The Power and Limits of Distributed Representations in Strategic Decision-Making*

Felipe A. Csaszar
Ross School of Business
University of Michigan
Ann Arbor, MI 48109
fcsaszar@umich.edu

Luke Rhee
Paul Merage School of Business
University of California, Irvine
Irvine, CA 92617
shyun.rhee@uci.edu

September 30, 2025

Forthcoming in Strategy Science

*We thank Gwen Lee for her editorial guidance and two anonymous reviewers for their insightful and constructive feedback. We are also grateful to Dylan Boynton, Mana Heshmati, and Ji-hyun Kim for valuable comments on earlier drafts. We benefited from helpful suggestions from seminar participants at Stanford Graduate School of Business, the University of St. Gallen, the Duke Strategy Conference, the Tuck Summer Camp, and the Academy of Management Meeting. All remaining errors are our own.

The Power and Limits of Distributed Representations in Strategic Decision-Making

Abstract

This paper develops a formal theory of distributed representations—collective cognitive models that emerge when organizations aggregate the simplified mental models of multiple members—and examines when they enhance or hinder strategic decision-making. We extend Brunswik's lens model to multiple decision-makers and introduce "decision boundaries" from machine learning to explain how aggregation structures interact with individual internal representations across varying task environments. Using a mathematical model of project screening, we compare two prototypical aggregation rules (Averaging and Unanimity) against individual Specialists (single-cue experts) and Generalists (multi-cue learners) across various environments and levels of experience. Our analysis reveals that effectiveness depends critically on the three-way interaction between internal representations, aggregation structure, and environmental conditions: Specialists excel when one cue dominates; Unanimity guards against errors when good projects are rare and decision-makers lack experience; Averaging delivers robust performance across most settings; while only highly experienced Generalists outperform distributed representations, though such individuals are scarce in practice. These findings advance microfoundations by linking individual cognition and organizational aggregation, enrich the attention-based view by showing how cognitive processing and aggregation matter beyond attention allocation, and offer actionable guidance for designing decision processes under strategic uncertainty.

Keywords: distributed representation; mental representation; aggregation; organization structure; decision boundary; evaluation

"All models are wrong but some are useful."

—George E. P. Box (1979:202)

1 Introduction

1.1 Distributed Representations in Strategic Decision-Making

Internal representations (also known as mental representations) are simplified models of the environment that boundedly-rational individuals hold in their minds and use to make predictions and guide decision-making (Craik 1943:61). Because of their simplified nature, internal representations inevitably blind individuals to aspects of the environment (Tripsas and Gavetti 2000, Martignoni et al. 2016, Csaszar and Ostler 2020) and result in a partial understanding due to the aspects ignored or imperfectly captured by the representation (Gavetti et al. 2007:530, Helfat and Peteraf 2015:834, Laureiro-Martínez and Brusoni 2018:1033, Ocasio et al. 2020:233).

Despite these individual limitations, organizations can in principle transcend any single simplification by fusing multiple internal representations into a *distributed representation*—the collective cognitive model that emerges when organizations aggregate the simplified internal representations of multiple members to guide decision-making. By integrating diverse opinions, organizations can encompass a broader range of perspectives and experiences (Csaszar 2012, Knudsen and Levinthal 2007, Keum and See 2017, Hong and Page 2004, Joseph and Gaba 2020), often through aggregation structures such as voting and averaging (Sah and Stiglitz 1986, Seshadri and Shapira 2003, Csaszar and Eggers 2013). Schematically: Distributed representation = internal representations + aggregation structure.[1]

Yet the relationship between distributed representations and organizational performance remains poorly understood. Consider a CEO deciding whether to launch a novel car design.

[1]Distributed representations can also include artifacts such as software, which may function either as a proxy for human internal representations or as part of the aggregation mechanism itself.

After consulting engineering and marketing managers—each a specialist with a distinct internal representation—the CEO receives contradictory assessments: one predicts great success, the other moderate failure. Should the CEO follow one manager's advice, require unanimity, average the opinions, or perhaps decide independently? The optimal approach likely depends on environmental factors such as the interdependence between engineering and marketing domains or the degree of experience of the different individuals. More broadly, this raises the question of whether organizations are better served by aggregating the views of multiple specialists or by relying on a generalist who integrates those perspectives internally—a choice that may appear equivalent at first glance but in fact produces markedly different outcomes across varying conditions.

Challenges like these pervade strategic decision-making. Top management teams, boards of directors, and investment committees routinely face situations where members hold divergent views shaped by their unique backgrounds and experiences. Start-up co-founders may disagree on technology choices, functional heads may value strategic alliances differently, and venture capital partners may assess investment opportunities through incompatible lenses. In each case, the organization must determine how to aggregate these diverse perspectives into a collective decision that hopefully produces the best possible performance.

Despite the ubiquity of distributed representations in strategic contexts, we lack a systematic understanding of how different aggregation structures interact with diverse internal representations across varying task environments. This gap limits our ability to predict when distributed representations contribute to organizational performance. Therefore, we ask: When do distributed representations enhance or hinder organizational decision-making, and how does this relationship depend on individuals' internal representations, the aggregation structure used, and the characteristics of the task environment?

1.2 Our Approach and Contribution

To address this question, we develop a theoretical framework examining how distributed representations affect organizational performance, contingent on the choice of aggregation structure, the characteristics of internal representations, and the conditions of the task environment. Our goal is to examine not only the strengths but also the limitations of distributed representations to clarify how and when they bolster or undermine organizational performance.

To do so, we develop a parsimonious mathematical model examining how organizations screen projects (e.g., acquisitions, new products, investments) when decision-makers possess limited internal representations. We focus on two fundamental aggregation structures: Averaging, where the organization equally weights different opinions, and Unanimity, where all decision-makers must agree. These structures, commonly observed in committees and management teams, aggregate the opinions of individuals whose internal representations differ based on their focus—for example, specialists who focus on individual cues and generalists who consider multiple cues simultaneously—and their experience, which shapes how accurately they interpret those cues.

Our model examines performance across environments that vary along four key dimensions: munificence (the proportion of good projects), complexity (interactions between project attributes), dominance (relative importance of different cues), and uncertainty (unpredictability in outcomes). By comparing distributed representations against individual specialists and generalists, we identify when aggregating diverse perspectives improves or impairs organizational performance.

Our paper contributes to the literature in several ways. First, recognizing that the effectiveness of distributed representations depends on the choice of aggregation structure and the characteristics of individual internal representations across different environmental conditions, our model describes the conditions under which different distributed representations are preferable. For example, our model shows that when the environment is munificent and dominated by a single cue, relying on the corresponding Specialist yields the best results. When good projects are rare and managers are inexperienced, Unanimity—requiring both specialists to agree—provides valuable protection against errors. In all other cases, Averaging the specialists' opinions offers the most robust performance. While a highly experienced Generalist can often outperform these distributed representations, such individuals are uncommon in practice.

Second, our paper contributes to the literatures examining the micro-to-macro processes that underpin organizations. We do so by elaborating and integrating the roles played by representation and aggregation in driving organizational performance. The separate roles of representation and aggregation have been emphasized by previous research—for example, Csaszar (2018:616) states that "nothing in strategy makes sense except in the light of representations" and Barney and Felin (2013:145) say that "aggregation is the sine qua non of microfoundations"—yet exactly how organizational performance depends on the joint effect of these two processes has remained undertheorized.

By formalizing the concept of distributed representation and systematically investigating its strengths and limitations across varying contexts, our work provides a more rigorous foundation for understanding organizational cognition. This formal approach allows us to make precise predictions about how different aggregation structures and internal representations interact to affect organizational performance, revealing patterns that aggregation rules alone cannot explain—such as when aggregating specialists creates a collective model capable of capturing complex environmental interactions that no specialist can perceive alone. In this way, we advance the view of organizations as systems of distributed representations.

Finally, we make a methodological contribution by introducing the concept of "decision boundary," borrowed from machine learning, to explain the mechanism by which an aggregation structure interacts with limited internal representations to determine organizational performance in a given environment. This approach allows us to provide a more formal understanding of distributed representations and how they emerge from the interaction of individual internal representations and aggregation structures.

The rest of the paper is structured as follows. Section 2 provides the background necessary to motivate and understand our theory. Sections 3 and 4 present and analyze our model. Section 5 discusses the broader theoretical and practical implications of our work.

2 Theoretical Background

The main ideas underpinning the theory in this paper are that (i) internal representations imply cognitive limitations, (ii) distributed representations may address these limitations, and (iii) such representations can be studied formally by building on ideas from the judgment and machine learning literatures. Sections 2.1, 2.2, and 2.3 develop these three themes, respectively. Finally, Section 2.4 shows how this study differs from and extends previous work.

2.1 Internal Representations and Cognitive Limits

In its aim to build a behaviorally plausible theory of organizations, the Carnegie tradition imported from the then-nascent field of cognitive science the idea that decisions are made based on an internal representation: "choice is always exercised with respect to a limited, approximate, simplified 'model' of the real situation" (March and Simon 1958/1993:160). An initial empirical demonstration of internal representations' relevance to organizations was provided by Dearborn and Simon (1958), who showed that executives' functional backgrounds affected their judgments. In their experiment, when presented with the same business case, managers with a sales background tended to identify the main problem as sales-related, while those with a production background saw it as production-related. Building on these insights, extensive research has shown that internal representation significantly influences managers' decisions. For instance, internal representations have been shown to affect managers' competitive assessments (Porac et al. 1995), search behavior (Gavetti and Levinthal 2000), strategic choices (Gavetti et al. 2005), opportunity recognition (Shepherd et al. 2017), and strategic foresight (Heshmati and Csaszar 2024).

Given the constraints of human cognition, internal representations recast an arbitrarily complex task environment into a form that is manageable for boundedly rational decision-makers. This process of simplification is analogous to how a map selectively depicts only certain features of a landscape (March 2006). Inevitably, such simplification leads to an incomplete understanding of the environment (Walsh 1995:280–281). All else being equal, individuals whose internal representations more accurately capture the critical aspects of the task environment are better positioned to make effective decisions—a relationship consistently supported by both empirical (Bourgeois 1985, Tripsas and Gavetti 2000, Gary and Wood 2011, Smith 2014, Csaszar and Laureiro-Martínez 2018) and theoretical research (Gavetti and Levinthal 2000, Denrell et al. 2004, Gavetti 2012, Puranam and Vanneste 2016, Martignoni et al. 2016, Csaszar and Ostler 2020).

Because incomplete internal representations can undermine decision quality, enhancing the internal representations of organizational decision-makers is an appealing lever for improving performance. However, this is far from straightforward, for at least three reasons. First, because internal representations comprise both explicit and tacit knowledge (Polanyi 1966, Winter and Szulanski 2001), individuals are often unaware of them and unlikely to seek improvement. Second, because tacit knowledge is difficult to transfer, internal representations typically develop through experience and take time to learn (Simon and Chase 1973, Ericsson et al. 1993). Finally, once developed, an internal representation functions as a “thought world” that filters out non-conforming information, reinforcing reliance on the current representation and limiting improvements or corrections (Dougherty 1992).

2.2 Distributed Representations as a Solution

While individual internal representations are inevitably limited by bounded rationality, organizations can transcend these limits by forming a *distributed representation*: a collective cognitive model that emerges from integrating the diverse, simplified mental models held by multiple members. Distributed representation, a concept rooted in cognitive science (e.g., Norman 1993, Hollan et al. 2000), recasts cognition as an emergent property of interactions among individuals and the aggregation mechanisms that combine their perspectives.

Unlike internal representations that exist within an individual's mind, distributed representations span multiple people. A classic example is landing a large airplane, which requires information shared between the pilot and copilot, with neither having complete information Hutchins (1995). Distributed representations are common in strategic contexts, where decision-making often requires input from various organizational members. They allow groups to collectively "see" more of the environment than any single member, potentially offsetting the blind spots inherent in individual internal representations.

A key element of distributed representations is their *aggregation structure*—the formal or informal rules by which individual perspectives are combined into an organizational-level decision. Classic research on aggregation traces back to the study of voting, averaging, and committee decision rules, with intellectual roots in the Condorcet jury theorem and early organizational theory (Condorcet 1785/1994, Simon 1955, Cohen et al. 1972). In management and economics, foundational work by Sah and Stiglitz (1986, 1988) formalized how different aggregation structures (e.g., unanimity, majority, polyarchy) influence organizational outcomes. More recent studies have systematically compared aggregation mechanisms—including voting, averaging, and delegation—emphasizing their impact on omission and commission errors, as well as on exploration and exploitation (Kerr and Tindale 2004, Knudsen and Levinthal 2007, Csaszar 2012, Csaszar and Eggers 2013). Other research has examined specific features of the aggregation process, such as the degree of centralization (Rivkin and Siggelkow 2003), the isolation of sub-groups (Fang et al. 2010), and the number of dimensions along which performance is evaluated (Ethiraj and Levinthal 2009), as well as broader questions about how aggregation shapes learning (Piezunka et al. 2022) and evolves over time (Sharapov and Dahlander 2025).

Although such studies have provided important insights into the relationship between aggregation structures and performance, they have overlooked the role of internal representations. Individuals are assumed to make imperfect decisions, but how these limitations depend on their internal representations has been outside the scope of this work. However, an aggregation structure alone cannot constitute a distributed representation; it is only through the interaction of aggregation structures with individual internal representations that distributed representations emerge. Aggregation, as a structural element, lacks the interpretative content that individuals with limited internal representations bring to decision-making. Recognizing internal representations, hence, makes it possible to explain why the same aggregation rule sometimes improves performance and at other times worsens it. For example, averaging specialists from distinct domains can approximate complex interactions that no individual perceives, whereas averaging specialists who share a similar representation will only reinforce their shared biases—a mechanism that remains invisible if aggregation is studied separately from internal representation. Because the research on aggregation has been largely acognitive, understanding how aggregation structures and limited internal representations interact to determine organizational performance—and when this either improves or worsens outcomes—has remained unexplored.

2.3 Linking Distributed Representations to Performance

To understand how distributed representations affect organizational performance, we extend Brunswik's (1952) lens model to include multiple individuals and organizational aggregation structures. This extension formalizes the concept of distributed representation and allows us to analyze how individual cognition and aggregation structure jointly determine organizational outcomes.

2.3.1 Internal representations and the lens model. Brunswik's (1952:19–21) "lens model" explains how individuals make judgments based on multiple characteristics or cues. The task captured by Brunswik's model is predicting a value when presented with some of these cues. Brunswik's model conceptualizes the individual and the environment symmetrically: the individual is represented by a function that connects the cues to a predicted value, and the environment is represented by a function that connects cues to a real value. For instance, a manager may predict this year's sales based on cues such as last year's sales (x_1) and the growth of the economy (x_2) according to $\hat{y} = x_1 + 0.5x_2$. In reality, however, the weights may differ, and another cue—say, inflation (x_3)—may also matter; in this case, the actual sales could be given by $y = x_1 + 0.3x_2 - 0.1x_3$.

The lens model is customarily drawn as shown in Figure 1, with the right-hand side depicting the internal representation of an individual and the left-hand side depicting the environment. Sometimes, the prediction ($\hat{y}$) is just an intermediate step to get to a decision (e.g., turning on the house heating if I predict the temperature will go below 15 °C overnight; such a decision rule could be written as $d = \mathbb{1}[\hat{y} < 15]$). One could interpret such a representation as predicting when it is "cold." In that case, the rule determining the decision is also part of the representation. More generally, a representation could be a complex, multi-stage prediction algorithm, such as a deep neural network. The lens model is a standard framework in the judgment and decision-making literature and has influenced the study of many phenomena, including perception, learning, and decision-making (for an overview of the lens model and its applications, see Stewart 2001; see also Hogarth and Karelaia 2007 for a survey of 270 lens model studies).

2.3.2 Extending the lens model for distributed representations. While the right-hand side of Figure 1 illustrates the internal representation of a *single individual*, we propose extending the lens model to encompass *multiple* individuals along with an aggregation structure. This extension offers a framework for describing and modeling distributed representations.

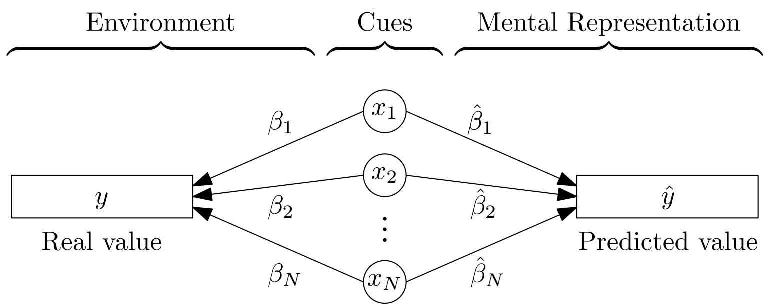

Figure 1: Brunswik's lens model.

Figure 2 illustrates this extension. First, this figure shows two individuals (A and B) who have different internal representations (i.e., they have different $\hat{\beta}$'s for each cue). Second, this figure includes the aggregation structure as a step that occurs after the individuals have formed their opinions (i.e., the aggregation structure takes the opinions of A and B, $\hat{y}^A$ and $\hat{y}^B$, and produces an organization-level prediction $\hat{y}$).[2]

Note that both internal and distributed representations serve the same function—taking inputs (denoted as x) and transforming them into predictions (denoted as $\hat{y}$)—which is why it is appropriate to refer to both as "representations." In addition to internal and distributed representations, strategy commonly uses "external representations," such as visual tools taught in MBA programs. For initial work on external representations in strategy, see Csaszar et al. (2024). For more distinctions among these three types of representation, see Section 2 in Csaszar (2018).

In the Model section, we will use this extension to Brunswik's framework as a starting point to develop our model to explore how distributed representations affect organizational performance.

[2]Although we do not explore it in this paper, this graphical notation could be extended to describe more complex distributed representations, such as including more individuals and more intricate aggregation processes.

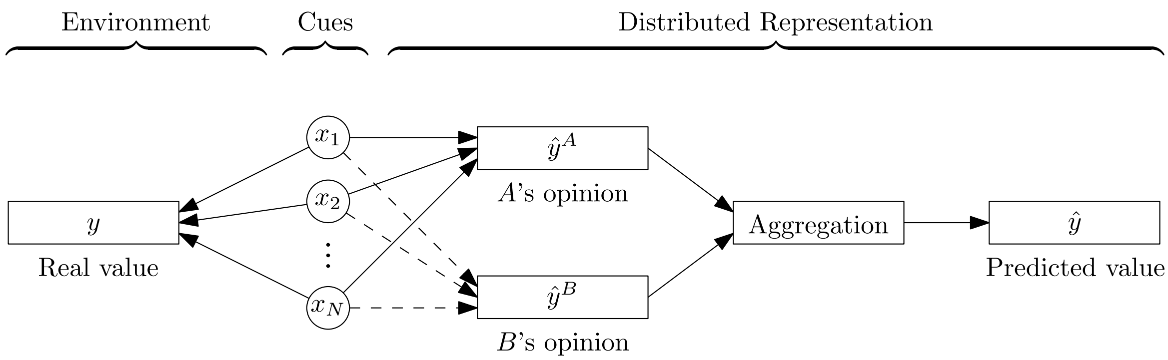

Figure 2: Extending Brunswik's lens model to depict a distributed representation.

2.4 Differences from Previous Work

Our study makes three key departures from existing research, enabling a novel theoretical framework for understanding organizational decision-making and the power and limits of distributed representations.

First, we theorize the joint effect of internal representations and aggregation structure on organizational performance. While extensive literature explores how these elements separately affect organizations, their interaction remains understudied (Ocasio 1997:187–188, Gavetti 2005:599, Posen et al. 2018:235). Recent empirical work has begun integrating these elements (Vuori and Huy 2016, Joseph and Wilson 2018, Rhee et al. 2019), but most discussions of their joint effect on performance remain meta-theoretical (Gavetti 2012, Csaszar 2018, Joseph and Gaba 2020). To the best of our knowledge, only Csaszar and Laureiro-Martínez (2018) and Lee and Csaszar (2020) explicitly study this interaction. Csaszar and Laureiro-Martínez (2018) use experiments to compare individuals' and groups' ability to predict start-up success, while Lee and Csaszar (2020) analyzes video game companies to show how cognitive and structural elements affect sales. Both studies suggest a joint effect but lack a comprehensive theory. We address this gap by formalizing how diverse internal representations and aggregation structures combine to form distributed representations. Our work demonstrates that distributed representations are not uniformly effective; rather, their impact on performance is highly contingent on the interplay between individual cognition, aggregation structures, and the task environment. This nuanced understanding allows us to provide a comprehensive framework with precise, testable predictions.

Second, we extend information aggregation research to complex, multi-dimensional environments. Most prior work has focused on single-dimensional settings, where the law of large numbers can simplify aggregation through averaging or voting to improve accuracy. However, multi-dimensional information cannot be simplified in the same manner. While a few studies, such as Csaszar and Eggers (2013), consider multiple dimensions (e.g., project quality and type), our model uses more broadly defined dimensions, capable of representing a more extensive and varied range of project attributes. Unlike Hong and Page (2004) and Csaszar and Eggers (2013), which examine single-cue environments, we model environmental complexity through interactions among multiple cues. Additionally, we explicitly model internal representations rather than assuming noisy perception or innate search preferences.

Third, we introduce methodological innovations for studying organizational cognition. We extend Brunswik's lens model, which traditionally focuses on individual decision-making, to encompass multiple individuals and aggregation structures (Figure 2), allowing us to analyze distributed representations across an organization. Furthermore, we incorporate the concept of "decision boundaries" from machine learning to explain the mechanism by which aggregation structures interact with limited internal representations. Together, these contributions allow us to formally describe distributed representations and predict their performance across different environments.

These departures provide a foundation for systematically analyzing how the interplay between internal representation and aggregation structure shapes the effectiveness of distributed representations—clarifying the conditions under which they enhance or hinder organizational performance.

3 Model

To study the power and limits of distributed representations, we compare the performance of internal and distributed representations under different types of individual and environmental conditions. The model describes an organization required to *screen* a stream of projects by approving good ones and rejecting the rest. Screening projects is common in settings such as deciding whether to acquire a firm (Haleblian et al. 2006), launch a new product (Masucci et al. 2021), or pick among potential investments (Csaszar 2012). The screening problem is familiar: should the firm acquire a target, green-light a prototype, or fund a research idea? We cast screening in the language of the lens model, which cleanly separates the *environment side*—how observable cues relate to the true quality of a project—from the *representation side*—how decision-makers perceive that relationship. We ask: under what conditions do distributed representations (Averaging or Unanimity) outperform internal representations (Specialist or Generalist)?

Four characteristics of the environment—its munificence, dominance, complexity, and uncertainty—shape the answer, as does the amount of experience each decision maker brings. We first formalize projects and environments, next describe internal and distributed representations, and finally define the performance metric used in the simulations that follow.

3.1 Projects and Environments

Every project is described by two observable cues, x_1 and x_2. Continuing the running automobile example, x_1 might measure exterior color on a scale from 0 (deep blue) to 1 (bright red) while x_2 records horsepower on a scale from 0 (low) to 1 (high). The project's unobservable true quality y—here, profit in millions of dollars—is generated by the environment-side lens equation

$$y = M + Dx_1 + x_2 + Kx_1x_2 + \varepsilon, \tag{1}$$

with noise $\varepsilon \sim \text{Normal}(0, U^2)$.

Interpreting the parameters. The constant M captures *munificence*: when M is large the firm operates in a rich hunting ground where most projects are profitable, while $M < 0$ reflects a barren landscape. The slope D expresses *dominance*: for instance, if color matters five times more than horsepower we set $D = 5$; if the cues are equally important $D = 1$. The coefficient K encodes *complexity.* When $K = 0$ the cues are purely additive; when $K > 0$ they interact so that, for instance, high horsepower pays off particularly well with red cars. Finally U measures *uncertainty*: larger U blurs the cue–quality link so that luck or unforeseeable shocks dominate. Varying (M, D, K, U) generates a rich menu of task environments.

3.2 Internal Representations

Decision makers do not observe Equation (1). Instead, each develops a mental model—an *internal representation*—from E previously seen projects. We study three such models.

Specialist A focuses only on cue 1: $\hat{y}^A = \hat{\beta}_0^A + \hat{\beta}_1^A x_1$. *Specialist B* mirrors her colleague and attends only to cue 2: $\hat{y}^B = \hat{\beta}_0^B + \hat{\beta}_2^B x_2$. A *Generalist* entertains both cues and their interplay: $\hat{y}^G = \hat{\beta}_0^G + \hat{\beta}_1^G x_1 + \hat{\beta}_2^G x_2 + \hat{\beta}_3^G x_1 x_2$.

After observing E historical cases, each person estimates their coefficients via ordinary least squares and obtains a representation $\hat{y}(\cdot)$. The quality of these representations improves with experience. For instance, an engineering specialist who has experienced twenty earlier car launches might estimate $\hat{y}^B = -3 + 5x_2$ and thus predict a profit of 2 for a prototype with horsepower $x_2 = 1$, regardless of color; this is likely more accurate than the estimate of a novice engineering specialist who has only experienced only a handful of car launches.

3.3 Distributed Representations

A *distributed representation* joins the two specialists with an aggregation rule. We consider two rules that are commonplace in organizations:

Unanimity. A project is approved only if both specialists predict $\hat{y}^A > 0$ *and* $\hat{y}^B > 0$. Think of two division heads who must jointly sign off on any investment.

Averaging. The project is approved when the mean opinion is positive, $(\hat{y}^A + \hat{y}^B)/2 > 0$. This mimics committees that blend members' views with equal weight. For example, if one manager forecasts a profit of \$100 million while the other foresees a loss of \$10 million, the averaged opinion is still positive, and the project proceeds.

Figure 3 gathers in one place the four internal and distributed representations described so far. Shaded rectangles correspond to internal representations—mental models held by individual decision makers—whereas dashed rectangles enclose distributed representations that yield a collective judgment. Panel (a) depicts Specialist A, panel (b) the Generalist, while panels (c) and (d) show how the two specialist opinions are combined under Unanimity and Averaging, respectively.

3.4 Performance Metric

We evaluate each structure on its ability to pick winners. Performance is the expected true quality per screened project. Concretely:

(i) For a given (M, D, K, U) environment, draw E historical projects for each specialist, and then fit their internal representations.

(ii) Draw a fresh project (x_1, x_2) from independent Normal$(0, 1)$ cues.

(iii) Plug in the fresh project into the different representations to decide acceptance.

(iv) For each representation, record the realized quality y if accepted, or zero otherwise.

Repeating these steps ten million times yields an accurate Monte-Carlo estimate of $\pi_s(M, D, K, U, E)$, where s in $\{\text{Specialist}, \text{Generalist}, \text{Unanimity}, \text{Averaging}\}$. We report performance on a 0–1 scale. A score of 1 belongs to an omniscient oracle that approves all projects with $y > 0$ and rejects the rest; 0 matches a rubber-stamp screener that approves everything at random.

The next section uses this apparatus to pinpoint when distributed representations can

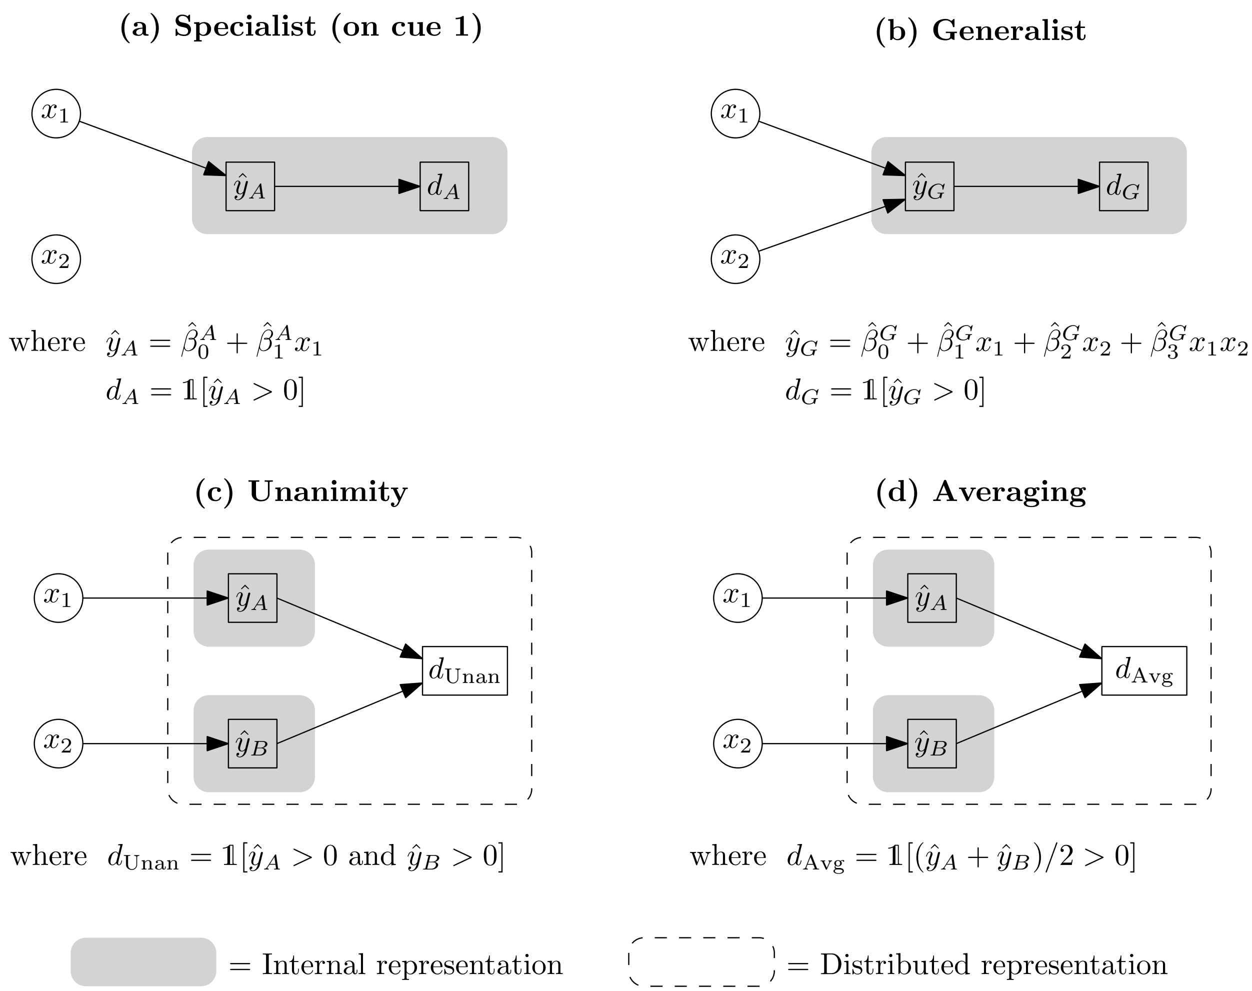

Figure 3: Synopsis of the internal and distributed representations examined in this article.

outperform internal representations—and when they do not.

4 Results

In this section, we analyze the performance of distributed representations by comparing them to internal representations—mental models held by individual decision makers operating in isolation. Our goal is to identify when distributed representations outperform or underperform these internal representations.

To ensure a clear comparison, we first focus exclusively on specialists: we compare the performance of a single specialist acting alone to distributed representations that aggregate the judgments of multiple specialists (using Unanimity and Averaging rules). This approach allows us to isolate how distributed representations add value beyond the bounded insights of specialists, each of whom bases decisions on the single cue aligned with their area of expertise. Later, we introduce generalists—individuals attending to multiple cues and their interaction. Generalists have broader but potentially shallower mental representations. By introducing the generalist later in our analyses, we can more precisely evaluate when distributed representations enhance or hinder organizational performance.

The presentation is organized around figures created after exhaustively exploring the model's behavior and selecting representative scenarios. Appendix C contains code to reproduce all results.

The figures use a common set of parameter values. Complexity (K) and munificence (M) affect performance most elaborately, so we vary these progressively: K from 0 to 8 and M from -5 to 5, as these are the ranges where interesting behavior occurs (for example, when M is outside of this range, projects are either so positive or so negative that they are respectively accepted or rejected regardless of the structure used; likewise with K). The remaining parameters (U, D, and E) affect performance in less elaborate ways and, hence, their effects can be understood by looking at a few parameter levels. Dominance D is examined when both cues are equally relevant ($D = 1$) and when cue 1 dominates ($D = 5$). Experience E is examined at high ($E = 64$) and low ($E = 4$) levels.[3] Varying uncertainty U affects the absolute performance of all representations but rarely changes their relative ranking; thus, we fix U at 2.[4]

[3]High experience ($E = 64$) describes asymptotic behavior, as at that level, in over 99% of the simulations, increasing experience does not improve performance. Low experience ($E = 4$) represents the minimum number of observations necessary to estimate the four coefficients that describe the environment (i.e., Equation (1)).

[4]Uncertainty affects performance ranking only when U is extremely high, making all representations perform equally poorly. See Appendix B for details.

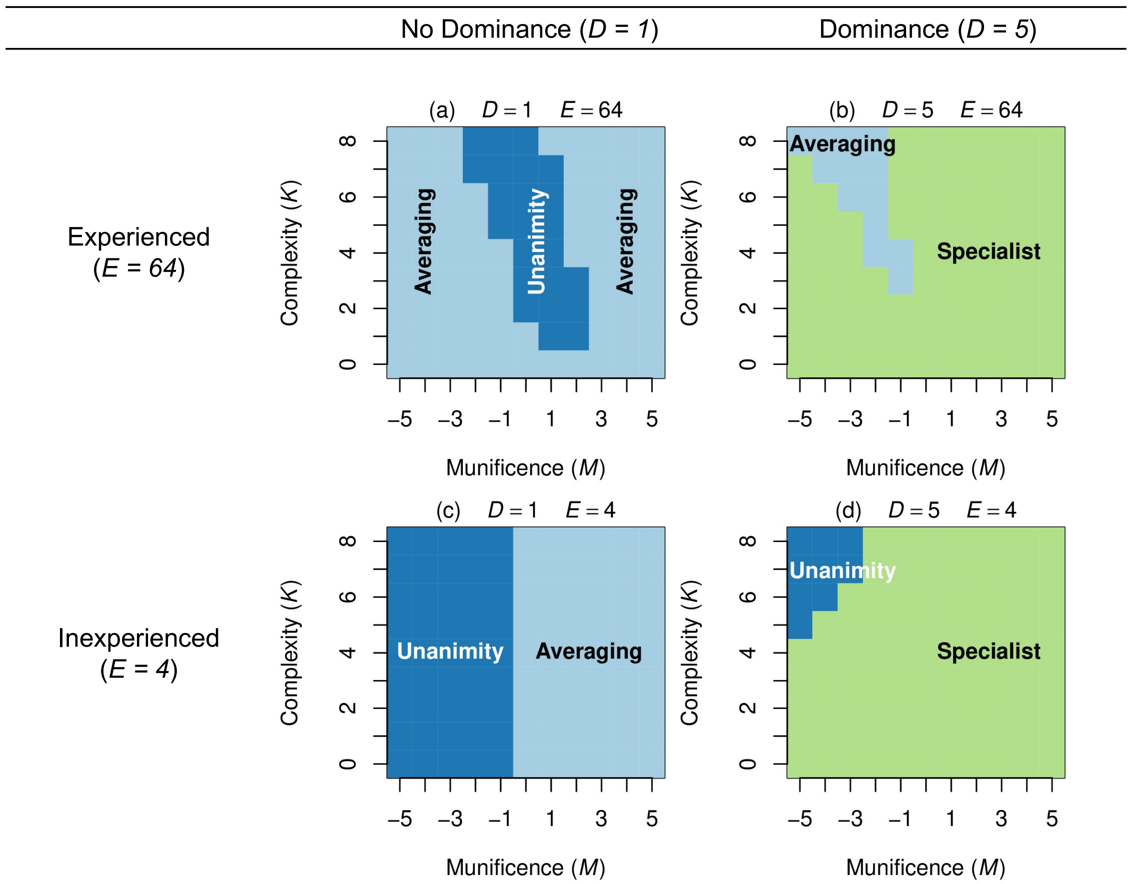

Figure 4: Best-performing structure as a function of munificence (M), complexity (K), dominance (D), and experience (E).

4.1 Overview of the Results

Figure 4 shows the best-performing representations as a function of dominance and experience (D and E, the columns and rows of the panel array) and munificence and complexity (M and K, the x- and y-axes of each panel).

A first look at Figure 4 reveals that the best-performing representation depends on the model parameters in intricate ways. The four panels differ substantially, with irregular patterns indicating complex parameter interactions. Some patterns seem counterintuitive. For instance, when dominance is high ($D = 5$; that is, performance depends mostly on one cue), one might expect delegating to the specialist on that cue to win. However, Averaging and Unanimity often outperform Specialist in panels (b) and (d), respectively. At first sight, then, the patterns in Figure 4 seem puzzling.

4.2 Understanding Performance Through Decision Boundaries

The key to understanding Figure 4's complex patterns is considering the discrimination task performed by representations. Machine learning's concept of "decision boundary" provides an intuitive yet precise explanation.

Nilsson's (1965:4–5) early textbook on machine learning introduced the concept of decision boundary. Nilsson's insight was that a geometric rationale could be used to understand what would otherwise be a very difficult task: predicting the performance of a given machine learning algorithm in a given environment. Since then, decision boundaries have been central to computational learning theory (Minsky and Papert 1969:169, Quinlan 1993:96, Vapnik 1995:4, Kearns and Vazirani 1994:2) and have served as intuitive tools for assessing different algorithms' suitability (Bishop 2006:204, Alpaydin 2014:214). Decision boundaries show how decision rules separate problem spaces into differently classified areas.

Figure 5 illustrates our decision boundary logic. Projects with two cues (x_1 and x_2) and one quality (y) can be visualized on a plane defined by the x_1 and x_2 axes, with quality represented by color (black for positive and red for negative values). Panel (a) illustrates this concept. The challenge for a representation is to predict a project's quality based solely on its position in the (x_1, x_2)-plane. Effective decision boundaries correctly identify most high-quality projects while excluding most low-quality projects.

For internal representations, decision boundaries capture the effect of individual cognition. For example, a specialist on cue 1 corresponds to the gray area in panel (b), including all points to the right of a vertical line because the specialist's prediction follows a linear inequality on cue 1 (i.e., specialist A approves only if $\hat{\beta}_0^A + \hat{\beta}_1^A x_1 > 0$). The precise position of this decision boundary depends on the specialist's experience: as specialists observe more projects, estimated coefficients ($\hat{\beta}$'s) become more accurate. For instance, if the specialist had observed only the two projects closest to the lower-right corner (see the points enclosed

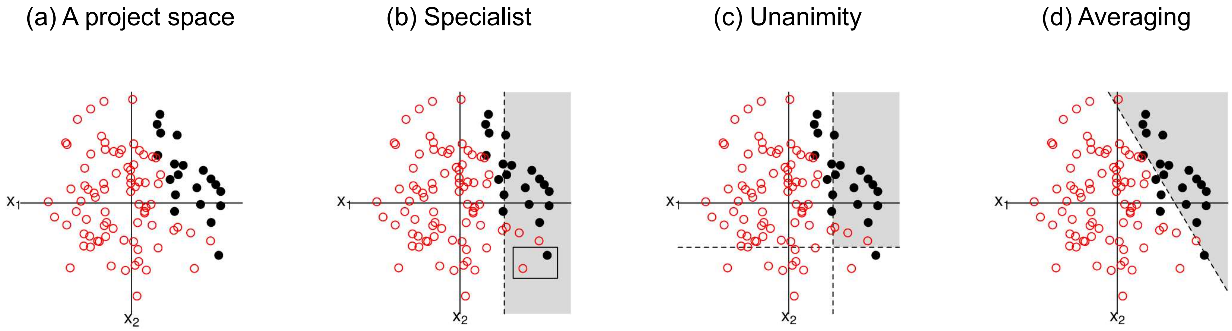

Figure 5: A project space (panel (a)) and the illustration of the decision boundaries of different representations (panels (b)–(d)).

in a box in panel (b)), the specialist would have set a decision boundary between them. This would shift the decision boundary to the right from its current position, restricting approval to projects lying further to the right of that new boundary. Thus, experience directly affects the accurate placement of these boundaries.

For distributed representations, decision boundaries capture the joint effect of aggregation structure and individual cognition. Different aggregation structures employ distinct approval rules, affecting boundary shape. For example, Unanimity, which approves only projects that both specialists deem positive, covers the intersection of the specialists' approval areas, as shown in panel (c). In contrast, Averaging creates a decision boundary that is not parallel to the axes, forming a diagonal boundary with any orientation (panel (d)).[5] While our model analyzes projects described by two cues, the concept of decision boundaries extends to n dimensions. In such cases, a linear boundary corresponds to a hyperplane rather than a line.

The exact placement of the decision boundary of distributed representations depends not only on the aggregation structure used but also on the internal representations used by their members. For example, less experienced specialists would produce less accurate boundaries in panels (c) and (d).

[5]This occurs because the quality predicted by each specialist lies on a plane sloped according to its cue (i.e., $\hat{y}^A = \hat{\beta}_0^A + \hat{\beta}_1^A x_1$ and $\hat{y}^B = \hat{\beta}_0^B + \hat{\beta}_2^B x_2$). Consequently, the quality predicted by Averaging lies on a plane sloped according to both cues (i.e., $\hat{y}^{\text{Avg}} = \frac{\hat{y}^A + \hat{y}^B}{2} = \frac{\hat{\beta}_0^A + \hat{\beta}_0^B}{2} + \frac{\hat{\beta}_1^A}{2} x_1 + \frac{\hat{\beta}_2^B}{2} x_2$). Therefore, the decision boundary of Averaging, where its plane intersects the quality $= 0$ plane, is a straight line that can have any orientation.

Decision boundaries are fundamental in machine learning for understanding how classification algorithms perform. As we demonstrate next, they also help us understand and compare distributed and internal representations under varying task environments.

4.3 The Contingent Nature of the Best-Performing Representation

Armed with the idea of decision boundaries, we can decode Figure 4's complex patterns. To understand which representation performs best, we must consider how effectively each decision boundary separates good from bad projects in the corresponding project space.

Figure 6 overlays twelve representative project spaces on panels (a) and (b) from Figure 4. These project spaces apply equally to panels (c) and (d), which differ only in experience—a property of the representation, not the environment. Each project space shows the distribution of projects at the corresponding arrow's starting point. For instance, panel (a1) shows a project space for $M = -4$, $K = 7$, $D = 1$, and $U = 2$.[6] In this scatterplot, most projects are bad (red dots) due to negative munificence. Good projects (black dots) cluster in the top-right corner, where equal dominance and high interaction yield higher values. The bottom-left corner contains good projects due to the interaction term, which becomes large and positive when both x_1 and x_2 are negative.[7]

4.3.1 Cases where one cue is dominant. We examine cases where one cue is dominant (panels (b) and (d) in Figure 4, where $D = 5$). Both panels use identical parameters except for experience ($E = 64$ versus $E = 4$).

The individual Specialist generally performs best, dominating most areas. This makes intuitive sense: in these environments, projects are unlikely to be good unless they rank high on cue 1, which the Specialist detects. Consider a film studio greenlighting movie projects. If "star power" (cue 1) is overwhelmingly dominant for predicting box office success in a particular genre (e.g., a summer blockbuster), relying on a specialist executive who excels at

[6] Thus, the environment in panel (a1) corresponds to $y = -4 + x_1 + x_2 + 7x_1x_2 + \varepsilon$, where $\varepsilon \sim \text{Normal}(0, 2)$.

[7] The exact position of good and bad projects in the (x_1, x_2)-plane can vary depending on the model parameters and how x_1 and x_2 are coded. For instance, by reverse coding one or both axes, the good projects could fall in any corner of the project space.

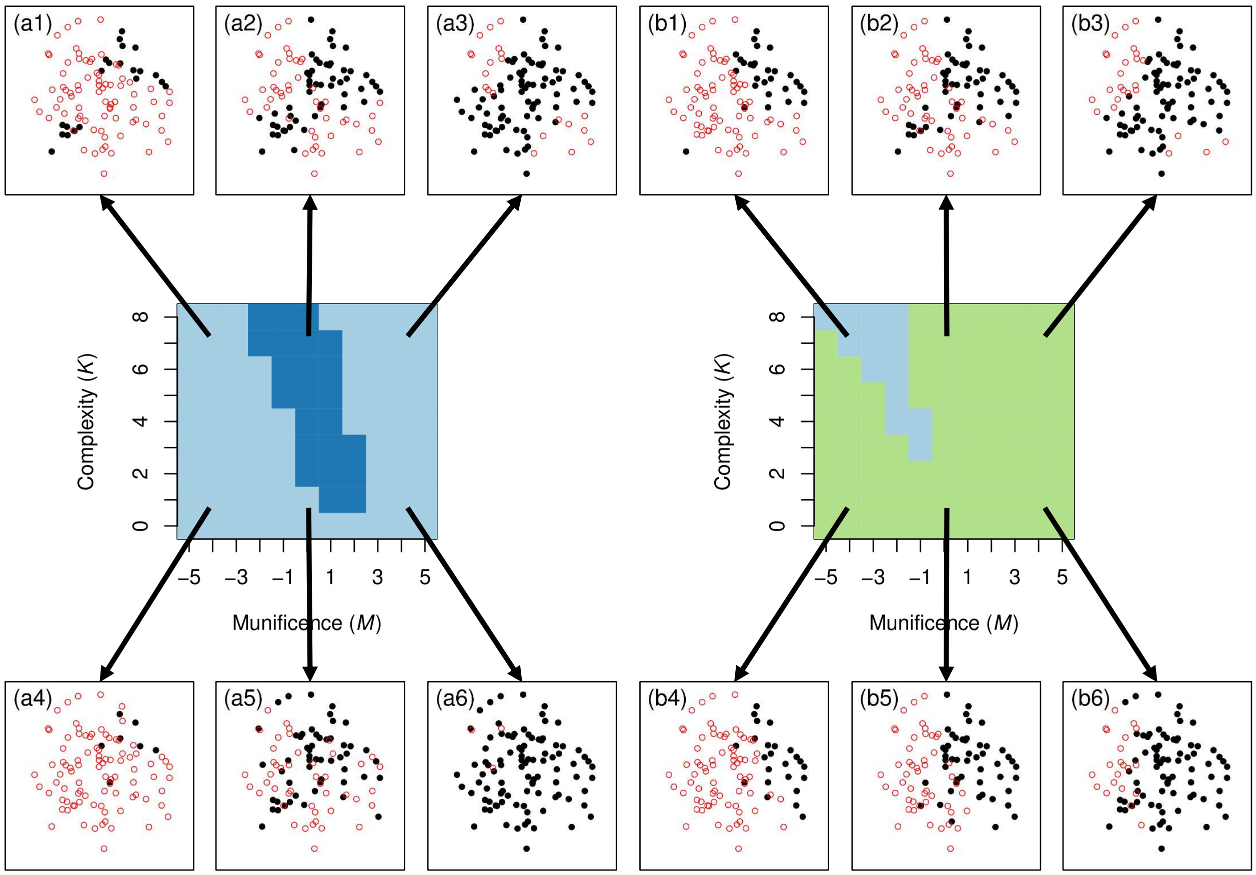

Figure 6: Sample project spaces. Black and red dots represent good ($y > 0$) and bad ($y \leq 0$) projects, respectively.

securing A-list actors might be the best strategy. Their judgment on this single, dominant cue could consistently yield profitable films.

However, the Specialist is not always preferable. The top-left corners of panels (b) and (d) are dominated by Averaging and Unanimity, respectively. That is, when munificence is low and complexity is high, even though one cue is dominant, it is better to use distributed representations that rely on the two specialists. Continuing the film example, if the market is saturated (low munificence) and success also depends on script novelty (which interacts with star power), then simply having a big star isn't enough. In such cases, a Unanimity rule (e.g., both the talent executive and the script development head must approve) or Averaging their assessments might better filter for truly promising projects, even if star power remains a very important factor. The choice between Averaging and Unanimity depends on the specialists' level of experience.

Consider project space (b1) in Figure 6. Good projects cluster in the top-right corner. Averaging or Unanimity can "slice" that corner without approving too many bad projects, while the Specialist would accept too many bad projects.

Whether Averaging or Unanimity is preferable depends on experience. Experienced specialists have accurate boundaries, making Averaging's diagonal boundary effective. Inexperienced specialists have more random boundaries; given negative munificence, Unanimity guards against errors by approving the fewest projects.[8] When the environment isn't munificent and complexity is high, experience becomes critical: lack of experience benefits from Unanimity's restraint; increased experience benefits from Averaging's better-fitting decision boundary. This case also demonstrates that the effectiveness of a decision boundary is determined not only by aggregation structures but also significantly by individual cognition.

This analysis highlights two deeper issues. First, it demonstrates the *power* of distributed representations: even without individuals fully understanding the environment, certain distributed representations performed as if they did. Although specialists' representations contained no interaction terms, Unanimity and Averaging performed well in environments with strong interactions (e.g., project space (b1)). A distributed representation enables an organization to operate *as if* it possessed a richer mental representation than any individual.

Second, it underscores the *perils* of distributed representations. Organizations must choose correctly, as the wrong choice can be costly. In an environment like panel (b3), using Unanimity instead of the Specialist would halve performance. The specialist on cue 2, who poorly understands this cue 1-dominated environment, would split the munificent region that the specialist on cue 1 accurately identifies. While distributed representations can enhance performance, they require careful consideration.

[8]This is because Unanimity approves only when both the specialist on cue 1 *and* the specialist on cue 2 like the project, which is a smaller region than when their average opinion is positive (compare panels (c) and (d) in Figure 5).

4.3.2 Cases where no cue is dominant. We now examine cases where both cues are equally relevant (panels (a) and (c), where $D = 1$). When munificence is high ($M > 2$), Averaging dominates, contrasting with panels (b) and (d) where the Specialist was preferable.

The high-M project spaces are (a3) and (a6) in Figure 6. In space (a3), a diagonal boundary works best, rejecting one of two bad project clusters (either top-left or bottom-right). Neither Unanimity nor the Specialist would be more effective. Averaging's superiority does not depend on experience because even imperfectly angled diagonal boundaries outperform axes-aligned boundaries that would reject many good projects in munificent environments.

For the remaining cases ($M < 2$), panels (a) and (c) differ. In panel (c), when munificence is low, Unanimity is best. The logic mirrors panel (d)'s top-left corner: few good projects exist and inexperienced specialists are inaccurate. The best strategy approves as few projects as possible, which Unanimity approximates by requiring both specialists' approval.

In panel (a), when munificence is low, Averaging is best. In these environments (project spaces (a1) and (a4)), good projects cluster near one corner, which Averaging selects well. Unanimity performs almost as well.[9]

The last unexplained area is medium munificence in panel (a). Consider project space (a2). Although projects appear ordered to us in two dimensions, from the viewpoint of a specialist, whose mental representation collapses the plane into a single dimension, projects look irremediably blended. In other words, the specialist on cue 1 will see that there are good and bad projects at any level of cue 1, and likewise for the specialist on cue 2. This causes specialists to set boundaries randomly. The best approach is conservative: use Unanimity.[10]

[9]The table in Appendix A shows that, in project space (a1), the absolute performance of Averaging is 0.779 versus 0.776 for Specialist.

[10]Averaging is preferable in the lower end of panel (a) (i.e., when $K = 0$), as in this case the project space becomes more intelligible to the specialists: because the interaction term in the environment disappears, projects become more likely to be good the closer they are to the top-right corner of the project space, something that can be selected using Averaging's diagonal decision boundary.

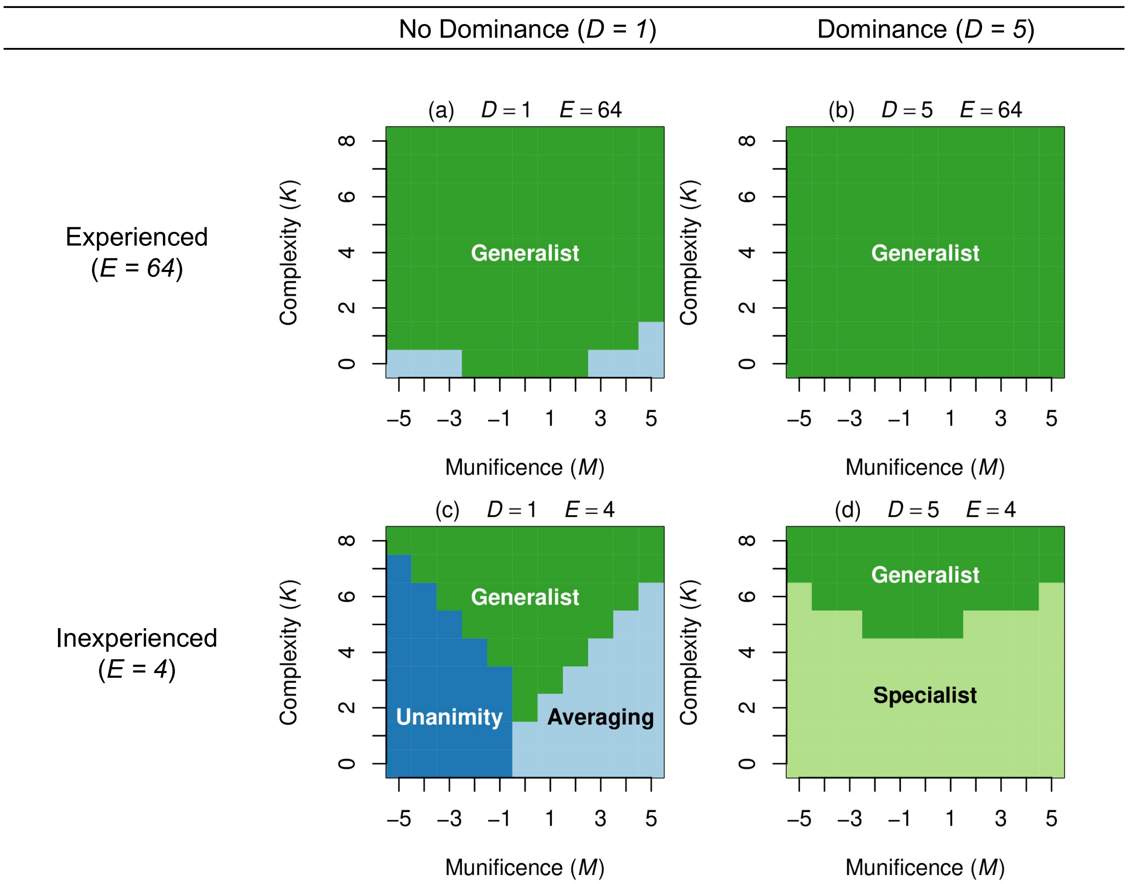

Figure 7: Similar analysis to Figure 4, but including the Generalist.

4.4 Comparing to Generalists

Now that we have a detailed understanding of how aggregation interacts with the internal representations of the specialists, we extend our analysis to include generalists. To facilitate this comparison, Figure 7 presents results similar to those in Figure 4, but adds the Generalist. Recall that the Generalist's internal representation includes both cues and their interaction, which matches the structure of the environment.

When experience is high (panels (a) and (b), where $E = 64$), the Generalist dominates all other representations. This is unsurprising, as with sufficient experience, generalists can estimate accurate coefficients for each cue and the interaction term, enabling a more accurate comprehension of the environment than specialist-based representations could provide. The careful observer may note that the lower corners of panel (a) are not dominated by the

Generalist; this occurs because the interaction term in the Generalist's representation is useless here, and instead, Averaging benefits from the experience of two specialists to more accurately estimate the main effects of x_1 and x_2.

When experience is low (panels (c) and (d), where $E = 4$), the Generalist outperforms others only under specific conditions: roughly speaking, when complexity is high and munificence is not extreme (upper-middle sections of panels (c) and (d)). This makes sense, as generalists are uniquely positioned to see interaction terms, and these regions of the parameter space are where the interaction term is most decisive for performance. When this is not the case, the explanation for why Unanimity, Averaging, and the Specialist dominate was already provided in the context of Figure 4.[11]

In sum, experienced generalists outperform distributed representations (see panels (a) and (b)). However, experienced generalists may be extremely rare, as acquiring deep expertise across multiple dimensions of a problem space requires extensive and varied experience—something few individuals can obtain given the time and attention constraints of organizational career paths. Moreover, unlike in our model, the real world has more than two dimensions, resulting in an exponentially larger number of interactions. The increasing specialization of knowledge and the accelerating pace of environmental change further compound the difficulty of developing broad yet deep expertise within a single individual.

The evolution of McKinsey & Company illustrates this challenge.[12] In the 1950s and 60s, under Marvin Bower's leadership, the firm notably favored generalists, believing their intellect and logic could solve any problem regardless of domain expertise. However, attitudes shifted considerably over time. As the head of McKinsey's German office later observed, presenting oneself as a generalist to clients in the modern era would be met with skepticism: "If you said that to a client today they would think you were in the wrong movie" (McDonald 2014:122). As knowledge grew increasingly complex and specialized, McKinsey correspondingly shifted

[11] An additional reason why the Specialist dominates in the low K part of panel (d) is that it avoids overfitting (Csaszar and Ostler 2020): while the generalist would use the little experience available ($E = 4$) to estimate too many coefficients, the Specialist focuses on just understanding the effect of one factor.

[12] We thank an anonymous reviewer for suggesting this example.

toward domain-specific expertise in the later decades of the twentieth century (McDonald 2014:142–146). Thus, while generalists represent an ideal type, distributed representations may be the more realistic alternative.

4.5 The Generalist Equivalence Point

Figure 7 examines only two levels of experience ($E = 4$ and $E = 64$) that are kept fixed across all representations. However, in the real world, specialists and generalists likely have different levels of experience. Thus, to clarify the relationship between distributed representations and generalists, we introduce the concept of the "equivalent generalist"—the level of experience a generalist needs to match the performance of averaging two highly experienced specialists.

Figure 8 shows this equivalence across a range of environments, varying in complexity (K) and uncertainty (U), while munificence and dominance are kept at neutral values ($M = 0$ and $D = 1$). The figure reveals that in environments with high uncertainty but low complexity, a generalist must be extremely experienced—often requiring eight to sixteen times as many cases as each specialist—to match the performance of Averaging experienced specialists. This happens because, under high uncertainty, the generalist's richer mental model (with more parameters to estimate) becomes harder to estimate reliably, while the value of capturing interactions is limited when complexity is low. Conversely, in environments with either high complexity or low uncertainty, even a relatively inexperienced generalist can match or exceed the performance of distributed representations based on Averaging. In these contexts, the generalist's ability to capture intricate cue interactions quickly becomes advantageous, and less data is needed for effective learning.

These results suggest that distributed representations can effectively substitute for highly experienced generalists in many real-world situations, especially when the environment is noisy but not deeply complex. This is particularly relevant given the scarcity of experienced generalists. As a result, organizations facing uncertain but relatively simple environments may rationally rely on distributed representations rather than seek the elusive "perfect"

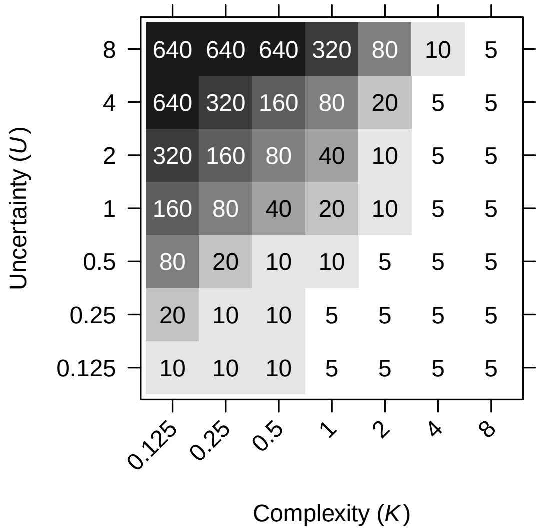

Figure 8: Equivalent Generalist—the level of experience E^G at which a generalist performs as well as Averaging experienced specialists with $E = 640$.

generalist. On the other hand, in settings where the environment is either highly predictable or deeply complex, investing in developing or recruiting generalists—even those with modest experience—can yield superior performance. This analysis helps explain why distributed representations are common in many firms, as well as why generalist-led organizations occasionally dominate in highly complex or stable industries.

5 Discussion

In this study, we developed a theoretical framework to examine how distributed representations influence organizational decision-making, exploring when and why these collective models enhance or hinder organizational performance across varying conditions. Unlike prior work that has treated cognition and aggregation independently, we explicitly modeled their interaction, demonstrating how organizations can transcend individual cognitive limitations by aggregating diverse internal representations. Through a mathematical model, we compared two fundamental distributed representations—Averaging and Unanimity—to the performance of individual Specialists and Generalists. We varied both the environmental factors—munificence, dominance, complexity, and uncertainty—and the experience levels of decision-makers to reveal the contingent nature of distributed representation effectiveness. Our findings highlight that the effectiveness of distributed representations depends critically on how aggregation structures interact with specialists' cognitive limitations, their level of experience, and environmental characteristics. This nuanced interplay underscores the importance of carefully selecting aggregation structures tailored to the organizational context, and demonstrates how even partial knowledge, when effectively integrated, can yield powerful collective decision-making outcomes.

5.1 Summary of Key Findings

Our analysis reveals both the power and limits of distributed representations, clarifying when they significantly enhance organizational decision-making.

Internal representations outperform distributed ones only under specific circumstances. When experienced generalists are available, they deliver the highest performance across all conditions, as shown in panels (a) and (b) of Figure 7. However, such individuals are rare due to the considerable difficulty of acquiring both deep and broad expertise. When only inexperienced managers are available, generalists outperform others when environmental complexity is high and munificence is not extreme (upper parts of panels (c) and (d) in Figure 7), while specialists work better in simpler contexts where one cue dominates (mid to lower part of panel (d) of Figure 7).

In all other situations, distributed representations yield superior performance. This is true when no cue dominates, as shown in panels (a) and (c) of Figure 4. Unanimity works better when munificence and experience are low, because it conservatively filters out low-quality projects, while Averaging serves as an effective default choice in most other circumstances, providing robust performance across diverse conditions.[13] Even in highly complex settings,

[13] The table presented in Appendix A shows that these three rules select the optimal structure in 21 out of

distributed representations remain valuable. Our analysis of the "equivalent generalist" shows that as uncertainty increases, generalists would need unrealistic levels of experience to match the performance of averaging experienced specialists. These findings highlight the practical power and versatility of distributed representations, offering a strong alternative to internal representations across many contexts.

Ultimately, the value of the different representations aligns with George E. P. Box's (1979:202) famous quip: "All models are wrong but some are useful." While none of the representations perfectly captures the real world, each provides useful guidance when chosen wisely. Our findings provide initial guidance on how to make that choice.

5.2 Theoretical Contributions

Our work advances the strategy and organizations literature in four ways. First, we provide a novel causal understanding of the drivers of organizational performance by uncovering a causal structure linking internal representations, distributed representations, and the task environment. The individuals' internal representations and the organization's aggregation structure produce a distributed representation, and its performance depends on the fit between this representation and the task environment. Specifically, we disentangle what would otherwise be a three-way interaction between individual cognition, aggregation structure, and the environment into two independent two-way interactions: one between internal representations and aggregation structure, and another between the distributed representation and the environment. The first interaction determines the organization's distributed representation, while the second interaction determines organizational performance (understood as making good decisions; that is, avoiding omission and commission errors). Figure 9 illustrates this logic. The levers available to the organization designer to improve the fit between the distributed representation and the environment are the constituting elements of the distributed representation: internal representations and aggregation structure. As shown in the Results

the 24 cases presented. In the remaining three cases, following the proposed rules results in a small error vis-à-vis the optimal structure (less than 0.07 in absolute performance).

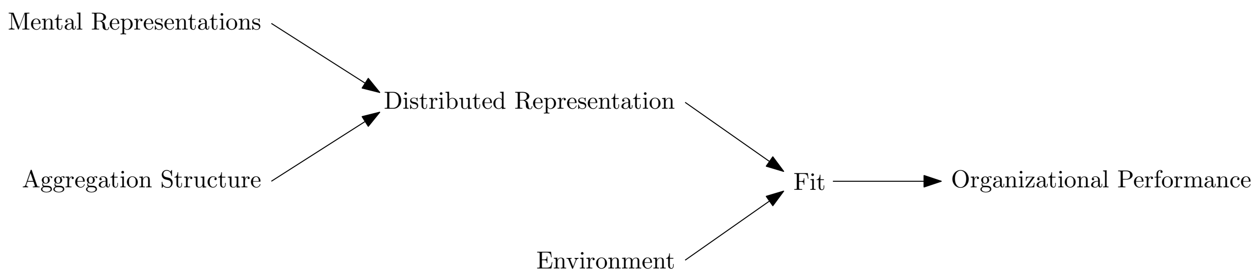

Figure 9: Disentangling how fit emerges from the relationship among internal representations, aggregation structure, and the environment.

section, although no one in the organization may have an ideal internal representation, in many cases it is possible to create a distributed representation that closely approximates that ideal.

Second, we contribute to the research on the microfoundations of organizations. Scholars have sought to understand how micro-level processes underpin organizational performance, particularly how aggregation mediates between individual cognition and organizational outcomes (Gavetti 2005, Foss and Pedersen 2016, Rindova and Martins 2021). Aggregation remains a less studied aspect in microfoundational approaches, often seen as a missing link between individual cognition and organizational performance (Barney and Felin 2013:145). As Finkelstein et al. (2009:115) state, "the gulf between executive characteristics and organizational outcomes is huge." Our study addresses calls for a nuanced understanding of how individuals aggregate their perspectives to generate organizational-level outcomes, clarifying the micro-to-macro process and illuminating the role of distributed representations.

Third, we contribute to the attention-based view (Ocasio 1997), which emphasizes the distribution of attention across individuals in the organization but is less explicit about how the content of attention is processed by individuals and then aggregated by the organization to produce collective outcomes (Ocasio et al. 2018:156–157, Joseph et al. 2024:10–11). By modeling internal representations and aggregation mechanisms explicitly, we show that it is not only the distribution of attention (i.e., what x's each manager sees) but also the internal representations (i.e., the $\hat{\beta}$'s of the different managers) and the aggregation of their perspectives (e.g., Unanimity and Averaging) that determine the organizational decision. Moreover, structures that attend to the same information can perform very differently (e.g., from Figure 3, it is clear that Averaging, Unanimity, and Generalist all attend to the same two cues; yet, as shown in Figure 7, they all perform differently). Thus, a fuller understanding of organizational decision making calls for expanding the focus from attention to distributed representation.

Finally, we advance the Carnegie tradition with two methodological innovations that together expand the analytical toolkit for studying organizational cognition. First, we extend Brunswik's lens model to describe multiple individuals and aggregation structures, providing a flexible language for distributed representations (see, e.g., Figure 3). Second, we introduce the concept of decision boundaries from machine learning to explain how representations fit their environment and impact organizational performance (see, e.g., Figure 5). Just as decision boundaries in machine learning determine which algorithms work best for specific problems (Vapnik 1995, Bishop 2006, Alpaydin 2014), in organizations they offer a criterion for evaluating both internal and distributed representations under different conditions. This concept also embodies Simon's (1990:7) notion of "scissors," demonstrating how performance stems from the interaction between organizational cognition and the environment. Together, these innovations underpin the junction nodes in our three-way interaction diagram (Figure 9) and respond to recent calls to bridge organization theory and AI (Csaszar and Steinberger 2022), thereby opening new avenues for interdisciplinary research.

5.3 A Research Agenda on Distributed Representations

The theory developed here necessarily abstracts away much of the complexity of organizational life in favor of analytical tractability. Our model fixes the number of cues, assumes linear relationships, treats representations as static, and considers only the screening of projects by groups of two. It also remains silent on the social processes through which representations are constructed, contested, and changed. These simplifications allow us to isolate the causal mechanisms linking internal and distributed representations to performance, but they also limit the generality of our conclusions. We therefore conclude by outlining four promising avenues for future research that could extend the frontier of distributed-representation scholarship well beyond the boundaries of this study.

1. Modeling more complex tasks and distributed representations. A natural next step is to relax some of our simplifying assumptions and build models that better reflect the richness of real organizations. Future work could examine screening tasks with more cues, more complex interactions, and interdependencies among projects (i.e., interactions within a portfolio). Models could also incorporate more elaborate distributed representations—for example, larger committees, delegation processes, or multi-layered screening structures where earlier stages shape the pool of projects considered by later ones. Integrating distributed representations with learning dynamics (e.g., the ability of different distributed representations to incorporate feedback) and with search (e.g., which distributed representations are better suited to represent a landscape and discover the high-fitness areas of the it) would further expand the relevance and realism of this line of research.

2. Empirically validating distributed representation theory. Empirical research is needed to test and refine the theory. Laboratory experiments can systematically vary cognitive diversity, environmental complexity, and aggregation rules to probe the boundaries of our findings, while ethnographic studies can uncover the micro-mechanisms by which opinions are aggregated in practice. Large-sample studies might exploit natural experiments—such as staggered changes in committee membership or decision protocols—to assess the impact of distributed representations on outcomes. We believe paying attention to distributed representations could be particularly fruitful in the study of top management teams, where future research could examine the three-way fit among the task environment, internal representations (proxied by managers' backgrounds and traits), and aggregation structures (proxied by leadership styles and decision practices).

3. From selection to generation: distributed creativity. While our study focused on idea selection, distributed representations are also relevant to idea generation. Different aggregation structures may foster divergent thinking by integrating diverse perspectives or, conversely, constrain creativity through consensus requirements or communication bottlenecks. Studying distributed representations in idea generation is inherently more complex, as it involves aggregating open-ended, unstructured ideas rather than simple votes or ratings. Moreover, the space of possible aggregation methods is vast, making traditional experiments with human subjects challenging. However, computational experiments using large language models to proxy for humans (both when generating and integrating ideas) offer a promising approach for exploring this important but difficult terrain.

4. Designing human–AI distributed representations. As AI becomes increasingly integrated into organizational decision-making, designing effective human–AI distributed representations will be increasingly important. Depending on the context, AI systems may function as generalists or specialists (as well as combinations in between), each with distinct strengths. For example, the "Moneyball" case (Lewis 2003) illustrates how data-driven systems can process a multitude of attributes more effectively than human managers, while in other domains, AI excels as a specialist (e.g., detecting anomalies in medical imaging; Anderson et al. 2024) but may miss contextual cues that humans readily perceive. Because AI systems develop fundamentally different internal representations—shaped by training data and algorithms rather than human experience and intuition—much research will be necessary to learn how to best design human–AI distributed representations for different tasks.

These research directions highlight the vast potential of distributed representations as a unifying framework for understanding and improving organizational cognition. By extending our model, empirically testing its predictions, exploring new domains such as creativity, and integrating advances in AI, future work can build a richer, more comprehensive theory of how organizations can harness distributed representations to achieve superior performance.

5.4 Conclusion

This paper has developed a novel perspective that, elaborating on the idea of distributed representations, addresses fundamental issues in organization design and strategic decision-making. From this perspective, individual-level internal representations and organization-level aggregation structures combine to form a distributed representation that determines which projects are approved. High performance is achieved when the distribution of projects in the environment is matched by decision boundaries that are capable of effectively distinguishing the good from the bad projects. By formalizing the concept of distributed representation, our work provides a more rigorous foundation for understanding organizational cognition and improving organizational performance.

Viewing organizations from an information processing perspective launched both the research on managerial cognition and organizational structure over 60 years ago (Joseph and Gaba 2020:267). For the most part, these two lines of research have moved in parallel tracks. The idea of distributed representation offers to unite these twin research traditions. We hope that the integration of these tracks—much like the aggregation processes we studied here—will produce a fuller understanding of organizations.

References

Alpaydin, E. 2014. *Introduction to Machine Learning*. MIT Press, Cambridge, MA.

Anderson, P. G., H. Tarder-Stoll, M. Alpaslan, N. Keathley, D. L. Levin, S. Venkatesh, E. Bartel, S. Sicular, S. Howell, R. V. Lindsey, R. M. Jones. 2024. Deep learning improves physician accuracy in the comprehensive detection of abnormalities on chest X-rays. *Scientific Reports* **14**(1).

Barney, J., T. Felin. 2013. What are microfoundations? *Academy of Management Perspectives* **27**(2) 138–155.

Bishop, C. M. 2006. *Pattern Recognition and Machine Learning*. Springer-Verlag, New York.

Bourgeois, L. J. 1985. Strategic goals, perceived uncertainty, and economic-performance in volatile environments. *Academy of Management Journal* **28**(3) 548–573.

Box, G. E. P. 1979. Robustness in the strategy of scientific model building. R. L. Launer, G. N. Wilkinson, eds., *Robustness in Statistics*. Academic Press, New York, 201–236.

Brunswik, E. 1952. *The Conceptual Framework of Psychology*. University of Chicago Press, Chicago, IL.

Cohen, M. D., J. G. March, J. P. Olsen. 1972. A garbage can model of organizational choice. *Administrative Science Quarterly* **17** 1–25.

Condorcet, J. A. N. 1785/1994. *Foundations of Social Choice and Political Theory*. E. Elgar, Brookfield, VT. Translated and edited by I. McLean and F. Hewitt.

Craik, K. J. W. 1943. *The Nature of Explanation*. Cambridge University Press, Cambridge, UK.

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

Csaszar, F. A. 2018. What makes a decision strategic? Strategic representations. *Strategy Science* **3**(4) 606–619.

Csaszar, F. A., J. P. Eggers. 2013. Organizational decision making: An information aggregation view. *Management Science* **59**(10) 2257–2277.

Csaszar, F. A., N. Hinrichs, M. Heshmati. 2024. External representations in strategic decision making: Understanding strategy's reliance on visuals. *Strategic Management Journal* **45**(11) 2191–2226.

Csaszar, F. A., D. Laureiro-Martínez. 2018. Individual and organizational antecedents of strategic foresight: A representational approach. *Strategy Science* **3**(3) 513–532.

Csaszar, F. A., J. Ostler. 2020. A contingency theory of representational complexity in organizations. *Organization Science* **31**(5) 1198–1219.

Csaszar, F. A., T. Steinberger. 2022. Organizations as artificial intelligences: The use of artificial intelligence analogies in organization theory. *Academy of Management Annals* **16**(1) 1–37.

Dearborn, D. C., H. A. Simon. 1958. Selective perception: A note on the departmental identifications of executives. *Sociometry* **21**(2) 140–144.

Denrell, J., C. Fang, D. A. Levinthal. 2004. From T-mazes to labyrinths: Learning from model-based feedback. *Management Science* **50**(10) 1366–1378.

Dougherty, D. 1992. Interpretive barriers to successful product innovation in large firms. *Organization Science* **3**(2) 179–202.

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

Ethiraj, S. K., D. A. Levinthal. 2009. Hoping for A to Z while rewarding only A: Complex organizations and multiple goals. *Organization Science* **20**(1) 4–21.

Fang, C., J. Lee, M. A. Schilling. 2010. Balancing exploration and exploitation through structural design: The isolation of subgroups and organizational learning. *Organization Science* **21**(3) 625–642.

Finkelstein, S., D. C. Hambrick, A. A. Cannella. 2009. *Strategic Leadership: Theory and Research on Executives, Top Management Teams, and Boards*. Oxford University Press, New York.

Foss, N. J., T. Pedersen. 2016. Microfoundations in strategy research. *Strategic Management Journal* **37**(13) E22–E34.

Gary, M. S., R. E. Wood. 2011. Mental models, decision rules, and performance heterogeneity. *Strategic Management Journal* **32**(6) 569–594.

Gavetti, G. 2005. Cognition and hierarchy: Rethinking the microfoundations of capabilities' development. *Organization Science* **16**(6) 599–617.

Gavetti, G. 2012. Toward a behavioral theory of strategy. *Organization Science* **23**(1) 267–285.

Gavetti, G., D. A. Levinthal. 2000. Looking forward and looking backward: Cognitive and experiential search. *Administrative Science Quarterly* **45**(1) 113–137.

Gavetti, G., D. A. Levinthal, W. Ocasio. 2007. Neo-Carnegie: The Carnegie School's past, present, and reconstructing for the future. *Organization Science* **18**(3) 523–536.

Gavetti, G., D. A. Levinthal, J. W. Rivkin. 2005. Strategy making in novel and complex worlds: The power of analogy. *Strategic Management Journal* **26**(8) 691–712.

Geman, S., E. Bienenstock, R. Doursat. 1992. Neural networks and the bias/variance dilemma. *Neural Computation* **4**(1) 1–58.

Haleblian, J., J. Kim, N. Rajagopalan. 2006. The influence of acquisition experience and performance on acquisition behavior: Evidence from the U.S. commercial banking industry. *Academy of Management Journal* **49**(2) 357–370.

Helfat, C. E., M. A. Peteraf. 2015. Managerial cognitive capabilities and the microfoundations of dynamic capabilities. *Strategic Management Journal* **36**(6) 831–850.

Heshmati, M., F. A. Csaszar. 2024. Learning strategic representations: Exploring the effects of taking a strategy course. *Organization Science* **35**(2) 453–473.

Hogarth, R. M., N. Karelaia. 2007. Heuristic and linear models of judgment: Matching rules and environments. *Psychological Review* **114**(3) 733–758.

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

Hong, L., S. E. Page. 2004. Groups of diverse problem solvers can outperform groups of high-ability problem solvers. *Proceedings of the National Academy of Sciences* **101**(46) 16385–16389.

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

Joseph, J., V. Gaba. 2020. Organizational structure, information processing, and decision-making: A retrospective and road map for research. *Academy of Management Annals* **14**(1) 267–302.

Joseph, J., D. Laureiro-Martínez, A. Nigam, W. Ocasio, C. Rerup. 2024. Research frontiers on the attention-based view of the firm. *Strategic Organization* **22**(1) 6–17.

Joseph, J., A. J. Wilson. 2018. The growth of the firm: An attention-based view. *Strategic Management Journal* **39**(6) 1779–1800.

Kearns, M. J., U. V. Vazirani. 1994. *An Introduction to Computational Learning Theory*. MIT Press, Cambridge, MA.

Kerr, N. L., R. S. Tindale. 2004. Group performance and decision making. *Annual Review of Psychology* **55**(1) 623–655.

Keum, D. D., K. E. See. 2017. The influence of hierarchy on idea generation and selection in the innovation process. *Organization Science* **28**(4) 653–669.

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

Laureiro-Martínez, D., S. Brusoni. 2018. Cognitive flexibility and adaptive decision-making: Evidence from a laboratory study of expert decision makers. *Strategic Management Journal* **39**(4) 1031–1058.

Lee, S., F. A. Csaszar. 2020. Cognitive and structural antecedents of innovation: A large-sample study. *Strategy Science* **5**(2) 71–97.

Lewis, M. M. 2003. *Moneyball: The Art of Winning an Unfair Game*. W. W. Norton, New York.

March, J. G. 2006. Rationality, foolishness, and adaptive intelligence. *Strategic Management Journal* **27**(3) 201–214.

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

Martignoni, D., A. Menon, N. Siggelkow. 2016. Consequences of misspecified mental models: Contrasting effects and the role of cognitive fit. *Strategic Management Journal* **37**(13) 2545–2568.

Masucci, M., S. C. Parker, S. Brusoni, R. Camerani. 2021. How are corporate ventures evaluated and selected? *Technovation* **99** 102–126.

McDonald, D. 2014. *The Firm: The Story of McKinsey*. Simon & Schuster, New York.

Minsky, M. L., S. A. Papert. 1969. *Perceptrons: An Introduction to Computational Geometry*. MIT Press, Cambridge, MA.

Nilsson, N. J. 1965. *Learning Machines: Foundations of Trainable Pattern-Classifying Systems*. McGraw-Hill, New York.

Norman, D. A. 1993. *Things That Make Us Smart*. Addison-Wesley, Reading, MA.

Ocasio, W. 1997. Towards an attention-based view of the firm. *Strategic Management Journal* **18**(Summer Special Issue) 187–206.

Ocasio, W., T. Laamanen, E. Vaara. 2018. Communication and attention dynamics: An attention-based view of strategic change. *Strategic Management Journal* **39**(1) 155–167.

Ocasio, W., L. Rhee, D. Boynton. 2020. March and the pursuit of organizational intelligence: The interplay between procedural rationality and sensible foolishness. *Industrial and Corporate Change* **29**(1) 225–239.

Piezunka, H., V. A. Aggarwal, H. E. Posen. 2022. The aggregation–learning trade-off. *Organization Science* **33**(3) 1094–1115.

Polanyi, M. 1966. *The Tacit Dimension*. Doubleday, Garden City, NY.

Porac, J. F., H. Thomas, F. Wilson, D. Paton, A. Kanfer. 1995. Rivalry and the industry model of Scottish knitwear producers. *Administrative Science Quarterly* **40**(2) 203–227.

Posen, H. E., T. Keil, S. Kim, F. Meissner. 2018. Renewing research on problemistic search: A review and research agenda. *Academy of Management Annals* **12**(1) 208–251.

Puranam, P., B. S. Vanneste. 2016. *Corporate Strategy: Tools for Analysis and Decision-Making*. Cambridge University Press, Cambridge, UK.

Quinlan, J. R. 1993. *C4.5: Programs for Machine Learning*. Morgan Kaufmann, CA.

Rhee, L., W. Ocasio, T.-H. Kim. 2019. Performance feedback in hierarchical business groups: The cross-level effects of cognitive accessibility on R&D search behavior. *Organization Science* **30**(1) 51–69.

Rindova, V. P., L. L. Martins. 2021. Shaping possibilities: A design science approach to developing novel strategies. *Academy of Management Review* **46**(4) 800–822.

Rivkin, J. W., N. Siggelkow. 2003. Balancing search and stability: Interdependencies among elements of organizational design. *Management Science* **49**(3) 290–311.

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

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

Seshadri, S., Z. Shapira. 2003. The flow of ideas and timing of evaluation as determinants of knowledge creation. *Industrial and Corporate Change* **12**(5) 1099–1124.

Sharapov, D., L. Dahlander. 2025. Selection regimes and selection errors. *Organization Science* .

Shepherd, D. A., J. S. McMullen, W. Ocasio. 2017. Is that an opportunity? An attention model of top managers' opportunity beliefs for strategic action. *Strategic Management Journal* **38**(3) 626–644.

Simon, H. A. 1955. A behavioral model of rational choice. *Quarterly Journal of Economics* **69**(1) 99–118.

Simon, H. A. 1990. Invariants of human behavior. *Annual Review of Psychology* **41** 1–19.

Simon, H. A., W. G. Chase. 1973. Skill in chess. *American Scientist* **61**(4) 394–403.

Smith, W. K. 2014. Dynamic decision making: A model of senior leaders managing strategic paradoxes. *Academy of Management Journal* **57**(6) 1592–1623.

Stewart, T. R. 2001. The lens model equation. K. R. Hammond, T. R. Stewart, eds., *The Essential Brunswik: Beginnings, Explications, Applications*, chap. 27. Oxford University Press, New York, 357–362.

Tripsas, M., G. Gavetti. 2000. Capabilities, cognition, and inertia: Evidence from digital imaging. *Strategic Management Journal* **21**(10–11) 1147–1161.

Vapnik, V. N. 1995. *The Nature of Statistical Learning Theory*. Springer, New York.

Vuori, T. O., Q. N. Huy. 2016. Distributed attention and shared emotions in the innovation process: How Nokia lost the smartphone battle. *Administrative Science Quarterly* **61**(1) 9–51.

Walsh, J. P. 1995. Managerial and organizational cognition: Notes from a trip down memory lane. *Organization Science* **6**(3) 280–321.

Winter, S. G., G. Szulanski. 2001. Replication as strategy. *Organization Science* **12**(6) 730–743.

Appendix A: Performance of the Aggregation Structures

Figure 4 only shows which is the best-performing structure as a function of the model's parameters. In contrast, Table A.1 presents the absolute performance of each structure under each of the scenarios detailed in Figure 6. The fact that some of the cells have negative values means that that aggregation structure performs worse than a random decision-maker. For example, when $M = 4$, several cells are negative. This indicates that when the great majority of projects are good, accepting all projects is better than any of the aggregation methods we examine.

(a) $D = 1$ $E = 64$

	M = -4			M = 0			M = 4		
	Spec	Avg	Unan	Spec	Avg	Unan	Spec	Avg	Unan
K = 7	0.758	0.779	0.776	0.067	0.095	0.147	-0.063	-0.041	-0.119
K = 1	0.974	0.974	0.974	0.355	0.500	0.478	-0.019	-0.001	-0.037

(b) $D = 5$ $E = 64$

	M = -4			M = 0			M = 4		
	Spec	Avg	Unan	Spec	Avg	Unan	Spec	Avg	Unan
K = 7	0.801	0.803	0.734	0.649	0.593	0.476	0.359	0.110	0.313
K = 1	0.958	0.916	0.838	0.908	0.907	0.538	0.826	0.452	0.809

(c) $D = 1$ $E = 4$

	M = -4			M = 0			M = 4		
	Spec	Avg	Unan	Spec	Avg	Unan	Spec	Avg	Unan
K = 7	0.625	0.662	0.748	0.023	0.027	0.027	-0.714	-0.558	-1.280
K = 1	0.930	0.956	0.973	0.140	0.201	0.154	-2.910	-1.370	-5.650

(d) $D = 5$ $E = 4$

	M = -4			M = 0			M = 4		
	Spec	Avg	Unan	Spec	Avg	Unan	Spec	Avg	Unan
K = 7	0.698	0.681	0.727	0.385	0.268	0.227	-0.058	-0.193	-0.707
K = 1	0.927	0.870	0.857	0.831	0.629	0.436	0.655	0.210	-0.395

Table A.1: Absolute performance of each structure under the scenarios illustrated in Figure 6.

Appendix B: Extended Analyses

This appendix provides supplementary analyses on how experience, uncertainty, and complexity affect the absolute performance of internal and distributed representations. The parameters of munificence and dominance are held constant, with values set at $M = 0$ and $D = 1$, respectively, as these factors have been thoroughly examined in the main body of the paper and do not yield additional noteworthy findings.

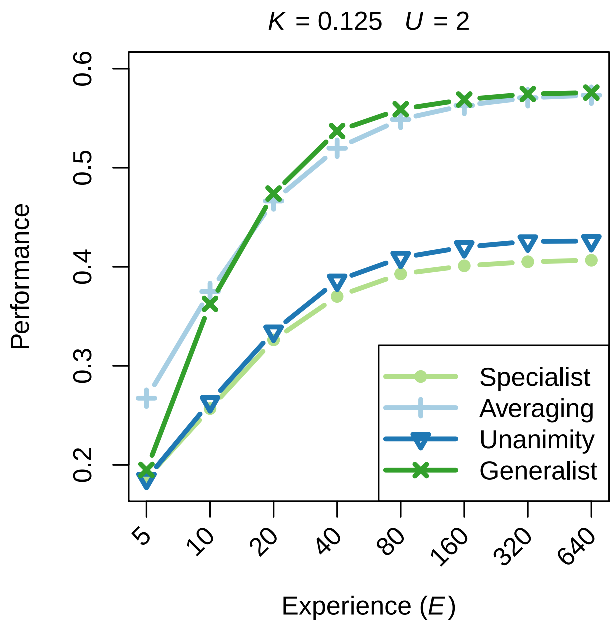

Figure B.1: Performance as a function of experience (E).

B.1 The Effect of Experience

Figure B.1 illustrates the effect of experience on performance in a given environment ($K = 0.125$, $U = 2$). A straightforward observation is that performance increases with experience for all internal and distributed representations. More experience allows individuals to develop more accurate mental representations and thus make better predictions. This finding aligns with empirical and theoretical research documenting that more accurate mental representations lead to fewer decision errors.

Another observation from Figure B.1 is that, with sufficient experience, the Generalist outperforms the Specialist and the two distributed representations. This superior performance stems from the Generalist's internal representation, which mirrors the environment's structure, characterized by two main effects and an interaction. While Averaging approaches the Generalist's performance, it falls short because it does not include an interaction term. In this specific scenario, the interaction term is small ($K = 0.125$), allowing Averaging to closely approximate the Generalist's results.

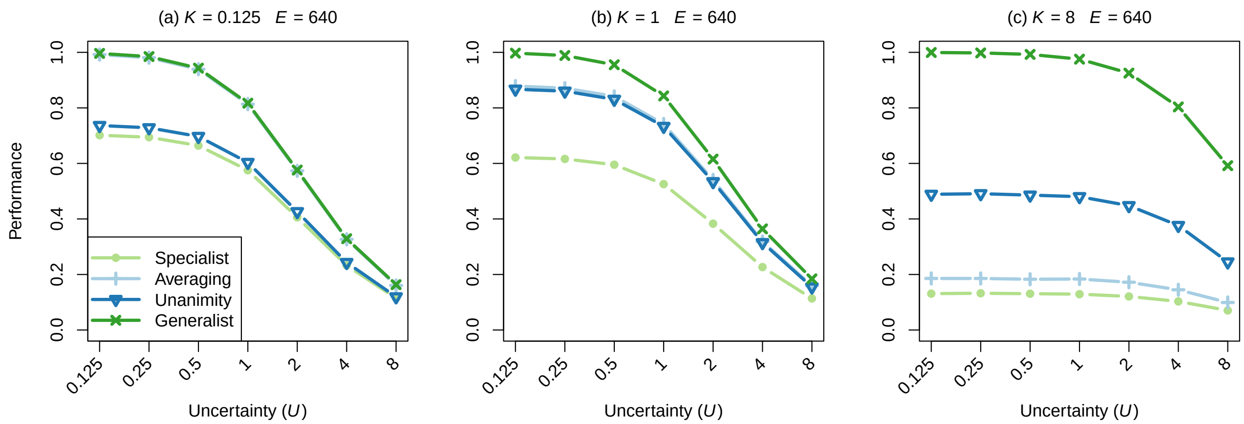

Figure B.2: Performance in different environments under experienced managers ($E = 640$).

B.2 The Effect of Uncertainty and Complexity

We now examine how performance changes as complexity increases. We analyze this relationship under both high and low experience conditions. The panels in Figure B.2 show the effect of uncertainty (U on the x-axes) and complexity (K in each panel) on performance under experienced managers ($E = 640$).

A first observation is that performance decreases as uncertainty increases. This occurs because higher uncertainty makes past observations less predictive of the future, limiting the manager's ability to learn about the environment.

Another observation is that the Generalist continues to be the best performer. When complexity is high (in panel (c)), Specialist and Averaging perform very poorly, as they completely miss the highly relevant interaction. When complexity is low, the Generalist performs similarly to averaging, as the interaction term is less consequential. In short, the value of an experienced Generalist increases as the environment becomes more complex.

Now, we turn to Figure B.3, where managers have low experience ($E = 5$). A striking difference with the previous analyses is that the Generalist is not always the best-performing case: panels (a) and (b) show crossovers between Averaging and the Generalist (occurring around $U = 0.5$ and $U = 2$, respectively). This happens because when experience is low and uncertainty is sufficiently high, the extra complexity of the Generalist's mental representation

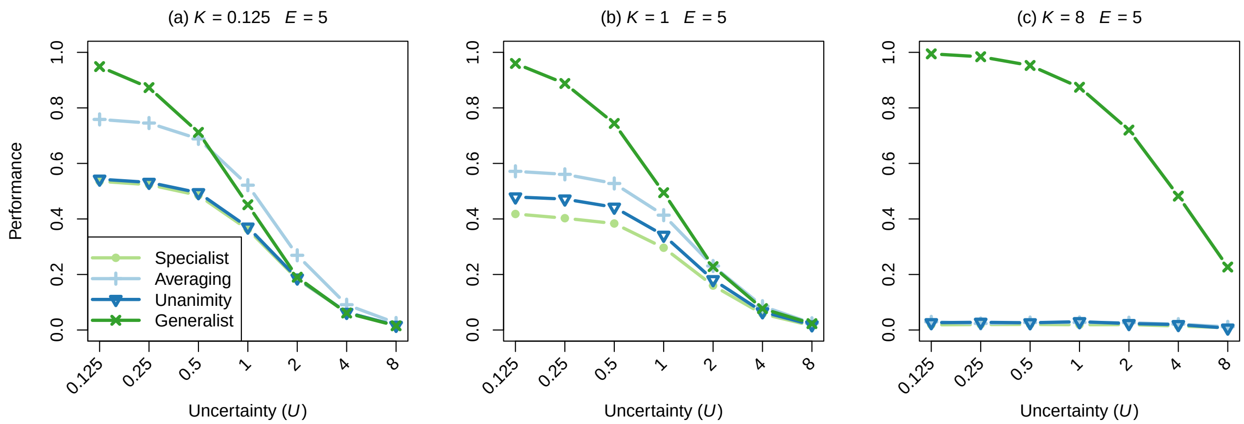

Figure B.3: Performance in different environments under inexperienced managers ($E = 5$).

comes at a cost: the coefficients are estimated with much noise, as very little data per coefficient is available. Moreover, in panels (a) and (b), as K is not high, the upside of including an interaction term in the representation is limited.

Under these conditions, averaging has two advantages over the Generalist. First, Averaging includes two specialists, collectively having twice the experience of the generalist. Second, each individual estimates only two parameters, thus having twice as much data per parameter.

A powerful way to understand when the Generalist is preferable to Averaging is in terms of a bias-variance trade-off (Geman et al. 1992; see also Csaszar and Ostler 2020). The Generalist's model matches the true structure of the environment, achieving lower bias but higher variance. In contrast, Averaging achieves lower variance but higher bias. The combined effect favors Averaging when complexity is not high and uncertainty is high, making the extra data available to the averaging structure particularly valuable.

Appendix C: Code

By modifying the R code below, it is possible to reproduce most plots in the paper. The code below reproduces Figure B.1.

```
library(tidyverse)
library(parallel)
```

```
# * HELPER FUNCTIONS ----------------------------------------------------

# Create E random past experiences in a given environment (M, D, K, U)
# Returns an E x 6 matrix with columns 1, x1, x2, x1x2, epsilon, y
createObs <- function(M, D, K, U, E) tibble(1, x1=rnorm(E), x2=rnorm(E), x1x2=x1*x2, e=
    rnorm(E), y=M+D*x1+x2+K*x1*x2+U*e)

# Runs OLS given Y (vector) and X (matrix)
ols    <- function(Y, X) tryCatch(qr.solve(X,Y), error=function(e) NA)

# Compute mental representations of A, B, and G given observations (Y, X)
mrA    <- function(D) ols(D$y, D[,c('1','x1')])
mrB    <- function(D) ols(D$y, D[,c('1','x2')])
mrG    <- function(D) ols(D$y, D[,c('1','x1', 'x2','x1x2')])

# Compute predictions of A, B, and G given observations (X)
predsA <- function(D, MR) as.matrix(D[,c('1','x1')]) %*% MR
predsB <- function(D, MR) as.matrix(D[,c('1','x2')]) %*% MR
predsG <- function(D, MR) as.matrix(D[,c('1','x1','x2','x1x2')]) %*% MR

# Normalize performance between 0 and 1
stdperf <- function(X, perfect, lazy) (mean(X)-lazy)/(perfect-lazy)

# * MAIN SIMULATION -----------------------------------------------------

# Given E experiences in a given environment,
# have all structures produce N.tests predictions
simObs <- function(M, D, K, U, E, N.tests) {
  # 1. Create past observations for A, B, and G
  D.A        <- createObs(M, D, K, U, E)
  D.B        <- createObs(M, D, K, U, E)
  D.G        <- createObs(M, D, K, U, E)
  # 2. Create test data
  D.test     <- createObs(M, D, K, U, N.tests)
  # 3. Compute mental representations
  mr.A       <- mrA(D.A)
  mr.B       <- mrB(D.B)
  mr.G       <- mrG(D.G)
  # 4. Compute predictions
  pred.A     <- predsA(D.test, mr.A)
  pred.B     <- predsB(D.test, mr.B)
  pred.G     <- predsG(D.test, mr.G)
  # 5. Create results matrix with test data + the predictions
  R          <- tibble(cbind(D.test, pred.A, pred.B, pred.G))
  # 6. + the decisions of the specialists and the aggregation structures
  R$d.A      <- ifelse(R$pred.A > 0, 1, 0)
  R$d.B      <- ifelse(R$pred.B > 0, 1, 0)
  R$d.deleg <- R$d.A
  R$d.avg    <- ifelse((R$pred.A + R$pred.B)/2 > 0, 1, 0)
  R$d.unan   <- ifelse((R$d.A==1) & (R$d.B==1), 1, 0)
  R$d.gen    <- ifelse(R$pred.G > 0, 1, 0)
  # 7. + the accrued performance (y when accepting; 0 otherwise)
  R$p.deleg <- R$y*R$d.deleg
  R$p.avg    <- R$y*R$d.avg
```

```
  R$p.unan  <- R$y*R$d.unan
  R$p.gen   <- R$y*R$d.gen
  # Return the results matrix
  R
}

# Repeat simulation N.runs times in a given environment
# (Hence, the total number of predictions will be N.tests*N.runs)
# Returns aggregate, normalized performance of each structure
simEnv <- function(M, D, K, U, E, N.tests, N.runs) {
  # 1. Repeat main simulation N.run times
  R <- bind_rows(mclapply(1:N.runs, function(i) simObs(M, D, K, U, E, N.tests), mc.cores
      = detectCores()))
  # 2. Compute bounds for normalization
  perfect <- mean(R$y*(R$y > 0))
  lazy    <- mean(R$y)
  # 3. Return normalized performance of the structures
  c(z.deleg = stdperf(R$p.deleg, perfect, lazy),
    z.avg   = stdperf(R$p.avg, perfect, lazy),
    z.unan  = stdperf(R$p.unan, perfect, lazy),
    z.gen   = stdperf(R$p.gen, perfect, lazy))
}

# * RUN THE SIMULATION ------------------------------------------------
Es  <- c(5,10,20,40,80,160,320,640)
RES <- lapply(Es, function(E) simEnv(M=0, D=1, K=0.125, U=2, E, N.tests=1000, N.runs
    =1000)) %>% do.call(rbind,.) %>% as_tibble %>% bind_cols(E=Es,.)

# * REPRODUCE FIGURE B.1 ----------------------------------------------
matplot(RES[,-1],type='l',lwd=2,xaxt='n',xlab='E',ylab='Performance')
axis(1, at=1:length(Es), Es)
legend('bottomright', legend=c('Del','Avg','Unan','Gen'), col=1:length(Es), lty=1:length(
    Es),lwd=2)
```