

Organizational Structure as a Determinant of Performance:
Evidence From Mutual Funds¹

Felipe A. Csaszar

The Wharton School, University of Pennsylvania

2000 Steinberg Hall-Dietrich Hall

Philadelphia, PA 19104

Tel: (215) 746 3112 - Fax: (215) 898 0401

Email: fcsaszar@wharton.upenn.edu

December 15, 2008

¹I would like to give special thanks to Dan Levinthal, Nicolaj Siggelkow, Jitendra Singh, and Sid Winter, for their insights throughout this project. I also thank Nick Argyres, Dirk Martignoni, and the participants at the 15th CCC Doctoral Colloquium and the Management PhD Brown Bag at Wharton for helpful comments. Errors remain the author's own.

Abstract

This paper develops and tests a model of how organizational structure influences organizational performance. Organizational structure, conceptualized as the decision-making structure among a group of individuals, is shown to affect the number of initiatives pursued by organizations, and the omission and commission errors (Type I and II errors, respectively) made by organizations. The empirical setting are over 150,000 stock-picking decisions made by 609 mutual funds. Mutual funds offer an ideal and rare setting to test the theory, as detailed records exist on the projects they face, the decisions they make, and the outcomes of these decisions. The independent variable of the study, organizational structure, is coded from fund management descriptions made by Morningstar, and the estimates of the omission and commission errors are computed by a novel technique that uses bootstrapping to create measures which are comparable across funds. The findings suggest that organizational structure has relevant and predictable effects on a wide range of organizations. Applications include designing organizations that compensate for individual's biases, and that achieve a given mix of exploration and exploitation.

Keywords: Organization Design, Exploration/Exploitation, Decision Making.

1 Introduction

There is a long standing concern that the strategy literature needs a better understanding of how organizational structure and decision-making affect organizational performance. This concern goes back at least to Cyert and March (1963:21), who used the following questions in motivating their theoretical enterprise: “What happens to information as it is processed through the organization? What predictable screening biases are there in an organization? [...] How do hierarchical groups make decisions?” But with a few exceptions, questions of this sort remain mostly unexplored in the strategy literature (Rumelt et al., 1994:42). This lack of knowledge regarding how decision-making structure affects organizational performance continually resurfaces in different areas of management—for example, in the context of ambidextrous organizations, Raisch and Birkinshaw (2008:380) note that “far less research has traditionally been devoted to *how* organizations achieve organizational ambidexterity,” and in the context of R&D organization, Argyres and Silverman (2004:929) show surprise “that so little research has addressed the issue of how internal R&D organization affects the directions and impact of technological innovation by multidivisional firms.” These observations are congruent with the view that organization design—the field specifically devoted to studying the linkages between environment, organizational structure, and organizational outcomes—despite its long history, is in many respects an emerging field (Daft and Lewin, 1993; Zenger and Hesterly, 1997; Foss, 2003).

This paper addresses this gap by developing and testing a model of how the structure of a decision-making organization affects organization-level outcomes. The model is built on two simple ideas from statistical decision theory present in every decision-making organization: that there are two types of errors (i.e., omission and commission errors, or Type I and II errors respectively), and that the way in which individual decisions are aggregated has implications on the overall magnitude of these two errors.

Omission and commission errors are natural measures of performance, as they directly impact performance in many organizations. For example, the profitability of a movie studio depends both on minimizing the number of acquired scripts that turn into box office flops (commission error), and on minimizing the number of *unacquired* scripts that turn into box office hits (omission error). In general, any organization whose task can be broadly defined as making decisions (e.g., top

management teams, boards of directors, venture capital firms, R&D teams, hiring committees), may err in two distinct ways: missing a good choice (omission error), or pursuing a bad choice (commission error). Because of its broad applicability, the interest in omission and commission errors in the management literature is long standing (interestingly, they are the subject of an article in the inaugural issue of the *Academy of Management Journal*, Schmidt, 1958), and a number of papers have used them to measure performance (see references in the next section). However, omission and commission errors have not become as widespread as sales- or profit-based measures, in part because they are considerably harder to observe than more aggregate measures.¹

Omission and commission errors are not only useful performance measures, but they also provide an opportunity to explore the implications of different ways of aggregating decisions. For instance, if two individuals (e.g., partners in a venture capital firm) had to agree on the quality of a project before investing in it, the probability that the project gets funded is lower than if the two individuals could approve it independently of one another. Sah and Stiglitz (1986, 1988) call these two ways of organizing a hierarchy and a polyarchy, respectively.² 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), at the expense of having a higher chance of investing in a bad project. Sah and Stiglitz mathematically formalized this intuition, showing that the hierarchy minimizes commission errors, while the polyarchy minimizes omission errors. Their work implies that these two alternate structures allow an organization to tradeoff one error for the other, hence which of the structures is ‘better’ is context-dependent (it depends on the relative cost of the two errors).

Sah and Stiglitz used their model to address the contrast between centralized planned economies and free markets, but their approach has a much broader applicability—in particular, it speaks to the contrast between centralized and decentralized organizations. The applicability of their model to this context stems from the fact that “decision makers generally base their actions on estimates formulated at other points in the organization” (Cyert and March, 1963:85), and that in centralized organizations these estimates must “flow up” through more decision-makers before reaching the final

¹Another reason for their relatively low use is that their connection to the “bottom line” is less direct—omission and commission errors are an antecedent of performance, but require to be complemented with other information (such as the cost of each error) for them to have a direct “bottom line” impact.

²Sah and Stiglitz’s use of the word ‘hierarchy’ is non-standard, as none of the two evaluators is a superior of the other. It should not be confused with the traditional meaning of the word in the organizations literature.

decision-maker than in decentralized organizations (Robbins, 1990:6). Thus, the information flow in centralized organizations resembles that of hierarchies, while the information flow in decentralized organizations resembles that of polyarchies. In sum, the Sah and Stiglitz framework captures the fact that information must pass more filters in centralized than in decentralized organizations. Interestingly, the predictions of Sah and Stiglitz have never been empirically tested.

This paper develops a model using the Sah and Stiglitz framework, and then tests its predictions on a large sample of mutual funds, as these firms represent the quintessential decision-making organization trying to detect opportunities in an uncertain environment (Kirzner, 1973; Amit and Schoemaker, 1993). From an empirical viewpoint, mutual funds are an ideal and rare setting in which to test these ideas, as funds are heavily scrutinized, very detailed records exist on the projects they face (possible investments), the decisions they make or do not make (buying or not buying each of these possible investments), and the outcomes of these decisions (the ex-post return of having bought or missed a given investment). Moreover, the organizational structure of mutual funds exhibit substantial variation and can be coded from fund management descriptions made by Morningstar.

This paper adds to the literature in several respects. First, it depicts a process that links organizational structure to organizational-level outcomes, which has implications for a broad range of decision-making organizations. Second, it is the first empirical examination of the predictions of Sah and Stiglitz regarding centralized and decentralized organizational structures. Third, by describing a pervasive mechanism by which individual decision-making aggregates into organizational level outcomes, this paper provides an answer to the long standing question of “do organizations have predictable biases?” (Cyert and March, 1963:21; Rumelt et al., 1994:42), and responds to calls for exploring how the behavioral aspects of decision-making affect strategic outcomes (Zajac and Bazerman, 1991) and how micro decisions turn into macro behaviors (Coleman, 1990:28). Finally, the results provide a basis to develop prescriptive guidelines for organization design.

2 Theoretical Motivation

What are the effects of organizational structure on organizational performance is among the fundamental questions of the strategy field (Rumelt et al., 1994:42) and organization theory (Thompson,

1967), hence it is no surprise that this question has been extensively attacked from several perspectives since old—even biblical (Van Fleet and Bedeian, 1977:357)—times. Thus, instead of attempting the impossible task of summarizing these literatures, this section attempts to present a broad overview with an emphasis on highlighting the differences between the current and previous approaches.

A first distinction when dealing with structure is what is the level of analysis. Broadly speaking, the basic unit of analysis used in the organizational structure literature has either been individuals (e.g., Cyert and March, 1963) or business-divisions (e.g., Chandler, 1962). This paper deals with the former type of structures. Under this view, organizational structure is “the pattern of communications and relations among a group of human beings, including the processes for making and implementing decisions.” (Simon, 1947/1997:18–19).

The modern interest in organizational structure as a pattern of communications among individuals can be traced back to Graicunas’ paper on the use of graphs to understand span of control, published as a chapter on Gulick and Urwick (1937). Simon’s (1947/1997) more elaborate view of organizations as information-processing devices composed of boundedly rational individuals, led him to make span of control contingent on contextual factors, and later to extend the work of Bavelas (1950) and Leavitt (1951) to determine how effective were different information processing structures at completing organization-level goals (Guetzkow and Simon, 1955). Subsequently, the role of organizational structure took a central place in the *Behavioral Theory of the Firm* (Cyert and March, 1963). However, with one exception (Cohen et al., 1972), the Carnegie tradition devoted most of its energies to decision-making in the absence of organizational structure concerns. In fact, in a recent article, Gavetti, Levinthal, and Ocasio (2007) call organizational structure a “forgotten pillar” of this tradition.

Contingency Theory (e.g., Burns and Stalker, 1961; Woodward, 1965; Lawrence and Lorsch, 1967; Thompson, 1967) shared with the Carnegie tradition a sensibility (borrowed from cybernetics and systems theory) that highlighted the role of information-processing constraints. Contingency Theory extended that sensibility by delving into the linkages between the environment and the organization, and seeking to find the patterns of organizational structure—such as formalization and administrative intensity—that are typically associated, or have the best ‘fit,’ with contextual factors—such as size and technological uncertainty. Most of this literature has not dealt with

individuals as level of analysis nor with establishing the processes that connect context to structure (Meyer et al., 1993).

Team Theory (Marschak and Radner, 1972) took a formal and information-theoretic approach to organizations, by mathematically modeling the effects of information decentralization (i.e., not all team members have access to the same information) under perfect alignment of incentives. Interestingly, the role of structure is almost absent in the initial version of the theory. More recently, Radner (1992) and Van Zandt (1999) extended the theory to account for process decentralization (i.e., different members perform different tasks) in hierarchical organizations (i.e., tree-like graphs). These models, which are almost solely focused on efficiency measures, analyze the number of operations it takes an organization to perform a given task.

Sah and Stiglitz (1986, 1988) contributed to the team-theoretic approach by introducing two new elements into it: modeling communication patterns as sequential or parallel, and measuring performance as omission and commission errors. They used this approach to mathematically analyze organizations with two members (1986) and committees (1988). An appealing characteristic of their approach is that it creates bridges between organization design and vast and distant literatures: parallel and sequential structures have been well studied in fields as disparate as reliability theory (Rausand and Høyland, 2004), circuit design (Moore and Shannon, 1956/1993; von Neumann, 1956), and machine learning (Hansen and Salamon, 1990); and omission and commission errors have been well studied in statistical decision theory (Berger, 1985), diagnostic testing (Hanley and McNeil, 1982), and signal detection theory (Green and Swets, 1966).

The literature that has built on the work of Sah and Stiglitz has focused mostly on analytical models of voting and committee decision-making. Few of the references to their work have come from the management domain; among the exceptions is work discussing M&As (Puranam et al., 2006), venture capital syndication (Lerner, 1994), technological choices (Garud, Nayyar, and Shapira, 1997), and the implications of alternative evaluation on search behavior (Knudsen and Levinthal, 2007). Interestingly, perhaps because of the empirical difficulties associated with collecting information on organizational structure and errors (particularly omissions), the predictions of Sah and Stiglitz have never been empirically tested.

During the 80s and 90s—a period dominated by content rather than process approaches to strategy research (Rumelt et al., 1994:545)—questions of structure became less central to the strat-

egy field. More recently, this tendency has began to reverse, with several researchers issuing calls to better understand the strategy process (see Zajac, 1992; Chakravarthy and White, 2002, and references therein), a topic that naturally leads to questions of organizational structure. Some examples of this renewed interest in organizational structure is the work exploring how problem-decomposition relates to organization-decomposition (Marengo et al., 2000; Ethiraj and Levinthal, 2004), how the search behavior of employees affects organization-level search (Rivkin and Siggelkow, 2003), how the network connectedness of the members of an organization determines the organization's innovative output (Lazer and Friedman, 2007), and how the location of R&D units within an organization (i.e., headquarters- or subsidiary-level) affects the type of innovations produced by the organization (Argyres and Silverman, 2004). With the exception of Argyres and Silverman (2004), these papers have used simulation as a research methodology.

Another literature that has contributed to the understanding of the interplay between structure and performance is the work by Bower (1970) on the resource allocation process, which has gained further development and attention with the development efforts of Burgelman, Christensen, Doz, Gilbert and others (for references see Bower and Gilbert, 2005). This line of research has described the complex and subtle processes whereby projects are identified, proposed, refined, and approved in large corporations.

Almost without connection to the previous literatures, a rich body of research on group decision-making rooted in psychology was developed. The theoretical work in this literature (e.g., Davis, 1973; Kerr et al., 1996) is remarkably similar to the work of Sah and Stiglitz, as it presents mathematical models of decision-schemes that predict group-level outcomes. This literature has produced empirical results (Stoner, 1961; Stasser and Titus, 1985; Hinsz et al., 1997), but because it has been primarily conducted in the laboratory, using small groups that meet for brief times, its results may not be generalizable to more complex, on-going organizations (Argote and Greve, 2007:344). Moreover, one of the most important issues to strategy researchers that is analyzed by the group decision-making literature, the issue of whether groups take more or less risks than its members, remains an open question (Connolly and Ordóñez, 2003:510).

Although the previous literatures have provided many important insights on what is the impact of structure on performance, the field of organizations lacks an empirically validated theory that starting from structure at the level of individuals is able to predict organization-level measures of

performance relevant to firm strategy. Generally, the previously reviewed literatures do not provide such a theory because of at least one of the three following reasons: not describing structure at the individual level of analysis, not predicting measures of performance useful to strategy research, or not having empirical support. While clearly limited, this paper empirically explores a theoretical development that meets these three criteria.

3 Model

This section describes a simple mathematical model which is used to rigorously derive all the hypotheses tested in this paper. The aim is to present a synthesis of results relevant for organization design, selected from a loose collection of models on fallible decision making. The model describes an organization which receives projects of various qualities, facing the task of screening them, i.e., to select those projects that surpass a given quality threshold or benchmark. This characterization is consistent with viewing the environment as a flow of opportunities (Kirzner, 1973; Shane, 2000), and the organization as deciding based upon uncertain information about these opportunities (Amit and Schoemaker, 1993).

An organization is represented by the number of individuals (N) it has, and by the decision making rule it uses to arrive to an organization-level decision. For simplicity, these rules are coded as one number (C) that denotes the minimum consensus level required to approve a project. Hence, for example an organization with five members which approves projects based on the majority rule is represented by $N = 5$ and $C = 3$, or simply $5/3$; likewise, a $2/2$ is a two member organization that only approves projects for which there is consensus; a $3/1$ is a three member organization that approves a project when any member decides to approve it. An organization of a single individual is denoted $1/1$.

The projects faced by the organization are assumed to have a true value which is imperfectly perceived, as a signal plus noise, by the members of the organization. For simplicity, all the members of the organization have the same ability to screen projects, i.e., their screening generates independent draws from the same noise distribution. Finally, the model assumes that each of the two types of errors the organization can make (missing a good project or approving a bad one, or omission and commission errors respectively) has a given cost (c_I and c_{II} , respectively).

As all models, this stylized description of organizations leaves outside of its scope many phenomena such as organizations whose task is different from screening projects, heterogeneity in ability, group dynamics such as herding (Bikhchandani et al., 1992) or groupthink (Janis, 1972), and more generally, organizational structures different from those describable in terms of N and C . Nonetheless, the model allows us to focus on some basic mechanisms which are pervasive to organizations: how centralized or decentralized is the decision process of an organization, and how many individuals are involved in it. Some examples can illustrate how the model captures these organizational characteristics.

For instance, a 3/3 could represent the decision making process occurring inside a venture capital firm in which the three partners must agree to invest in a firm, or a three-level hierarchy in which projects received by a low-level employee must escalate up to the CEO in order to be approved. In both examples, three out of three individuals must concur about the goodness of the project for it to be approved by the organization. On the other hand, a 3/1 could represent the following decentralized structures: a firm with three research engineers, anyone of whom may independently decide to pursue further research on a new technology; or a mutual fund with three autonomous fund managers, anyone of whom may authorize the purchase of a security. In these last two examples, it suffices that one out of the three individuals likes the project, for the project to be approved.

Mathematically, the model is described as follows. An individual approves a project if her perception of its quality is above a benchmark b , hence the probability that an individual approves a project of a given quality q is $p(q) = \Pr\{q + \tilde{n} > b\}$, where \tilde{n} is a random draw from a noise distribution. An organization with N members and consensus level C will approve the project if at least C of its members approve it, which happens with probability

$$P(q; N, C) = \sum_{i=C}^N \binom{N}{i} p(q)^i (1 - p(q))^{N-i}.$$

Based on this formula, several organization-level metrics can be computed.

The probability that an organization will accept a project of an unknown quality is the marginal

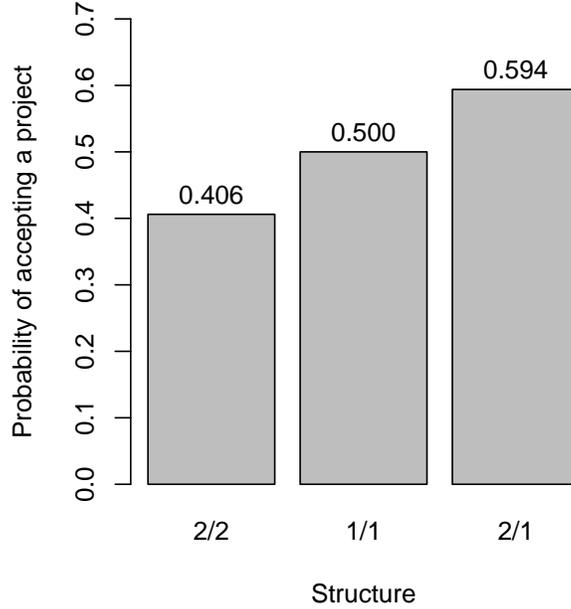


Figure 1: Probability of accepting a project for three organizational structures.

distribution of $P(q; N, C)$ with respect to q (with pdf $f_{\tilde{q}}(q)$),

$$P_A(N, C) = \int_{-\infty}^{\infty} f_{\tilde{q}}(q) P(q; N, C) dq. \quad (1)$$

Figure 1 shows the expected probability of accepting a project by three organizational structures (a centralized structure, 2/2; an individual manager, 1/1; and a decentralized structure, 2/1), assuming $\tilde{q} \sim U[-3, 3]$, $\tilde{n} \sim N(0, 1)$, and $b = 0$. Note that the centralized firm is the one that approves the least projects, the decentralized firm is the one that approves the most, and the individual manager lies in between the other two structures. Even if the actual values shown in the figure depend on the probability distributions used, their relative ordering is always the same.³

The probability that a given N/C organization will miss a good project (a Type I or omission error) is the probability of rejecting $(1 - P(q; N, C))$ a good project ($q > b$) weighted by the

³To show that the ordering does not change, imagine how two individuals, A and B, would operate under centralized (2/2) and decentralized (2/1) organizations. In the centralized organization both A AND B must agree, and hence the final probability of acceptance is the conjunction of both approval probabilities ($p_{2/2} = p^2$). In the decentralized organization either A OR B must accept the project, and hence the final probability of acceptance is the disjunction of the individuals' probabilities ($p_{2/1} = p + p - p^2$). The acceptance probability of the individual manager is simply $p_{1/1} = p$. It is easy to show that for any possible value of p (i.e., from 0 to 1), $p_{2/2} \leq p_{1/1} \leq p_{2/1}$.

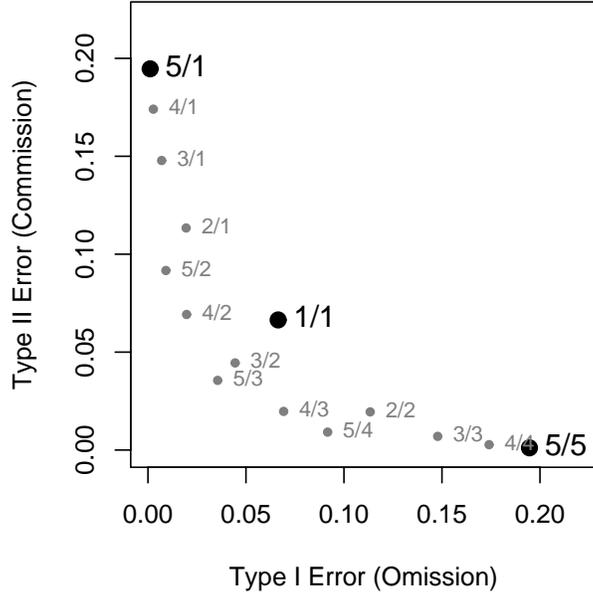


Figure 2: Type I versus Type II error for all N/C -structures with up to five members.

probability of receiving a project of that quality ($f_{\bar{q}}(q)$), i.e.,

$$P_I(N, C) = \int_b^{\infty} f_{\bar{q}}(q)(1 - P(q; N, C))dq. \quad (2)$$

Similarly, the probability of accepting a bad project (a Type II or commission error) is

$$P_{II}(N, C) = \int_{-\infty}^b f_{\bar{q}}(q)P(q; N, C)dq. \quad (3)$$

How organizational structure affects the types of errors made by the organization becomes more evident when plotted. Figure 2 plots all the organizations with up to five individuals (1/1, 2/1, 2/2, 3/1, 3/2, ...) according to their Type I and Type II errors, under the same probability distribution assumptions used for the previous figure. As before, the exact positions of the organizations vary depending on the probability distributions used, but not their relative ordering along both dimensions.

Figure 2 illustrates several results of the model: (a) centralized structures minimize the commission error (e.g., structure 5/5 appears at the bottom right); (b) decentralized structures minimize the omission error (e.g., structure 5/1 appears at the top left); (c) for a fixed organization size,

intermediate structures (e.g., $5/4$, $5/3$, $5/2$) offer tradeoffs between the two extremes; and (d) larger organizations, allow decreasing both errors at the same time (see for example in Figure 2 how $1/1$, $3/2$, and $5/3$, are successively better along both axes).

Note that, a priori, no organization is better than any other: the ‘right’ organization depends on the task the organization must perform (i.e., the costs of the errors, c_I and c_{II}), the cost of the organization (e.g., the number of decision-makers), the characteristics of the individuals (their noise distribution, $f_{\tilde{n}}$), and characteristics of the environment (the probability distribution of the projects’ quality, $f_{\tilde{q}}$). For example, if the cost of accepting a bad project is very high (this may be the case of a high reliability organization), it pays off to choose a structure close to the bottom-right of Figure 2; on the other hand, if the cost of missing a good project is high (this could be the case of an R&D lab in a highly competitive industry), it pays off to choose a structure on the top-left of the figure; and if both errors are equally relevant (this could be the case of a group of investors, for which both not investing in a good asset is as costly as investing in a bad one), the best is to minimize both errors jointly. This conditional view of organization design is consistent with the concept of ‘fit’ which pervades structural contingency theories (Donaldson, 2001; Siggelkow, 2002).

Even though the logic presented so far is not new, it has only recently received attention from management scholars. To the best of my knowledge, currently the only management papers that have tried to elaborate the broader organizational implications of this logic are Garud, Nayyar, and Shapira (1997), Christensen and Knudsen (2002, 2007), Knudsen and Levinthal (2007), and Csaszar (2007), all of which have used a theoretical approach. The present paper attempts to empirically test hypotheses derived from the previous model, using a large dataset of mutual fund investment decisions. Empirically validating the model is important, as it describes in a stylized manner centralization and decentralization, two basic properties of organization structure.

4 Hypotheses

The independent variable of the study is organizational structure, which from the dataset used can reliably be coded into three non-overlapping categories: organizations managed by one individual, decentralized organizations, and centralized organizations. These categories, in terms of the previous model, correspond to structures $1/1$, $N/1$, and N/N , respectively (where N represents any

integer greater than 1). Because of data limitations later discussed, the labelling scheme does not need to account for organizations with intermediate levels of consensus (i.e., N/C with $1 < C < N$ such as $3/2$ or $7/6$).

The dependent variables of the study are three outcomes predicted by the model: number of approved projects, omission errors, and commission errors. Because the model predicts that these three outcomes are most different for centralized versus decentralized structures, the hypotheses are stated as comparisons between these two structures. It would be also possible to write down six other hypotheses comparing individual managers to centralized structures, and individual managers to decentralized structures, but to avoid a litany of hypotheses, these comparisons are discussed in the results section without being formally enumerated here.

The first hypothesis asserts that the number of projects accepted behaves as predicted by Equation 1.

Hypothesis 1 *Decentralized organizations accept more projects than centralized organizations.*

Similarly, the second hypothesis purports that omission errors behave as Equation 2 predicts.

Hypothesis 2 *Decentralized organizations make fewer omission errors than centralized organizations.*

Finally, the third hypothesis states that commission errors behave as predicted by Equation 3.

Hypothesis 3 *Decentralized organizations make more commission errors than centralized organizations.*

5 Empirical setting and approach

Before delving into the specifics of the dataset and the statistical methods, it is important to understand the structure of the empirical problem. To test the previous hypotheses, all of the following must be observed: (i) organizations making decisions about projects, (ii) a measure of the quality of each project decided upon, (iii) the decision that each organization made with respect to every project it faced, and (iv) the organizational structure of each organization. Point (i) exists in many settings (e.g., firms deciding who to hire, where to expand, what to sell, etc.). Point (ii)

is also readily available in settings where the ex-post value of the project is visible and can proxy for the true quality of the project (e.g., in the venture capital context it could be a function of the IPO value of a startup a VC considered investing in, or in the R&D context it could be the number of citations accrued by a patent after a firm had the opportunity to buy it).

Points (iii) and (iv) from the previous list pose serious hurdles to the empirical researcher. First, typically there is no track record of the projects an organization considered but decided not to pursue (e.g., all the firms a venture capitalist screened but did not invest in). Secondly, organizational structure is not tracked in public databases. What may be available to some extent are organizational charts, but these do not tell how centralized or decentralized a given decision making process is (e.g., by looking at an organizational chart it is not possible to know the decision process used to set the direction of R&D, perform M&As, or decide on IT investments).

Mutual funds offer a rare window into the implications of organization design on organizational performance, as in this setting the four necessary ingredients previously mentioned are observable: (i) managing a mutual fund is essentially about making decisions (i.e., deciding what to buy); (ii) the ex-post return of each investment is an adequate measure of the quality of each decision; (iii) by regulation, funds must disclose their holdings periodically, allowing the researcher to discern the projects the fund accepted (e.g., stocks that were bought) from the projects the fund rejected (e.g., stocks that were not bought); and (iv) organizational structure is observable from descriptions of the fund management prepared by Morningstar. Additionally, there are thousands of mutual funds, and the typical fund makes dozens of decisions per quarter. All these considerations make mutual funds an exceptional vehicle to study the effects of organization design on organizational performance—mutual funds would make a good aspirant for the “fruit fly” of organization design.

Despite the virtues of mutual funds as an empirical setting, there is a strong tradition in the finance literature that would predict that organizational structure should not matter as a determinant of fund performance. In a nutshell, the efficient market hypothesis (EMH) (Fama, 1970) purports that all available information is already reflected in asset prices, and hence future returns are unpredictable. If that is true, organizational structure should not predict mutual fund performance. However, the EMH no longer holds the invulnerable position it once did, as in the last fifteen years a vast literature on market anomalies has emerged. See Malkiel (2003) and Barberis and Thaler (2003) for arguments and references coming from the two opposing camps.

Several anomalies have been reported in the context of mutual funds. For example, Grinblatt and Titman (1992) and Goetzmann and Ibbotson (1994) showed that differences in performance between funds persist over time, Chevalier and Ellison (1999) found that managers who attended higher-SAT undergraduate institutions have systematically higher risk-adjusted excess returns, Makadok and Walker (2000) identified a forecasting ability in the money fund industry, and Cohen et al. (2007) presented evidence that fund managers place larger and more profitable investments on firms they are connected to through their social network.

As the variance explained by market anomalies is small (e.g., the typical R^2 of an anomaly is below 1%), even if the EMH is not true, from a pragmatic point of view, a large portion of asset returns *is* random. For the effects of this paper, this implies that the variance explained by organizational structure—if any—is not expected to be large. It also implies, that if the model has some explanatory power, this is likely to increase in settings where the link between cause and effect is more deterministic. Given that stock picking is possibly one of the most random task environments, then mutual funds can be seen as a stringent testing arena, and the results of this paper as conservative estimates.

5.1 Independent variable: mutual fund organizational structure

A mutual fund is a type of investment that pools money from many investors to buy a portfolio of different securities such as stocks, bonds, money market instruments, or other securities. US mutual funds are regulated by the Securities and Exchange Commission (SEC), which among other requirements, forces funds to report their portfolio holdings at the end of the last trading day of every quarter (Form 13F), and also to periodically report who their fund managers are (Form 487). Mutual funds are heavily scrutinized not only by the SEC, but by institutional investors and investment research firms.

Morningstar, one of the leading investment research firms, offers information about mutual funds to investors and financial advisors. By using public sources and periodically meeting fund managers, Morningstar’s analysts produce a one-page report, densely packed with statistics and analysis, for each fund they track. For the present study, what is important about these profiles is that they contain a section called “Governance and Management” that presents a short biography of the managers and describes how they manage the portfolio. This section of the report contains

Structure (N/C)	Excerpts from Morningstar’s mutual fund description
1/1	“Ron Baron has been at the helm since the fund’s inception ... He’s the driving force behind this portfolio ... buys companies he thinks can ...” (BPTRX)
2/1	“Managers Scott Glasser and Peter Hable each run 50% of the portfolio ...” (CSGWX)
3/1	“Three management firms select 10 stocks apiece for this fund’s portfolio.” (SFVAX)
5/1	“[the fund] divvies up assets among five subadvisors, and each picks eight to 15 stocks according to his own investing style.” (MSSFX)
8/1	“The fund used to divide the assets among five different subadvisors, but it added another three ... Each subadvisor has a separate sleeve that it manages in a particular style ...” (AVPAX)
2/2	“Teresa McRoberts and Patrick Kelly became comanagers of this fund in late September 2004 ... They don’t pay too much attention to traditional valuation metrics such as ...” (ACAAX)
7/7	“All investment decisions are vetted by the entire seven-person team ... Management populates the fund with 30–50 stocks ...” (CBMDX)

Table 1: Examples of how organizational structure is coded from Morningstar’s fund descriptions. The ticker symbol of each fund appears in parenthesis.

enough information to code organizational structure as modeled in this paper (in terms of number of managers, N ; and level of consensus required, C). To understand how the coding was done, consider the excerpts shown in Table 1, which illustrate typical descriptions. To increase consistency, coding was done using the following four rules:

1. If the description mentions managers’ names, then N is set to the number of people mentioned as manager or co-manager, with the exception of people that is described in an explicit secondary role (for example, if one manager is described as subordinate, performing administrative tasks, not participating in the day-to-day management, or recently ascended but retaining his/her analyst tasks).
2. If the description is explicit about the number of ‘sleeves,’ subadvisors, or describes how managers split their portfolios, then N is set to the number of divisions of the portfolio, and C to 1, as this is a decentralized fund.
3. If two or more managers are mentioned, but nothing is said about how they coordinate (e.g., they are addressed as a plurality, as in “they invest in ...”) it is assumed that the fund uses consensus ($N = C$). This is reasonable, as this is the default structure of co-managed

funds, and because if managers work separately, they do not have incentives to be reported as working in tandem (managers want to create their own reputations).

4. If no specific manager names are mentioned (e.g., the description only talks about a generic “the management”), or if the description says that the fund is run by an algorithm (e.g., some funds that track indices operate like this), then the fund is left unclassified.

Less than 4% of the funds fell in the unclassified bucket, and less than 1% of the funds had a consensus level different from 1 or N . These two classes of funds were eliminated from the dataset.

Because fund descriptions do not include nuances such as the relative sizes of each sleeve of a decentralized fund, the organizational structure of the subadvisor of each sleeve, or the share of power of each manager in a centralized fund, the funds were aggregated into three broader categories: ‘1/1’ (managed by an individual), ‘ $N/1$ ’ (decentralized), ‘ N/N ’ (centralized). This decision insures against over interpreting the results.

All the funds were coded both by the author and by one research assistant. The percentage of agreement between both categorizations was 96%. The results here presented use the author’s categorization, but all the results are robust to using the other categorization as well.

5.2 Dependent variables: omission and commission errors

The main intuition behind the measures of omission and commission error here developed is the following: In hindsight, a commission error occurred whenever a fund bought an asset that turned out to have a poor performance, i.e., whose ex-post return fell below a given benchmark; similarly, an omission error occurred whenever a fund failed to buy an asset which turned out to have a good performance.⁴ To observe these errors, two types of data are required: the list of assets that a fund did and did not buy, and the returns of these assets. Good data sources exist for both elements.

In order to make the discussion more precise, some notation is useful. For a given mutual fund F at time t , let $A = \{a_1, a_2, \dots, a_n\}$ be the set of assets that F bought during time period t (subscript

⁴Omission and commission errors can also be measured using sell (instead of buy) decisions (i.e., not selling a stock that tumbles, and selling a stock that rises in price). Informal interviews with fund managers pointed to the fact that while the process of buying a stock is quite deliberative, the sale of stocks is a semi-automatic process, guided by stop-loss orders, and tax and liquidity considerations. In fact, the coefficients for organizational structure cease to be significant if the regressions on §6.2 are re-run using errors on sales as dependent variable. Further research may explore if this represents a bias toward over-studying buy decisions, maybe at the expense of sell-decisions.

t is omitted for convenience). As the best available information on mutual fund holdings is quarterly reported, from here on a unit of time is one quarter. Let $U = \{u_1, u_2, \dots, u_N\}$ represent the assets in which F can invest, or F 's investment universe at time t . The number of assets bought by F at period t is n , and the number of assets in its investment universe at time t is N . By definition, the assets bought by a fund are a subset of the fund's investment universe, $A \subseteq U$.

Asset returns are computed by comparing the end-of-period prices, i.e., $r(a)$ represents the return of asset a from the end of period t to the end of period $t + 1$. The study uses a per fund benchmark, defined as the average return of the assets in the fund's investment universe at time t , i.e., $b = \frac{1}{N} \sum_{i=1}^N r(u_i)$. An asset is catalogued as 'good' if its return in a given period is equal or above the benchmark b . The subset of good assets that the fund bought during period t is denoted $A^+ = \{a | a \in A \wedge r(a) \geq b\}$, and its cardinality is denoted n^+ . Similarly, the bad assets bought are $A^- = \{a | a \in A \wedge r(a) < b\}$ with cardinality n^- .

At first sight, several measures may capture the commission error of a fund. Two examples could be the number of bad assets bought (n^-), and the total negative return ($\text{TNR} = - \sum_{\{a \in A^-\}} r(a)$; the initial minus sign makes the measure increase in the right direction). But an important problem with these metrics is that as different funds invest in a different number of assets and in different investment universes, these metrics are not comparable among funds, and are thus unsuitable for the purposes of the present study.

One way to address the issue of comparability, is to infer the probability distribution of the errors and to report them in terms of how likely or unlikely was their occurrence; in other words, to report errors as probabilities. If the probability distribution used already accounts for the specifics of the situation (i.e., the assets that were available and the number of assets that were picked), then the measures are comparable across funds.⁵

The hypergeometric distribution serves as a first approach to create a probability-adjusted measure of the errors of a fund. The hypergeometric distribution, whose probability mass function is $f(k; N, m, n) = \binom{m}{k} \binom{N-m}{n-k} / \binom{N}{n}$, is typically illustrated in terms of the probability of getting exactly k red marbles after drawing n marbles (without replacement) from an urn with N marbles

⁵An example may clarify the idea further: imagine you want to compare who is better at games of chance, someone who flipped one thousand coins and got 600 heads, or someone who threw two thousand dices and got 400 ones. By putting a probability distribution on the outcomes ($\Pr\{\text{Head}\} = \frac{1}{2}$ and $\Pr\{\text{One}\} = \frac{1}{6}$), it does not matter that they both played different games, in both cases it is possible to compute a statistic (in this case a chi-squared), and compare the players in terms of how unlikely were their results.

out of which m are red. Thus, replacing ‘marble’ for ‘stock,’ and ‘red’ for ‘bad,’ gives rise to a function that computes the probability of getting a given number of bad stocks, which is already adjusted for portfolio and universe size, and the number of bad stocks in the investment universe.⁶

But a limitation of the hypergeometric approach is that it weighs all bad decisions equally, regardless of the size of the errors (i.e., a stock that slightly underperformed the benchmark gets counted the same as a stock whose price collapsed). To avoid discarding the valuable information contained in the size of the errors, the probability distribution of the errors must be estimated via bootstrapping (Efron, 1979; Efron and Tibshirani, 1993). The bootstrap consists in creating an arbitrarily good approximation of a population by means of Monte Carlo simulations, and using this new population to compute the exact value of a statistic. In this case, the population to be estimated is the set of all the possible portfolios of a given size that can be drawn from a given investment universe.

An example clarifies how the bootstrap can be used to measure commission errors. Suppose the returns of the assets in the investment universe of Fund F are $\{-5\%, -2\%, -1\%, 1\%, 3\%, 4\%\}$, the benchmark is $b = 0$, and of these assets, the fund bought the three assets which ended up returning $\{-2\%, -1\%, 4\%\}$. Hence F ’s total negative return (TNR) is 3% ($= -[-2\% + -1\%]$). To assess how large or small this number is, it has to be compared to the TNRs of the population of funds that can draw three stocks from the same investment universe as F . In this example 20 ($= \binom{6}{3}$) other portfolios could have been bought, but in realistic cases the space of possible portfolios cannot be exhaustively explored,⁷ hence the method relies in randomly sampling the space of possible portfolios. With the exception of some well known pathological cases (Davison and Hinkley, 1997, Sec. 2.6), a statistic computed via bootstrap converges to the real statistic as the number of random draws increases. For the case of the data used in this paper, by making each fund ‘compete’ against

⁶Under this approach, the commission error of Fund F is the cumulative distribution of the hypergeometric evaluated at the number of bad assets that the fund picked ($\sum_{i=0}^{k-1} f(i; N, m, n) + \frac{1}{2}f(k; N, m, n)$). This sum represents the proportion of all the possible portfolios that can be drawn from the fund’s investment universe that contain at most k bad assets. The larger the sum, the larger the error. Conceptually, what this measure does is to determine how well Fund F would stack up in a competition against all the possible funds that could have existed in the same environment. For example, if an investment universe has three bad and three good stocks ($m = 3, N = 6$), the portfolio size is two ($n = 2$), and the fund picked one bad asset ($k = 1$), then it is clear that the fund did an average job, and in fact, the measure says exactly that: $0.5 (= f(0; 6, 3, 2) + \frac{1}{2}f(1; 6, 3, 2) = \binom{3}{0} \binom{6-3}{2-0} / \binom{6}{2} + \frac{1}{2} \binom{3}{1} \binom{6-3}{2-1} / \binom{6}{2} = 0.2 + 0.3 = 0.5)$. This measure is now comparable to that of a fund that could have bought a completely different number of stocks in a different investment universe. A measure of the omission error can be defined similarly.

⁷The average fund in the dataset buys 16 stocks from a universe of 195, which creates a space of $\binom{195}{16} \approx 10^{23}$ possible portfolios.

100,000 simulated portfolios, the standard error introduced by the bootstrap procedure is less than 0.003.

Once the population of comparable portfolios is created, the measure of commission error is simply a measure of how deviant is F 's error with respect to the commission errors of that population. Given the Central Limit Theorem and the large number of simulations, the normal distribution is a very good approximation for the TNRs of the population. Thus, errors are reported in terms of standardized scores, where the higher the score, the higher the error.

The omission error can be defined analogously to the commission error, but instead of measuring TNR, measuring the total unbought positive returns (TUPR). That is, the sum of the good assets that belong to the investment universe of Fund F , but were not bought in the current period. Mathematically, $TUPR = \sum_{\{a \in U \wedge a \notin A\}} r(a)$. Following the previous example, the TUPR of Fund F is 4% (= 1% + 3%). As before, the bootstrap is then used to compute a probability-adjusted measure that is expressed as a standardized score.

Finally, to increase reliability, the omission and commission errors of each fund were averaged using errors computed for ten quarters (from 2004Q4 to 2007Q1). More formally, if $E_{F,t}^I$ is the omission error of Fund F at quarter t , then the dependent variable used was $E_F^I = \frac{1}{10} \sum_{t=1}^{10} E_{F,t}^I$, and similarly for the commission error. The logic behind this, is that by averaging, the systematic information contained in the quarterly errors remains, but part of the noise of the measure is canceled out.

5.3 Data preparation and limitations of the datasets

The content and format of Morningstar's one-page mutual fund reports has changed repeatedly over the years, and only since 2007 it started including a "Governance and Management" section with enough information to code organizational form for a large sample of funds. This implies that the data on organizational structure is only available as a snapshot for December 2007. Because organizational structure is only available for December 2007, while the dependent variables are computed using errors from 2004Q4 to 2007Q1, funds that changed their organizational structure after 2004Q4 but before December 2007 are partially misclassified in the analysis. Fortunately, changes to the organizational structure of funds are rare. There are no official statistics, but a good estimate of change to the organizational structure of mutual funds can be gathered from Morningstar

(2008). Apart from 500 fund reports, Morningstar (2008) also includes a brief description of all the management changes occurred to these funds during 2007 (p. 29). From the 500 reported funds, 32 experienced some sort of management change (the most typical change is the replacement of a manager), and only four funds experienced a change of organizational structure as coded in this paper. This amounts to a 0.8% yearly probability of change in organizational structure.

In December 2007, Morningstar kept organizational descriptions for 1687 funds. To increase comparability, only the funds that were primarily devoted to stocks (and not other asset classes such as bonds or options) were selected. To do so, funds were chosen if its asset composition (according to the CRSP dataset “Mutual Fund Profiles and Monthly Asset Data”) were at least 60% stocks in 12 out of the 16 quarters from 2003Q2 to 2007Q1. This narrowed down the list to 1087 funds. I then used the CRSP datasets “Portfolio Holding Information,” and “Monthly Stocks” to choose only those funds for which CRSP had the returns of the individual stocks owned by the fund for at least 50% of its portfolio value for at least 6 out of the 10 quarters from 2004Q4 to 2007Q1 (these are the periods used to compute the error measures), and at least 3 out of the 6 quarters from 2003Q2 to 2004Q3 (these periods are used to define investment universes, which is addressed later). This narrowed down the list to 642 funds. The drop is primarily explained because CRSP only tracks the returns of the stocks traded in NYSE, NASDAQ, and AMEX, while many funds invest in international stocks, and to a lesser extent because the CRSP portfolio holdings dataset has missing observations. Finally, funds for which the Morningstar description did not allow one to infer an organizational structure were dropped, leaving the final count in 609 funds, which are owned by 154 different parent firms. Collectively, in the ten quarters from 2004Q4 to 2007Q1, these funds invested in 5833 distinct stocks (as identified by their CUSIP number), made 153,457 buy decisions, and had \$1.6 trillion under management at the end of the period. The range of dates used is due to data limitations: before 2003Q2 the CRSP holdings database becomes sparse, and by December 2007, CRSP had not yet uploaded the holdings information for the quarters after 2007Q1.

Which stocks a fund bought during the quarter ending at date t was determined by looking at the stocks added to the portfolio since the last reported quarterly holdings. The quarterly holdings were gathered from the CRSP dataset “Portfolio Holdings Information,” which is itself gathered from the Forms 13F mutual funds report to the SEC. One intrinsic limitation of the data is that if

a stock is bought and sold during the same quarter, that buy decision is unobserved. This would only pose a problem if the error measures of the unobserved and observed trades differ in a way which is dependent on organizational structure. A priori, there are no reasons to believe that this might be the case.

The returns used to determine if an investment was a good or a bad one, were the quarterly returns of each stock from the end of quarter t to the end of quarter $t + 1$, as gathered from the CRSP dataset “Monthly Stocks” using the field “holding period return,” which adjusts for stock splits and dividends. Note that as the exact date at which assets are bought is unknown (i.e., the holdings dataset has quarterly resolution), then another intrinsic limitation of the dataset is that the return accrued since a stock is bought until the end of that quarter is not accounted for. This lack of data should affect the results of the study in a conservative way, because if managers have an ability to minimize the errors they make, this ability should be more noticeable closer to the decision, and not later when more unpredictable events may affect the price of what they bought.

The investment universe of a fund at time t was defined as all the stocks available to be bought at time t from the union of all the holdings reported by the fund in a trailing window of seven quarters including the current one (i.e., using the last seven Forms 13F reported by the fund). There are at least three other ways to define the investment universe, but these alternatives present conceptual and practical problems that make them less preferable to the trailing-period definition. The first alternative is to use the investment objective each fund typically reports in its prospectus; but this information is imprecise⁸ and not always available, and hence defining the investment universe would have a subjective quality. A second alternative considered is to include all the 5833 stocks ever bought by all the funds. This approach was discarded because it is unfair to count as omissions not buying stocks that would never be bought by a fund (e.g., a utilities fund does not buy high-tech firms). A third alternative is to use the union of all the stocks ever bought by the funds that share the same Morningstar investment category. Similarly to the previous alternative, this method creates very loose investment universes, leading to a similar, albeit less serious, unfairness problem. In sum, letting the deeds of the fund speak for itself seemed the most appropriate choice.

⁸For example, a fund may say that it attempts to track a broad index like the S&P500, but this does not imply that it only invests in stocks part of the index, many of its investments may fall outside of it; another fund may say it invests in small caps (which is a broad category with thousands of stocks), while its investments consistently fall within a group of less than one hundred stocks.

In robustness checks (available from the author) the third alternative definition produced results which are qualitatively similar to those reported here using the trailing-period definition.

6 Results

Given that the data and the measures used are to some degree novel, the statistical tests of hypotheses are accompanied by exploratory data analysis (Tukey, 1977) aimed at uncovering the structure of the data and gaining insights that would pass undetected by only running regressions. Each hypothesis is tested using OLS regressions of the form

$$\text{dependent variable} = \text{structure dummies} + \text{controls} + \text{error},$$

where the dependent variable is the logarithm of the portfolio size to test H1, omission error to test H2, and commission error to test H3. The independent variable of the study, organizational structure, is coded as two dummies representing the decentralized and the individual structure (the centralized structure is the omitted dummy).

The controls used, which are in line with those used in the mutual fund literature, are: (a) the risk profile of the fund, as measured by its Beta with respect to the S&P500; (b) a measure of the experience of the parent firm (the firm owning the fund), as proxied by the logarithm of the number of mutual funds that the parent firm owns (within the universe of 1087 stock mutual funds tracked by Morningstar); (c) the size of the fund, as measured by the logarithm of the net assets managed by the fund (in millions of dollars); and (d) seven investment category dummies, as coded by Morningstar (Large Growth, Large Blend, Large Value, Mid-Cap Growth, Small Growth, Small Blend, and Mid-Cap Blend). Eighty percent of the funds fell in any of these seven categories, while the remaining twenty percent was consolidated in an ‘Other’ class grouping thirteen smaller categories, and used as the omitted dummy in the regressions.

To avoid a source of endogeneity, all the controls were measured at the beginning of the period used to compute the dependent variables (beginning of 2004Q4). To counter the effects of heteroscedasticity, and because observations coming from funds that belong to the same parent firm may not be independent, the standard errors are computed using cluster-robust estimation

		mean	sd	min	max	1	2	3	4	5
1	Log(Portfolio Size)	4.40	0.75	2.92	8.15	1.00				
2	Omission Error	-0.16	0.48	-2.77	1.09	-0.02	1.00			
3	Commission Error	0.14	0.47	-1.03	2.35	0.20	0.33	1.00		
4	Beta	1.15	0.27	-0.09	2.77	0.04	0.01	0.05	1.00	
5	Log(Parent Experience)	2.14	1.10	0.00	4.68	0.32	-0.11	0.06	-0.04	1.00
6	Log(Net Assets)	6.17	1.68	-0.24	11.25	0.27	0.06	0.08	-0.14	0.32

Table 2: Descriptive statistics and correlations ($N = 609$).

(Williams, 2000), with clusters defined according to the parent firm. All the p-values reported correspond to two-tailed tests—this is a conservative decision, as the model presented in §3 predicts relations of the form $a < b$, which just call for running one-tailed tests on the independent variables.

Table 2 displays summary statistics and correlations for the controls and dependent variables. The correlations show no evidence of multicollinearity, which is reaffirmed by the variance inflation factors, none of which was larger than 1.6, a number well below the customary threshold of 10. Of the 609 funds in the dataset, the most common structure is the individual manager (324 funds), followed by the centralized structure (233 funds), and the decentralized structure (52 funds).

6.1 Number of projects accepted

Table 3 shows descriptive statistics disaggregated per organizational structure, for portfolio size (row 1), number of stocks bought (row 2), and investment universe size (row 3). The average fund held 123.3 stocks, bought 16.5 stocks per quarter, and had an investment universe of 194.7 stocks per quarter. These averages display great dispersion. Portfolio sizes varied from 18.6 to 3455.2 (the decimals are because these are ten-period averages), and had a standard deviation of 230.6, or almost twice the mean.

An interesting relationship becomes apparent when looking at the dispersion in the portfolio sizes of the different structures. As seen in row 1 of Table 3, the centralized and decentralized structures have their means similar to their standard deviations (N/N : $\mu = 91.5$, $\sigma = 105.9$; $N/1$: $\mu = 171.1$, $\sigma = 168.3$), which is the signature of the exponential distribution. This finding reinforces the validity of the categorization used, as it suggests that the N/N and $N/1$ structures represent distinctive populations, each one captured by a well-defined data generating process. Conversely, the overdispersion of the $1/1$'s, hints that this class may be a mixture of populations. In fact, it is

	N/N Centralized ($n = 233$)	1/1 Individual ($n = 324$)	N/1 Decentralized ($n = 52$)	Total ($n = 609$)
	avg (sd) [min, max]	avg (sd) [min, max]	avg (sd) [min, max]	avg (sd) [min, max]
1) #stocks in portfolio	91.5 (105.9) [18.6, 1220.3]	138.6 (293.7) [20.3, 3455.2]	171.1 (168.3) [25.6, 990.9]	123.3 (230.6) [18.6, 3455.2]
2) #stocks bought per quarter	15.9 (27.9) [1.2, 280.0]	15.3 (16.7) [1.3, 131.2]	26.1 (27.2) [3.4, 148.2]	16.5 (22.7) [1.2, 280.0]
3) #stocks in investment universe	161.3 (209.4) [22.0, 2143.3]	205.0 (317.0) [22.7, 3547.9]	280.3 (254.1) [49.8, 1433.5]	194.7 (276.9) [22.0, 3547.9]

Table 3: Descriptive statistics—number of stocks per organizational structure.

likely that the funds that use an algorithmic investment approach (e.g., those tracking indices like the S&P500) are overrepresented in the population of 1/1's, because the algorithm greatly reduces the need for managers no matter the number of stocks in which it invests. The great dispersion in the portfolio size of the 1/1's summed to the fact that the predictions for this structure fall in a middle ground between the predictions for the two other structures, suggests that the tests involving structure 1/1 should not exhibit high statistical significance.

The average portfolio size of each structure (row 1 of Table 3) is distributed as predicted by Equation 1: structures N/N , $1/1$, and $N/1$, each has an increasingly larger portfolio (for a synoptic comparison between the predicted and the actual results, juxtapose figures 1 and 3). This finding seems logical, as it is easy to imagine that, for example, two managers that must agree on what to buy should end up buying fewer stocks than two managers that can act independently.

Interestingly, the finance literature has not identified a relationship between organizational structure and portfolio size, probably because researchers have measured organization simply as number of managers. For example, even if Chen et al. (2004) had data on number of managers (p. 1297) and portfolio size (p. 1290), they do not report if there is a relationship between these numbers. Moreover, had they reported the relationship, it would have probably contradicted their statement that funds hire more managers to invest in additional stocks (p. 1290), as the current

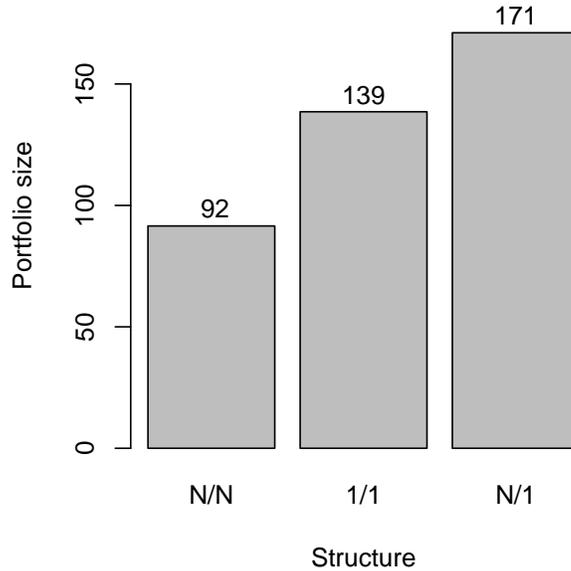


Figure 3: Average portfolio size for each organizational structure. Compare to Figure 1.

dataset shows that the average portfolio size of structures with more than one manager is less than the average portfolio size of the funds with one manager (i.e., Table 3 implies that the weighted average of the portfolio sizes of structures N/N and $N/1$ is 106.0, while the average portfolio size of structure $1/1$ is 138.6).

To test if the relationship between organizational structure and portfolio size is statistically significant, five models were tested (Table 4). Given that the distribution of portfolio sizes is highly skewed, the logarithm of portfolio size was used as dependent variable. In all the models, the decentralized structure was associated with a significantly larger portfolio size than the centralized structure (the effect size corresponds to a 30%–50% increase in portfolio size, depending on the model and the value of the controls). No significant relationship is present for the structure $1/1$, yet the sign of the coefficients associated to it has the predicted direction in all the models.

The coefficients associated with the controls tell stories which are interesting per se. Models A3 to A5 show that funds belonging to more experienced firms hold more stocks, even after controlling for the size of the mutual fund and investment category. One possible interpretation is that more experienced firms have better support structures, allowing managers to track more stocks. The regressions also show that the larger a fund (in net assets), the more stocks it will invest in, which

	Dependent Variable: Log(Portfolio Size)				
	A1	A2	A3	A4	A5
Decentralized (Structure N/1)	0.541*** (0.134)	0.539*** (0.135)	0.485*** (0.128)	0.431** (0.134)	0.436*** (0.128)
Individual (Structure 1/1)	0.169 (0.119)	0.166 (0.119)	0.111 (0.090)	0.106 (0.084)	0.121 (0.083)
Beta		0.086 (0.140)	0.122 (0.148)	0.183 (0.163)	-0.151 (0.162)
Log(Parent Size)			0.209*** (0.056)	0.172*** (0.043)	0.201*** (0.047)
Log(Net Assets)				0.079* (0.030)	0.085** (0.025)
Category effects (joint test)					***
Constant	4.269*** (0.058)	4.172*** (0.164)	3.716*** (0.233)	3.247*** (0.367)	3.384*** (0.385)
Observations	609	609	609	609	609
Adjusted R^2	0.035	0.034	0.125	0.151	0.265

Note. Robust standard errors between parenthesis.

+ $p < 0.10$, * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$ (two-tailed tests).

Table 4: Results of regression analysis of portfolio size.

may reflect that large funds are more likely to run into the liquidity limits of the underlying stocks. Finally, model A5 shows that there is a significant category effect, which gives an additional support to the liquidity explanation, as the categories that have the largest positive coefficients are those involving small companies (categories Small Growth and Small Blend were the only statistically significant categories, with coefficients 0.55 and 0.91 respectively).

Models A1 to A5 were rerun using number of stocks bought per quarter instead of portfolio size, and all the results were qualitatively the same. This increases the confidence on the results, as it shows that what was true for a stock variable (portfolio size) is also true for its corresponding flow variable (number of stocks bought). In all, the large and significant coefficients accompanying the decentralized structure provide evidence that decentralized funds accept more projects than centralized funds (H1).

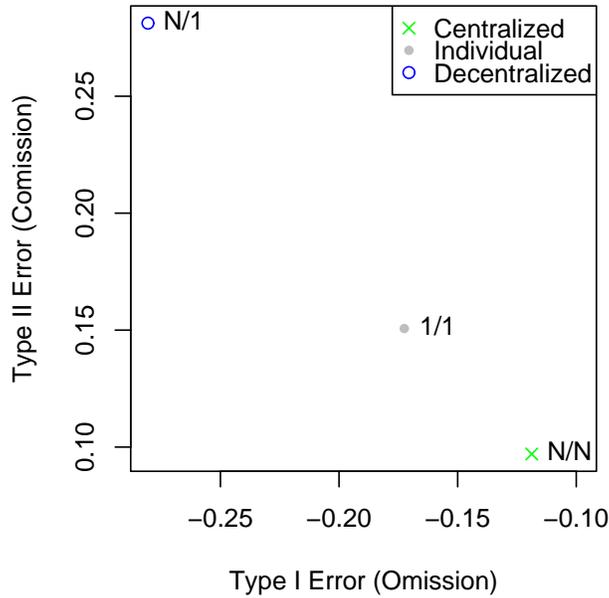


Figure 4: Average (centroids) omission and commission errors of the three organizational structures. Compare to Figure 2.

6.2 Omission and commission errors

Figure 4 displays the average omission and commission error made by each organizational structure. The axes of the figure correspond to the standardized measures described in §5.2. Interestingly, the figure looks exactly as expected, with the centralized fund at the bottom-right, the decentralized fund at the top-left, and the individual manager in between. For a graphic comparison with the expected results, contrast this figure with Figure 2.

All the models in Table 5 support H2, by showing that a decentralized fund makes significantly fewer omissions than a centralized one. The magnitude of the coefficients associated to the decentralized structure is sizable, as it can be shown that decreasing an error by 0.15 points of standardized score is associated to a 13% increase in annual performance (relative to the current performance, e.g., a 10% annual return would become 11.3%).⁹ As in the previous set of regressions, the coefficients accompanying the individual manager have the right sign, but are not statistically significant.

⁹To compute the effect on a fund's annual return, a simulation was run using parameters representative of the average fund. This fund buys 16 stocks from a universe of 195 stocks each quarter, the stock's quarterly returns are drawn from a $N(0.0339, 0.2042)$, the portfolio turnover is one year, and the effect due to superior stock-picking is only effective (this is a conservative assumption) on the quarter after the stock was bought.

	Dependent Variable: Omission Error				
	B1	B2	B3	B4	B5
Decentralized (Structure N/1)	-0.162*	-0.162*	-0.150*	-0.172*	-0.161*
	(0.077)	(0.077)	(0.069)	(0.074)	(0.078)
Individual (Structure 1/1)	-0.054	-0.054	-0.042	-0.044	-0.043
	(0.046)	(0.046)	(0.046)	(0.046)	(0.047)
Beta		0.019	0.011	0.037	0.034
		(0.071)	(0.071)	(0.071)	(0.078)
Log(Parent Size)			-0.047**	-0.062***	-0.064***
			(0.016)	(0.017)	(0.018)
Log(Net Assets)				0.033**	0.035**
				(0.012)	(0.013)
Category effects (joint test)					not sig.
Constant	-0.119**	-0.140	-0.039	-0.235+	-0.217+
	(0.039)	(0.089)	(0.100)	(0.127)	(0.130)
Observations	609	609	609	609	609
Adjusted R^2	0.005	0.004	0.014	0.024	0.022

Note. Robust standard errors between parenthesis.

+ $p < 0.10$, * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$ (two-tailed tests).

Table 5: Results of regression analysis of omission error.

Among the controls, parent experience and net assets appear to be significant determinants of omission errors. The fact that funds owned by more experienced firms make fewer omissions, may suggest that part of the skills to avoid missing investment opportunities may reside in routines which are more likely to exist in larger firms, as could be research support services, fund manager training, or knowledge-sharing among managers of different funds. Conversely, the finding that funds managing more assets make more omission errors may be due to large funds having low incentives to exploit small, yet profitable investment opportunities because their relative contribution to the overall profitability of the fund would be tiny.

All the models on Table 6 support H3, by showing that a decentralized fund makes significantly more commission errors than a centralized one. As before, the coefficients for the individual manager have the predicted sign but are not significant. The fact that parent experience and net assets, which were significant controls in the regressions of omission error, are not significant predictors of commission error may mean that the market is more efficient with respect to commission rather than omission errors. This may happen because not all market participants may agree that a

	Dependent Variable: Commission Error				
	C1	C2	C3	C4	C5
Decentralized (Structure N/1)	0.184** (0.068)	0.183** (0.067)	0.177* (0.071)	0.164* (0.075)	0.146* (0.073)
Individual (Structure 1/1)	0.054 (0.044)	0.051 (0.044)	0.045 (0.042)	0.044 (0.041)	0.045 (0.041)
Beta		0.079 (0.079)	0.083 (0.080)	0.098 (0.088)	0.082 (0.103)
Log(Parent Size)			0.023 (0.019)	0.014 (0.018)	0.014 (0.018)
Log(Net Assets)				0.019 (0.014)	0.018 (0.014)
Category effects (joint test)					not sig.
Constant	0.097** (0.033)	0.008 (0.103)	-0.042 (0.118)	-0.156 (0.178)	-0.183 (0.207)
Observations	609	609	609	609	609
Adjusted R^2	0.008	0.008	0.010	0.012	0.013

Note. Robust standard errors between parenthesis.

+ $p < 0.10$, * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$ (two-tailed tests).

Table 6: Results of regression analysis of commission error.

fund made an omission error (as these depend on agreeing on an investment universe), while a commission error is an unquestionable event. Hence, fund managers may be more motivated to focus on what they are more likely to be assessed, that is, commission errors. An illustration that commission errors are more observed is that after the Internet bubble burst, some investment banks were sued for having recommended Internet stocks to their customers, but it is unheard of a bank being sued for not having recommended a given stock.

Two controls which are typically significant in studies of investment performance—the fund’s Beta and its investment category—are not significant predictors of either omission or commission errors. This occurs because the Monte Carlo mechanism used to compute the errors already controls for these parameters, as each fund is compared in standardized terms against a large number of funds that draw stocks from the same investment universe, and hence on average, have the same Beta and investment category as the focal fund.

7 Discussion

The current study has used mutual funds as a rich data source to explore how organizational structure affects organizational performance. In perfect accordance with the predictions of the model of fallible decision-making presented early in the paper, decentralized structures accept more projects (H1), make fewer omission errors (H2), and make more commission errors (H3) than centralized structures. This section looks at these results in perspective.

7.1 Mutual Funds and Organizational Structure

Two questions that come to mind regarding the organizational design of mutual funds are: is there an optimal organizational structure for mutual funds? and why is the individual manager the most common structure? (53.2 percent of the funds in the dataset used this structure).

For a mutual fund only concerned about maximizing returns, omission and commission errors are equally costly, because not buying a stock that would contribute a 1% of extra return, is as costly as buying a stock that subtracts a 1% of extra return—both cases imply a loss of 1% of returns with respect to a competing fund that did not make that error. Hence, the structure this hypothetical fund should choose is the one that minimizes the sum of both errors. Strikingly, the sum of the omission and commission errors (measured as standardized scores) for each of the three structures is statistically indistinguishable from zero (i.e., if the coordinates of the points on Figure 4 are added, the results are $-0.28 + 0.28 = 0.00$, $-0.17 + 0.15 = -0.02$, and $-0.12 + 0.10 = -0.02$, for structures $N/1$, $1/1$, and N/N , respectively). Given this equivalency in overall errors, it seems natural that most funds choose the structure that is the least expensive. The existence of funds with structures different than $1/1$ may speak about other concerns such as securing continuity against manager turnover, offering promotion opportunities to junior employees, or creating a differentiated product.

The fact that the overall error of each structure is not different from the overall error of picking stocks at random has a special beauty to it: the unpredictability of returns stated by the Efficient Market Hypothesis holds when looking at the overall error, even if each error measured independently is partly predictable.

7.2 Generalizability and Further Work

Since the model used to derive the hypotheses is built on basic information-processing and probabilistic arguments, none of which is specific to mutual funds, it is reasonable to expect the model to generalize to other decision-making settings such as top management teams, boards of directors, venture capital firms, or R&D organizations. One important difference between mutual funds and these other settings is that organizations whose Type I and II errors are equally costly are probably more the exception rather than the rule, hence different organizational structures should not just trade one error for the other, but carry tangible performance differences. Examples of organizations facing unbalanced error costs are juries, which are more concerned with the commission error (e.g., avoid convicting the innocent); the typical IT department, which is presumably more concerned with minimizing commissions (e.g., not leaking sensitive information) rather than minimizing omissions (e.g., implementing every good IT innovation); or a well-funded R&D lab in an industry where first-mover advantages matter, which is more likely concerned with avoiding omissions.

Further research could use alternative settings, or perhaps experiments, to explore the generalizability of the findings, as well as the predictions of the model that the current dataset does not allow testing. Some questions open to empirical examination relate to the position of N/C structures other than the three studied, assessing the consequences of correlation among decision makers' perceptions, and the effects of changing the probability distribution of the incoming projects (\tilde{q}) and the noisiness of decision makers (\tilde{n}). Another line of inquiry, very much in the spirit of contingency theory, could explore if firms that exhibit a better structure–environment fit achieve a higher performance or survival rate. For example, in industries requiring more conservative decision making (i.e., where commissions are costlier than omissions), one would expect firms using structure N/N to surpass those using structure $N/1$.

In more general terms, this paper also suggests that decomposing performance into omission and commission errors can reveal phenomena otherwise unobservable when using standard performance measures. Hence, future research on organizations may benefit from including omission and commission errors as alternative measures of performance.

7.3 Conclusions

From a theoretical point of view, this research presents a mechanism by which micro decisions are aggregated into macro behaviors and links to important questions of strategy research such as “do organizations have predictable biases?” (Cyert and March, 1963:21), “what do we know about the relationships between organizational size (or other stable characteristics) and behavior?” (Rumelt et al., 1994:42), and what is the relationship between decision making and decision outcomes (Zajac and Bazerman, 1991:37).

This research also speaks to the unexplored question of what are the processes that link organizational structure to exploration and exploitation (Siggelkow and Levinthal, 2003:650; Argyres and Silverman, 2004:929; Raisch and Birkinshaw, 2008:380). A relevant observation to address this question is that omission and commission errors are another way of looking at exploration and exploitation (Garud, Nayyar, and Shapira, 1997:33; Garicano and Posner, 2005:157). The logic of this argument is that, on the one hand, firms in unstable or fermenting environments must try to avoid omissions because these curtail the extent of exploration of new high-fitness positions. Illustrations of this behavior are Bill Gates saying that “the real sin is if we [Microsoft’s R&D] miss something” (Hawn, 2004), or Andy Grove’s quip “miss the moment [for change in a high-tech firm such as Intel] and you start to decline” (Stratford, 1993). On the other hand, firms facing stable or incrementally changing environments try to avoid commission errors, as these may disrupt their currently efficient exploitative operations. Examples of these phenomena include Procter and Gamble, where new product proposals are often reviewed more than 40 times before reaching the CEO (Herbold, 2002:74), or IBM’s mainframe-era inspired “non-concur policy,” which enabled any department to veto projects initiated anywhere in the firm (Gerstner, 2003:192–199). Hence, given that this paper has shown how organizational structure can influence the omission and commission errors made by organizations and that previous research has shown that these errors control the degree to which organizations can explore and exploit, this research exposes a mechanism by which organizational structure can influence exploration and exploitation.

From a practical standpoint, this research sheds light on how to use organizations to compensate for shortcomings of individuals, and allows several managerial concerns to be addressed, such as: What organization is needed to avoid exceeding a given error level? Is it true that hierarchy hampers

innovation? What organizational structures can lead to more innovation? In regard to this last question, an important application area is how to enable established organizations to exhibit traits that are usually associated to 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 [...] it 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)

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.” This paper has aimed to shed light on this topic by advancing a small step towards understanding how organizational structure aggregates individual decisions into strategic outcomes.

References

- AMIT, R. AND P. J. H. SCHOEMAKER (1993): “Strategic assets and organizational rent,” *Strategic Management Journal*, 14, 33–46.
- ARGOTE, L. AND H. R. GREVE (2007): “A Behavioral Theory of the Firm—40 years and counting: Introduction and impact,” *Organization Science*, 18, 337–349.
- ARGYRES, N. S. AND B. S. SILVERMAN (2004): “R&D, organization structure, and the development of corporate technological knowledge,” *Strategic Management Journal*, 25, 929–958.
- BARBERIS, N. AND R. H. THALER (2003): “A survey of Behavioral Finance,” in *Handbook of the Economics of Finance, Volume 1B*, ed. by G. M. Constantinides, M. Harris, and R. M. Stultz, Amsterdam: Elsevier, chap. 18, 1053–1124.
- BAVELAS, A. (1950): “Communication patterns in task-oriented groups,” *Journal of the Acoustical Society of America*, 22, 723–730.
- BERGER, J. O. (1985): *Statistical Decision Theory and Bayesian Analysis*, New York, NY: Springer-Verlag, 2nd. ed.
- BIKHCHANDANI, S., D. HIRSHLEIFER, AND I. WELCH (1992): “A theory of fads, fashion, custom, and cultural-change as informational cascades,” *Journal of Political Economy*, 100, 992–1026.
- BOWER, J. L. (1970): *Managing the Resource Allocation Process: A Study of Corporate Planning and Investment*, Boston, MA: Division of Research, Graduate School of Business Administration, Harvard University.
- BOWER, J. L. AND C. G. GILBERT, eds. (2005): *From Resource Allocation to Strategy*, Oxford, UK: Oxford University Press.
- BURNS, T. AND G. M. STALKER (1961): *The Management of Innovation*, London: Tavistock Publications.
- CHAKRAVARTHY, B. S. AND R. E. WHITE (2002): “Strategy process: Forming, implementing and changing strategies,” in *Handbook of Strategy and Management*, ed. by A. M. Pettigrew, H. Thomas, and R. Whittington, London, UK: SAGE Publications, chap. 9, 182–205.
- CHANDLER, A. (1962): *Strategy and Structure: Chapters in the History of American Industrial Enterprise*, Cambridge, MA: MIT Press.
- CHEN, J., H. HONG, M. HUANG, AND J. D. KUBIK (2004): “Does fund size erode mutual fund performance? The role of liquidity and organization,” *American Economic Review*, 94, 1276–1302.
- CHEVALIER, J. AND G. ELLISON (1999): “Are some mutual fund managers better than others? Cross-sectional patterns in behavior and performance,” *Journal of Finance*, 54, 875–899.
- CHRISTENSEN, M. AND T. KNUDSEN (2002): “The architecture of economic organization: Toward a general framework,” LINK Working Paper.
- (2007): “The human version of Moore-Shannon’s theorem: The design of reliable economic systems,” Available at SSRN: <http://ssrn.com/paper=996311>.
- COHEN, L., A. FRAZZINI, AND C. J. MALLOY (2007): “The small world of investing: Board connections and mutual fund returns,” NBER Working Paper No. W13121 Available at SSRN: <http://ssrn.com/abstract=988932>.
- COHEN, M. D., J. G. MARCH, AND J. P. OLSEN (1972): “A garbage can model of organizational choice,” *Administrative Science Quarterly*, 17, 1–25.
- COLEMAN, J. S. (1990): *Foundations of Social Theory*, Cambridge, MA: Belknap Press.
- CONNOLLY, T. AND L. ORDÓÑEZ (2003): “Judgment and Decision Making,” in *Handbook of Psychology: Industrial and Organizational Psychology, Vol. 12*, ed. by W. C. Borman, D. R. Ilgen, and R. J. Klimoski, New York, NY: Wiley, 493–517.
- CSASZAR, F. A. (2007): “An efficient frontier in organization design,” Available at SSRN: <http://ssrn.com/abstract=985322>.
- CYERT, R. M. AND J. G. MARCH (1963): *A Behavioral Theory of the Firm*, Englewood Cliffs, NJ: Prentice-Hall.

- DAFT, R. L. AND A. Y. LEWIN (1993): "Where are the theories for the new organizational forms: An editorial essay," *Organization Science*, 4, i–vi.
- DAVIS, J. H. (1973): "Group decision and social interaction: Theory of social decision schemes," *Psychological Review*, 80, 97–125.
- DAVISON, A. C. AND D. V. HINKLEY (1997): *Bootstrap Methods and Their Application*, New York, NY: Cambridge University Press.
- DONALDSON, L. (2001): *The Contingency Theory of Organizations*, Thousand Oaks, CA: SAGE Publications.
- EFRON, B. (1979): "Bootstrap methods: Another look at the jackknife," *Annals of Statistics*, 7, 1–26.
- EFRON, B. AND R. J. TIBSHIRANI (1993): *An Introduction to the Bootstrap*, New York, NY: Chapman & Hall.
- ETHIRAJ, S. K. AND D. A. LEVINTHAL (2004): "Modularity and innovation in complex systems," *Management Science*, 50, 159–173.
- FAMA, E. F. (1970): "Efficient capital markets: Review of theory and empirical work," *Journal of Finance*, 25, 383–423.
- 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, 331–349.
- GARICANO, L. AND R. A. POSNER (2005): "Intelligence failures: An organizational economics perspective," *Journal of Economic Perspectives*, 19, 151–170.
- GARUD, R., P. R. NAYYAR, AND Z. B. SHAPIRA (1997): "Technological choices and the inevitability of errors," in *Technological Innovation: Oversights and Foresights*, ed. by R. Garud, P. R. Nayyar, and Z. B. Shapira, New York: Cambridge University Press, chap. 3, 20–40.
- GAVETTI, G., D. A. LEVINTHAL, AND W. OCASIO (2007): "Neo-Carnegie: The Carnegie School's past, present, and reconstructing for the future," *Organization Science*, 18, 523–536.
- GERSTNER, L. V. (2003): *Who Says Elephants Can't Dance?: Leading a Great Enterprise Through Dramatic Change*, New York, NY: HarperBusiness.
- GOETZMANN, W. N. AND R. G. IBBOTSON (1994): "Do winners repeat," *Journal of Portfolio Management*, 20, 9–18.
- GREEN, D. M. AND J. A. SWETS (1966): *Signal Detection Theory and Psychophysics*, New York, NY: Wiley.
- GRINBLATT, M. AND S. TITMAN (1992): "The persistence of mutual fund performance," *Journal of Finance*, 47, 1977–1984.
- GUETZKOW, H. AND H. A. SIMON (1955): "The impact of certain communication nets upon organization and performance in task-oriented groups," *Management Science*, 1, 233–250.
- GULICK, L. H. AND L. F. URWICK, eds. (1937): *Papers on the Science of Administration*, New York, NY: Institute of Public Administration, Columbia University.
- HANLEY, J. A. AND B. J. MCNEIL (1982): "The meaning and use of the area under a receiver operating characteristic (ROC) curve," *Radiology*, 143, 29–36.
- HANSEN, L. K. AND P. SALAMON (1990): "Neural network ensembles," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 12, 993–1001.
- HAWN, C. (2004): "Microsoft skills," *Fast Company*, 89 (December 2004), 70.
- HERBOLD, R. J. (2002): "Inside Microsoft: Balancing creativity and discipline," *Harvard Business Review*, 80, 73–79.
- HINSZ, V. B., R. S. TINDALE, AND D. A. VOLLRATH (1997): "The emerging conceptualization of groups as information processors," *Psychological Bulletin*, 121, 43–64.
- JANIS, I. L. (1972): *Victims of Groupthink: A Psychological Study of Foreign-Policy Decisions and Fiascoes*, Boston, MA: Houghton, Mifflin.
- KERR, N. L., R. J. MACCOUN, AND G. P. KRAMER (1996): "Bias in judgment: Comparing individuals and groups," *Psychological Review*, 103, 687–719.

- KIRZNER, I. M. (1973): *Competition and Entrepreneurship*, Chicago, IL: University of Chicago Press.
- KNUDSEN, T. AND D. A. LEVINTHAL (2007): "Two faces of search: Alternative generation and alternative evaluation," *Organization Science*, 18, 39–54.
- LAWRENCE, P. R. AND J. W. LORSCH (1967): *Organization and Environment: Managing Differentiation and Integration*, Boston, MA: Graduate School of Business Administration, Harvard University.
- LAZER, D. AND A. FRIEDMAN (2007): "The network structure of exploration and exploitation," *Administrative Science Quarterly*, 52, 667–694.
- LEAVITT, H. J. (1951): "Some effects of certain communication patterns on group performance," *Journal of Abnormal and Social Psychology*, 46, 38–50.
- LENER, J. (1994): "The syndication of venture capital investments," *Financial Management*, 23, 16–27.
- MAKADOK, R. AND G. WALKER (2000): "Identifying a distinctive competence: Forecasting ability in the money fund industry," *Strategic Management Journal*, 21, 853–864.
- MALKIEL, B. G. (2003): "The efficient market hypothesis and its critics," *Journal of Economic Perspectives*, 17, 59–82.
- MARENGO, L., G. DOSI, P. LEGRENZI, AND C. PASQUALI (2000): "The structure of problem-solving knowledge and the structure of organizations," *Industrial and Corporate Change*, 9, 757–788.
- MARITAN, C. A. AND D. E. SCHENDEL (1997): "Strategy and decision processes: What is the linkage?" in *Strategic Decisions*, ed. by V. Papadakis and P. Barwise, London: Kluwer Academic Publishers.
- MARSCHAK, J. AND R. RADNER (1972): *Economic Theory of Teams*, New Haven, CT: Yale University Press.
- MEYER, A. D., A. S. TSUI, AND C. R. HININGS (1993): "Configurational approaches to organizational analysis," *Academy of Management Journal*, 36, 1175–1195.
- MOORE, E. F. AND C. E. SHANNON (1956/1993): "Reliable circuits using less reliable relays (I and II)," in *Claude Elwood Shannon: Collected Papers*, ed. by N. J. A. Sloane and A. D. Wyner, New York, NY: Wiley-IEEE Press, 796–830.
- MORNINGSTAR (2008): *Morningstar Funds 500*, Hoboken, NJ: Wiley.
- 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*, New York, NY: Norton.
- PURANAM, P., B. C. POWELL, AND H. SINGH (2006): "Due diligence failure as a signal detection problem," *Strategic Organization*, 4, 319–348.
- RADNER, R. (1992): "Hierarchy: The economics of managing," *Journal of Economic Literature*, 30, 1382–1415.
- RAISCH, S. AND J. BIRKINSHAW (2008): "Organizational ambidexterity: Antecedents, outcomes, and moderators," *Journal of Management*, 34, 375–409.
- RAUSAND, R. AND A. HØYLAND (2004): *System Reliability Theory: Models, Statistical Methods, and Applications*, Hoboken, NJ: Wiley-Interscience, 2nd. ed.
- RIVKIN, J. W. AND N. SIGGELKOW (2003): "Balancing search and stability: Interdependencies among elements of organizational design," *Management Science*, 49, 290–311.
- ROBBINS, S. P. (1990): *Organization Theory: Structure, Design, and Applications*, Englewood Cliffs, NJ: Prentice Hall, 3rd ed.
- RUMELT, R. P., D. E. SCHENDEL, AND D. J. TEECE (1994): *Fundamental Issues in Strategy: A Research Agenda*, Boston, MA: Harvard Business School Press.
- SAH, R. K. AND J. E. STIGLITZ (1986): "The architecture of economic systems: Hierarchies and polyarchies," *American Economic Review*, 76, 716–727.
- (1988): "Committees, hierarchies and polyarchies," *Economic Journal*, 98, 451–470.
- SCHMIDT, R. N. (1958): "Executive decision-making," *Academy of Management Journal*, 1, 36–44.
- SHANE, S. (2000): "Prior knowledge and the discovery of entrepreneurial opportunities," *Organization Science*, 11, 448–469.

- SIGGELKOW, N. (2002): "Evolution toward fit," *Administrative Science Quarterly*, 47, 125–159.
- SIGGELKOW, N. AND D. A. LEVINTHAL (2003): "Temporarily divide to conquer: Centralized, decentralized, and reintegrated organizational approaches to exploration and adaptation," *Organization Science*, 14, 650–669.
- SIMON, H. A. (1947/1997): *Administrative Behavior*, New York, NY: The Free Press, 4th. ed.
- STASSER, G. AND W. TITUS (1985): "Pooling of unshared information in group decision-making: Biased information sampling during discussion," *Journal of Personality and Social Psychology*, 48, 1467–1478.
- STONER, R. A. F. (1961): "A comparison of individual and group decisions involving risk," Master's thesis, MIT, School of Industrial Management.
- STRATFORD, S. (1993): "Andy Grove: How Intel makes spending pay off," *Fortune*, February 22, 58.
- THOMPSON, J. (1967): *Organizations in Action: Social Science Bases in Administrative Theory*, New York, NY: McGraw-Hill.
- TUKEY, J. W. (1977): *Exploratory Data Analysis*, Reading, MA: Addison-Wesley.
- VAN FLEET, D. D. AND A. G. BEDEIAN (1977): "A history of the span of management," *Academy of Management Review*, 2, 356–372.
- VAN ZANDT, T. P. (1999): "Decentralized information processing in the theory of organizations," *Contemporary economic issues, Volume 4: Economic design and behavior*, 125–160.
- VON NEUMANN, J. (1956): "Probabilistic logics and the synthesis of reliable organisms from unreliable components," in *Automata Studies*, ed. by C. E. Shannon and J. McCarthy, Princeton, NJ: Princeton University Press, 43–98.
- WILLIAMS, R. L. (2000): "A note on robust variance estimation for cluster-correlated data," *Biometrics*, 56, 645–646.
- WOODWARD, J. (1965): *Industrial Organization: Theory and Practice*, London, UK: Oxford University Press.
- ZAJAC, E. J. (1992): "Relating economic and behavioral perspectives in strategy research," *Advances in Strategic Management*, 8, 69–96.
- ZAJAC, E. J. AND M. H. BAZERMAN (1991): "Blind spots in industry and competitor analysis: Implications of interfirm (mis)perceptions for strategic decisions," *Academy of Management Review*, 16, 37–56.
- ZENGER, T. R. AND W. S. HESTERLY (1997): "The disaggregation of corporations: Selective intervention, high-powered incentives, and molecular units," *Organization Science*, 8, 209–222.