Gerstner, L. V. 2003. *Who Says Elephants Can't Dance? Leading a Great Enterprise Through Dramatic Change*. HarperBusiness, New York.

Green, D. M., J. A. Swets. 1966. *Signal Detection Theory and Psychophysics*. Wiley, New York.

Janis, I. L. 1972. *Victims of Groupthink: A Psychological Study of Foreign-Policy Decisions and Fiascoes*. Houghton Mifflin, Boston, MA.

Kahneman, D., A. Tversky. 1979. Prospect theory: Analysis of decision under risk. *Econometrica* **47**(2) 263–291.

Knudsen, T., D. A. Levinthal. 2007. Two faces of search: Alternative generation and alternative evaluation. *Organization Science* **18**(1) 39–54.

Knuth, D. E. 2005. *The Art of Computer Programming, Volume 4, Fascicle 2: Generating All Tuples and Permutations*. Addison-Wesley, Reading, MA.

Langlois, R. N. 1997. Cognition and capabilities: Opportunities seized and missed in the history of the computer industry. R. Garud, P. R. Nayyar, Z. Shapira, eds., *Technological Innovation: Oversights and Foresights*, chap. 3. Cambridge University Press, New York, 71–94.

Lawrence, P. R., J. W. Lorsch. 1967. *Organization and Environment: Managing Differentiation and Integration*. Graduate School of Business Administration, Harvard University, Boston, MA.

Lerner, J. 1994. The syndication of venture capital investments. *Financial Management* **23**(3) 16–27.

Levinthal, D. A. 1997. Adaptation on rugged landscapes. *Management Science* **43**(7) 934–950.

Macmillan, N. A., C. D. Creelman. 2004. *Detection Theory: A User's Guide*. 2nd ed. Lawrence Erlbaum Associates, Mahwah, NJ.

March, J. G. 1991. Exploration and exploitation in organizational learning. *Organization Science* **2** 71–87.

March, J. G., H. A. Simon. 1958/1993. *Organizations*. John Wiley & Co., New York.

Maritan, C. A., D. E. Schendel. 1997. Strategy and decision processes: What is the linkage? V. Papadakis, P. Barwise, eds., *Strategic Decisions*. Kluwer Academic Publishers, London, UK.

Marschak, J., R. Radner. 1972. *Economic Theory of Teams*. Yale University Press, New Haven, CT.

Miles, R. E., C. C. Snow. 1978. *Organizational Strategy, Structure, and Process*. McGraw-Hill, New York.

Moore, E. F., C. E. Shannon. 1956/1993. Reliable circuits using less reliable relays (I and II). N. J. A. Sloane, A. D. Wyner, eds., *Claude Elwood Shannon: Collected Papers*. Wiley-IEEE Press, New York, 796–830.

National Commission on Terrorist Attacks upon the United States. 2004. *The 9/11 Commission Report: Final Report of the National Commission on Terrorist Attacks Upon the United States*. Norton, New York.

Nickerson, J. A., T. R. Zenger. 2002. Being efficiently fickle: A dynamic theory of organizational choice. *Organization Science* **13**(5) 547–566.

Noyce, R. N. 1978. Innovation: The fruit of success. *Technology Review* **80**(4) 24–27.

O'Brien, J P. 2003. The capital structure implications of pursuing a strategy of innovation. *Strategic Management Journal* **24**(5) 415–431.

Perrow, C. 1984. *Normal Accidents: Living With High-Risk Technologies*. Basic Books, New York.

Pettigrew, A. M. 1992. On studying managerial elites. *Strategic Management Journal* **13**(Special issue) 163–182.

Pettigrew, A. M., E. M. Fenton. 2000. *Innovating Organization*. Sage, London, UK.

Puranam, P., B. C. Powell, H. Singh. 2006. Due diligence failure as a signal detection problem. *Strategic Organization* **4**(4) 319–348.

Radner, R. 1993. The organization of decentralized information processing. *Econometrica* **61**(5) 1109–1146.

Raisch, S., J. Birkinshaw. 2008. Organizational ambidexterity: Antecedents, outcomes, and moderators. *Journal of Management* **34**(3) 375–409.

Rausand, R., A Høyland. 2004. *System Reliability Theory: Models, Statistical Methods, and Applications*. 2nd ed. Wiley-Interscience, Hoboken, NJ.

Roberts, K. H. 1990. Some characteristics of one type of high reliability organization. *Organization Science* **1**(2) 160–176.

Romme, A. G. L. 2004. Unanimity rule and organizational decision making: A simulation model. *Organization Science* **15**(6) 704–718.

Rothaermel, F. T., D. L. Deeds. 2004. Exploration and exploitation alliances in biotechnology: A system of new product development. *Strategic Management Journal* **25**(3) 201–221.

Sah, R. K., J. E. Stiglitz. 1986. The architecture of economic systems: Hierarchies and polyarchies. *American Economic Review* **76**(4) 716–727.

Sah, R. K., J. E. Stiglitz. 1988. Committees, hierarchies and polyarchies. *Economic Journal* **98**(391) 451–470.

Sheremata, W. A. 2000. Centrifugal and centripetal forces in radical new product development under time pressure. *Academy of Management Review* **25**(2) 389–408.

Siggelkow, N. 2001. Change in the presence of fit: The rise, the fall, and the renaissance of Liz Claiborne. *Academy of Management Journal* **44**(4) 838–857.

Siggelkow, N., D. A. Levinthal. 2003. Temporarily divide to conquer: Centralized, decentralized, and reintegrated organizational approaches to exploration and adaptation. *Organization Science* **14**(6) 650–669.

Simon, H. A. 1947/1997. *Administrative Behavior*. 4th ed. The Free Press, New York.

Tindale, R. S., T. Kameda, V. B. Hinsz. 2003. Group decision making. M. A. Hogg, J. Cooper, eds., *The SAGE Handbook of Social Psychology*. Sage, London, UK, 381–406.

Tushman, M. L., C. A. O'Reilly. 1996. Ambidextrous organizations: Managing evolutionary and revolutionary change. *California Management Review* **38**(4) 8–30.

Van Zandt, T. 1999. Decentralized information processing in the theory of organizations. *Contemporary economic issues, Volume 4: Economic design and behavior* 125–160.

Vaughan, D. 1996. *The Challenger Launch Decision: Risky Technology, Culture, and Deviance at NASA*. University of Chicago Press, Chicago, IL.

von Neumann, J. 1956. Probabilistic logics and the synthesis of reliable organisms from unreliable components. C. E. Shannon, J. McCarthy, eds., *Automata Studies*. Princeton University Press, Princeton, NJ, 43–98.

Weick, K. E., K. H. Roberts. 1993. Collective mind in organizations: Heedful interrelating on flight decks. *Administrative Science Quarterly* **38**(3) 357–381.

Zenger, T. R., W. S. Hesterly. 1997. The disaggregation of corporations: Selective intervention, high-powered incentives, and molecular units. *Organization Science* **8**(3) 209–222.

A Appendix: Mathematical Derivation of Organization-Level Metrics

A.1 Individual's Screening Function

An individual perceives a project of quality q as a noisy signal $q' = q + \tilde{n}$, where $\tilde{n}$ is distributed according to probability distribution f. The individual will approve a project if the quality he perceives is positive (i.e., $q' > 0$), which happens with probability

$$p_A(q) = \int_0^\infty f(x - q)\,\mathrm{d}x;$$

he will reject the project with probability $p_R(q) = 1 - p_A(q)$. For succinctness, these probabilities can be rewritten in terms of F, the cumulative distribution of f, as

$$p_A(q) = 1 - F(-q) \quad \text{and} \quad p_R(q) = F(-q).$$

Furthermore, if the noise $\tilde{n}$ is symmetric around zero, then

$$p_A(q) = F(q) \quad \text{and} \quad p_R(q) = 1 - F(q). \tag{1}$$

In the paper $f \sim N(\mu, \sigma)$, so $F(q) = \Phi(\mu, \sigma)\big|_q$.

A.2 Organization's Screening Function

To compute the probability that a grid organization will approve a project (also known as the organization's screening function), a Markov chain is used. The transition probabilities of the Markov chain are given by the possible paths of acceptances and rejections than connect the agents (see Section 3.3). For instance, form (i.e., the structure in Figure 4 and in the transition table in Section 3.3) has transition matrix

$$\mathbf{T} = \begin{pmatrix} 0 & 1-F(q) & 0 & F(q) & 0 \\ 0 & 0 & F(q) & 0 & 1-F(q) \\ 0 & 0 & 0 & F(q) & 1-F(q) \\ 0 & 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 0 & 1 \end{pmatrix}.$$

Because each individual of a grid organization can either accept or reject a project in each time period, after N time periods a grid organization with N individuals must have definitely accepted or rejected a project. Thus, the probability that a project of quality q is accepted by a given grid organization is

$$P_A(\mathbf{T}, q) = (\mathbf{T}^N)_{1,N+1} \tag{2}$$

In other words, cell $(1, N+1)$ of matrix $\mathbf{T}^N$ contains the screening function of the organization described by transition matrix $\mathbf{T}$. Analogously, the probability that organization $\mathbf{T}$ ends up rejecting a project of quality q is $P_R(\mathbf{T}, q) = 1 - P_A(\mathbf{T}, q)$. For the example described in the preceding transition matrix, $P_A(\mathbf{T}, q) = -F(q)^3 + F(q)^2 + F(q)$. Appendix C lists the screening functions of all the grid organizations with up to five individuals.

A.3 Organization's Omission and Commission Errors

If we assume that an organization $\mathbf{T}$ receives a project of quality q randomly drawn from the range $[\underline{q}, \overline{q}]$, then the probabilities of the project becoming an omission error or a commission error are (respectively)

$$P_{OE}(\mathbf{T}) = \frac{1}{\overline{q} - \underline{q}} \int_0^{\overline{q}} P_R(\mathbf{T}, q)\,\mathrm{d}q \quad \text{and} \quad P_{CE}(\mathbf{T}) = \frac{1}{\overline{q} - \underline{q}} \int_{\underline{q}}^{0} P_A(\mathbf{T}, q)\,\mathrm{d}q. \tag{3}$$

A.4 Other Organization-Level Metrics

If each error has an associated cost, then it is possible to derive the expected value of the errors as well as their standard deviations. If the cost of the errors is constant (the cost per omission and commission error is β_1 and β_2, respectively), then the expected value of the errors is

$$E_{OE}(\mathbf{T}) = \beta_1 P_{OE}(\mathbf{T}) \quad \text{and} \quad E_{CE}(\mathbf{T}) = \beta_2 P_{CE}(\mathbf{T});$$

and the standard deviation around these expected errors is

$$\sigma_{OE}(\mathbf{T}) = \beta_1\sqrt{P_{OE}(\mathbf{T})(1 - P_{OE}(\mathbf{T}))} \quad \text{and} \quad \sigma_{CE}(\mathbf{T}) = \beta_2\sqrt{P_{CE}(\mathbf{T})(1 - P_{CE}(\mathbf{T}))}.$$

Other metrics that can be useful for extending the current paper are the expected quality of the accepted projects $\mathbb{E}[q_i|\text{project } i \text{ is approved}]$, as well as order statistics of the accepted projects (i.e., assuming the firm receives M projects, but can only accept M' of them, what is the quality of the i-th accepted project, as well as the probability of it being an omission or a commission error).

B Appendix: Generating All Possible Organizations of a Given Size

All the possible grid organizations of exactly N agents were created by the combinatorial algorithm described here. This algorithm represents organizations as square binary matrices (e.g., organization is represented as $\left[\begin{smallmatrix}1 & 1\\ 0 & 0\end{smallmatrix}\right]$). The algorithm has three phases: (1) create all possible matrices with exactly N cells set to 1; (2) prune the matrices that do not represent valid grid organizations; and (3) group all the remaining organizations into equivalency classes.

Phase 1: Combinatorial generation of matrices. Since the longest grid organization of size N can have either N agents in a row or N agents in a column, it follows that any grid organization of size N can fit into an $N \times N$ matrix. By definition, grid organizations have an agent in the top left; hence any possible grid organization must belong to the set of matrices of size $N \times N$ with cell (1,1) occupied and $N-1$ other cells occupied. Thus, the total number of these matrices is $\binom{N^2-1}{N-1}$. In order to generate these matrices, an algorithm to create all the permutations of a multiset was used (see sec. 7.2.1.2 in Knuth, 2005).

Phase 2: Pruning the set of matrices to obtain the set of valid grid organizations. Phase 1 generates more matrices than there are organizations; for example, it generates a matrix with all the diagonal elements occupied. This is not a valid organization because, except for the first agent, the agents are "disconnected": they do not have an agent from which to receive projects (i.e., agent 1 accepts or rejects, but no other agent ever receives projects). Hence, all the matrices with at least one disconnected agent are removed from the set of generated matrices, which yields the set of all valid grid organizations of a given size.[13]

Phase 3: Grouping the grid organizations into equivalency classes. In order to focus the analysis on essentially different organizational forms, the organizations created in Phase 2 are grouped according to their screening functions. For example, forms , , , , , , , , and  are trivially equivalent; the reason is that adding blank rows or columns to an organization does not alter its transition matrix and so its screening function remains the same. Some organizations that at first sight look different can also be equivalent; for example, and both have the

[13]Formally, an agent in cell (i, j) is disconnected if there are no agents above it (in cells $(1, j), (2, j), \ldots, (i-1, j)$) nor to its left (in cells $(i, 1), (i, 2), \ldots, (i, j-1)$).

N	1	2	3	4	5	6	7	8	9	10
Explored matrices	1	3	28	455	10626	324632	$1.23 \cdot 10^7$	$5.53 \cdot 10^8$	$2.90 \cdot 10^{10}$	$1.73 \cdot 10^{12}$
Grid organizations (size N)	1	2	5	15	45	130	416	1368	4735	17370
Grid organizations (size $\leq N$)	1	3	8	23	68	198	614	1982	6717	24087

Table 1: Number of grid organizations as a function of the number of agents.

screening function $3F^2 - 2F^3$. For the purposes of this paper, all organizations that fall into a given equivalency class are treated as the same organization.

Observe that, when using the algorithm here described, finding all the grid organizations of a given size rapidly becomes a computationally intensive task because the number of matrices to be explored grows as $\binom{N^2-1}{N-1}$. For example, finding all the grid organizations with 10 agents requires exploring 1.73 trillion matrices of size 10×10. Table 1 shows the number of explored matrices and the number of distinct grid organizations of up to 10 agents.

An interesting subclass of organizations are the rectangular grid organizations—namely, organizations that exactly fit into a $n \times m$ grid (, , , , , , , etc.). These organizations have the same screening function as a committee of $n + m - 1$ members on which m votes is the minimum level of consensus required to either approve or reject a project (for example, is equivalent to a 4-member committee with a minimum consensus level of 3). This is because, in a rectangular grid organization, only decision paths with at least m approvals are accepted (exit through the right) and no path can have more than $m + n - 1$ decisions; and that is exactly the behavior of a committee of $m + n - 1$ members and consensus level m. The previous logic implies that the square grid organizations (i.e., , , , etc.) represent the majority-rule committees. The square grid organizations lie along the diagonal in Figure 6 as they require the exact same number of votes to approve as to reject a project, and thus have the same probability of making an omission as of making a commission error.

C Appendix: List of All Possible Grid Organizations with $N \leq 5$ and Their Associated Screening Functions

This table reports the probability of omission errors and commission errors (P_{OE} and P_{CE}, respectively), in addition to the screening functions, as computed using $\tilde{n} \sim N(0,1)$ and $q \sim U[-3,3]$; here F is shorthand for $F(q)$.

ID	Graph(s)	Screening Function	N	$\mathbf{P_{OE}}$	$\mathbf{P_{CE}}$
1		F	1	0.0664	0.0664
2		$2F - F^2$	2	0.0195	0.1134
3		F^2	2	0.1134	0.0195
4		$F^3 - 3F^2 + 3F$	3	0.0070	0.1478
5		$-F^3 + F^2 + F$	3	0.0320	0.0789
6		$2F^2 - F^3$	3	0.0789	0.0320
7		F	3	0.0664	0.0664
8		F^3	3	0.1478	0.0070
9		$-F^4 + 4F^3 - 6F^2 + 4F$	4	0.0027	0.1740
10		$F^4 - 2F^3 + 2F$	4	0.0112	0.1216
11		$F^4 - 3F^3 + 2F^2 + F$	4	0.0237	0.0872
12		$2F - F^2$	4	0.0195	0.1134
13		$-F^4 + F^3 + F$	4	0.0402	0.0707
14		$F^4 - 3F^3 + 3F^2$	4	0.0707	0.0402
15		$-F^4 + F^3 + F^2$	4	0.0872	0.0237
16	,	$-F^3 + F^2 + F$	4	0.0320	0.0789
17		$F^3 - 2F^2 + 2F$	4	0.0539	0.1009
18	,	$2F^2 - F^3$	4	0.0789	0.0320
19		$3F^2 - 2F^3$	4	0.0445	0.0445
20		$2F^3 - F^4$	4	0.1216	0.0112
21		F^2	4	0.1134	0.0195
22		$F^3 - F^2 + F$	4	0.1009	0.0539
23		F^4	4	0.1740	0.0027
24		$F^5 - 5F^4 + 10F^3 - 10F^2 + 5F$	5	0.0011	0.1946
25		$-F^5 + 3F^4 - 2F^3 - 2F^2 + 3F$	5	0.0043	0.1534
26		$-F^5 + 4F^4 - 5F^3 + F^2 + 2F$	5	0.0086	0.1272
27		$F^3 - 3F^2 + 3F$	5	0.0070	0.1478
28		$F^5 - 2F^4 + F^3 - F^2 + 2F$	5	0.0139	0.1160
29		$-F^5 + 4F^4 - 6F^3 + 3F^2 + F$	5	0.0211	0.0928
30		$F^5 - 2F^4 + F^2 + F$	5	0.0264	0.0816
31	,	$F^4 - 2F^3 + 2F$	5	0.0112	0.1216
32		$-F^4 + 3F^3 - 4F^2 + 3F$	5	0.0152	0.1396
33		$2F^4 - 5F^3 + 3F^2 + F$	5	0.0155	0.0954
34		$F^5 - 3F^4 + 2F^3 + F$	5	0.0346	0.0733
35		$-F^4 + 2F^3 - 2F^2 + 2F$	5	0.0277	0.1051

36		$-F^5+F^4+F$	5	0.0458	0.0680
37		$2F^2-F^4$	5	0.0527	0.0362
38		$-F^5+4F^4-6F^3+4F^2$	5	0.0680	0.0458
39		$F^5-2F^4+2F^2$	5	0.0733	0.0346
40		$F^4-4F^3+4F^2$	5	0.0362	0.0527
41		$F^5-3F^4+2F^3+F^2$	5	0.0816	0.0264
42		$-F^5+F^4+F^2$	5	0.0928	0.0211
43	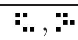	$-F^4+F^3+F$	5	0.0402	0.0707
44	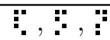	$2F-F^2$	5	0.0195	0.1134
45		$-F^4+4F^3-5F^2+3F$	5	0.0497	0.1271
46	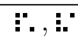	$F^4-2F^3+F^2+F$	5	0.0582	0.0747
47	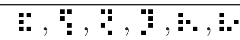	$F^4-3F^3+2F^2+F$	5	0.0237	0.0872
48		$F^4-3F^3+3F^2$	5	0.0707	0.0402
49		$3F^2-2F^3$	5	0.0445	0.0445
50		$2F^4-6F^3+5F^2$	5	0.0280	0.0610
51		$-F^3+F^2+F$	5	0.0320	0.0789
52		$-2F^4+2F^3+F^2$	5	0.0610	0.0280
53		$F^5-3F^4+3F^3$	5	0.1160	0.0139
54		$-F^5+F^4+F^3$	5	0.1272	0.0086
55		$F^4-2F^3+2F^2$	5	0.1051	0.0277
56		$2F^3-F^4$	5	0.1216	0.0112
57		$3F^3-2F^4$	5	0.0954	0.0155
58		$-F^4+2F^3-F^2+F$	5	0.0747	0.0582
59		F	5	0.0664	0.0664
60		$2F^3-3F^2+2F$	5	0.0884	0.0884
61		F^2	5	0.1134	0.0195
62	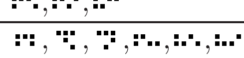	$-F^4+F^3+F^2$	5	0.0872	0.0237
63	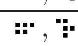	$2F^2-F^3$	5	0.0789	0.0320
64		$2F^4-F^5$	5	0.1534	0.0043
65		F^3	5	0.1478	0.0070
66		$F^4-F^3+F^2$	5	0.1396	0.0152
67		F^4-F^2+F	5	0.1271	0.0497
68		F^5	5	0.1946	0.0011