

**Search Heuristics Under Multiple Objectives:
The Case of Corporate Social Responsibility***

Daniel Albert	Felipe A. Csaszar
LeBow College of Business	Ross School of Business
Drexel University	University of Michigan
Philadelphia, PA 19104	Ann Arbor, MI 48109
daniel.albert@drexel.edu	fcsaszar@umich.edu

June 12, 2026

Forthcoming in *Academy of Management Review*.

*The authors thank Associate Editor Sergio Lazzarini and three anonymous reviewers for their guidance and constructive feedback throughout the review process. The authors are also grateful to Jay Barney, Aneel Karnani, Cha Li, Kate Odziemkowska, Luke Rhee, Chandler Rock, and Adrian Zicari for their insightful comments on earlier drafts. Finally, the authors thank seminar and conference participants at the Academy of Management Annual Meeting, the SMS Annual Meeting, the STR Paper Development Workshop for Mexico, Central America, and the Caribbean, the Troesh Research Conference, the University of Illinois at Urbana–Champaign, and the University of Texas at Austin for their helpful suggestions. All remaining errors are the authors' own.

Search Heuristics Under Multiple Objectives: The Case of Corporate Social Responsibility

Abstract

Organizations increasingly pursue multiple objectives, yet we know little about how boundedly rational firms search for better strategies when performance is multidimensional. We use corporate social responsibility (CSR)—pursuing social and financial outcomes simultaneously—as a motivating context. We distinguish multi-objective *decision making* (choosing among known alternatives) from multi-objective *search* (discovering alternatives through path-dependent local moves), and argue that preferences matter through the heuristics that implement them during search. In a dual-landscape NK simulation, we compare five prominent heuristics: *Maximize* (improve financial performance only), *Combine* (raise the sum of financial and social performance), *Alternate* (pursue one goal until stuck, then switch), *Penalize* (maximize financial performance while deducting shortfalls below a social threshold), and *Satisfice* (prioritize financial performance only after meeting a social threshold). Across varying complexity, goal correlation, and thresholds, these heuristics produce systematically different trajectories and joint outcomes. Importantly, by escaping local financial optima, *Alternate* often discovers “oblique” strategies that match or exceed the financial performance achieved by profit-only local search while also improving social performance. We conclude that implementing the same multi-goal preferences through different search heuristics steers firms toward different regions of the social–financial frontier, shaping both the compromises they reach and the opportunities they discover.

Keywords: behavioral strategy; corporate social responsibility; multiple goals; search; NK model

INTRODUCTION

Searching Under Multiple Goals: The Behavioral Challenge

Strategic decision-making is often presented as a problem of ranking alternatives on a single performance metric—such as return on investment or net present value—and selecting the highest-scoring option. Yet many strategic choices are inherently multidimensional: managers must evaluate alternatives that simultaneously affect cost and quality, speed and safety in product development, mission and margin in public and nonprofit services, and (increasingly) social and financial outcomes in corporate strategy. In such settings, comparing alternatives is not simply a matter of asking whether one scalar is greater than another; it requires comparing vectors. For example, consider two projects: Project *A* yields \$100 million in profits and two units of social value, while Project *B* yields \$200 million but only one unit of social value. Which is better? Unlike scalars, vectors lack a total order (Halmos 1974). Consequently, even defining what it means to move “up” or “down” in performance becomes ambiguous—and this ambiguity is not a peripheral complication but a fundamental behavioral challenge whenever organizations pursue multiple goals.

This challenge is especially salient in corporate social responsibility (CSR), where firms aim to “do well and do good” by improving both corporate financial performance (CFP) and corporate social performance (CSP). These goals are conceptually distinct and, empirically, only weakly correlated on average (Vishwanathan, van Oosterhout, Heugens, Duran, & van Essen 2020). Meta-analyses summarize hundreds of studies reporting correlations between CSP and CFP that range from modestly positive to strongly negative depending on context (Margolis & Walsh 2003, Orlitzky, Schmidt, & Rynes 2003, Vishwanathan et al. 2020). In some industries, technological and contracting conditions produce systematic trade-offs (e.g., cost–quality or environmental trade-offs), making it difficult to raise one dimension without lowering the other (Hart, Shleifer, & Vishny 1997). In other settings, stakeholder reactions and shared-value opportunities can create positive associations between CSP and CFP (Dorobantu, Henisz, & Nartey 2017, Flammer 2015). Thus, the relevant strategic problem is not merely whether CSP and CFP are aligned or in conflict, but how organizations pursue improvements in environments where the relationship between objectives varies.

Most research addressing such multidimensionality has focused on *decision-making*: given a set of known alternatives, how should actors evaluate and select among them under multiple criteria? The multi-objective decision-making literature formalizes this problem using concepts such as domination and the efficient frontier (Ehrgott 2005, Emmerich & Deutz 2018, Keeney & Raiffa 1976). However, this one-shot framing treats the set of alternatives as given. By contrast, the behavioral theory of the firm emphasizes *search*: under bounded rationality, alternatives are not given but must be discovered over time through local, iterative exploration (Cyert & March 1963, Levinthal 1997, Simon 1955). Search is therefore logically prior to choice and path dependent: the alternative adopted today becomes the starting point for discovering what is feasible tomorrow. This distinction is consequential in a multi-goal setting because the rules used to implement preferences are not only decision criteria applied to known options; they are *search heuristics* that steer which options the firm ever discovers. Yet we still lack a clear understanding of how different multi-goal search heuristics shape search trajectories and, ultimately, the joint social and financial outcomes organizations achieve.

Our approach and contribution

We address this gap by studying search under multiple objectives. Using CSR as a motivating and empirically rich context, we ask: *How do different search heuristics affect the social and financial outcomes organizations achieve when performance is multidimensional?* To answer this question, we model organizations as boundedly rational searchers on rugged performance landscapes and examine how different rules for evaluating nearby alternatives generate different search trajectories and, consequently, different outcomes. We focus on five heuristics drawn from the CSR and multiple-goals literatures. *Maximize* represents a profit-first rule that ignores social performance (Friedman 1970, Jensen 2002). *Combine* aggregates social and financial outcomes into a single composite metric, reflecting additive or utilitarian-like evaluation (Bentham 1789/1996, Mill 1848/2004). *Alternate* sequentially shifts attention between objectives, consistent with the behavioral argument that organizations often attend to one goal at a time (Cyert & March 1963, Greve 2008). *Penalize* maximizes financial performance while imposing explicit penalties for falling

below a social threshold (McWilliams & Siegel 2001). *Satisfice* prioritizes financial performance only after meeting a minimum social aspiration, reflecting problemistic search and threshold-based evaluation (Cyert & March 1963, March & Simon 1958/1993). These heuristics are not merely different preference profiles; some encode similar preferences but implement them differently (e.g., Combine and Alternate can both reflect joint concern for CSP and CFP), allowing us to isolate how *implementing* preferences via different search heuristics changes what firms discover and achieve.

We implement these heuristics in an NK simulation model extended to dual performance landscapes (Adner, Csaszar, & Zemsky 2014, Levinthal 1997). Each strategy is a vector of interdependent decisions that jointly determines both financial and social performance, and we vary the correlation between the two performance landscapes to represent environments with trade-offs, independence, or complementarities. This approach lets us examine how the relative effectiveness of heuristics depends on both internal search processes and external performance structure—a multi-goal instantiation of Simon’s scissors (Simon 1990).

Our analysis makes four contributions. First, we advance behavioral strategy by showing that, under multiple objectives, the heuristics that implement preferences are consequential not only for *choice* among known alternatives but also for *search* itself: they shape which alternatives organizations discover through path-dependent exploration, clarifying why the search/decision-making distinction matters in multi-goal settings (Cyert & March 1963, Keeney & Raiffa 1976). Second, we provide a systematic comparative analysis of five prominent multi-goal heuristics—Maximize, Combine, Alternate, Penalize, and Satisfice—and show how their relative effectiveness depends on the structure of the environment (e.g., whether social and financial outcomes are in tension or mutually reinforcing). Third, we uncover an important implication for CSR: alternating attention between social and financial goals can, in many conditions, improve social outcomes *without sacrificing* financial performance—and can even lead firms to *higher* financial performance than a narrow focus on profit maximization. By shifting between objectives, firms can escape a local financial optimum at which a profit-maximizing firm settles, and thereby potentially discover strategies that perform better on both dimensions. This is consistent with the logic of “oblique

strategies” (Eno & Schmidt 2001, Kay 2010), that is, the notion that high-performing solutions are discovered through detours rather than a persistent focus. This finding also challenges strong versions of the shareholder-primacy argument (Friedman 1970, Jensen 2002). Fourth, we distill the modeled heuristics into a parsimonious “grammar” for describing and contrasting multi-objective search heuristics, providing a template that can be extended to other multi-goal contexts beyond CSR.

The remainder of the paper is organized as follows. Section 2 develops the theoretical foundation by (i) formalizing the challenge of multidimensional performance, (ii) summarizing multi-objective decision-making concepts, (iii) shifting to the behavioral logic of multi-objective search, and (iv) motivating the five heuristics. Section 3 presents the dual-landscape NK model and our simulation design. Section 4 reports comparative results across environments, including conditions under which multi-objective heuristics match or outperform Maximize. Section 5 discusses implications for behavioral strategy and CSR, including the practical relevance of heuristic choice when organizations face multiple objectives.

THEORETICAL BACKGROUND

This section establishes the theoretical foundation for our analysis of how firms search for strategies that satisfy multiple objectives. We focus on the specific, yet broadly relevant, context of corporate social responsibility (CSR), where firms face the dual imperatives of financial and social performance. We proceed in four parts. First, we frame the core challenge: the inherent ambiguity of comparing strategic alternatives when performance is multidimensional. Second, we introduce key ideas from the literature on multi-objective decision making, which provides a formal framework for comparing and selecting among alternatives when objectives are multidimensional. Third, we discuss the role of multi-objective search within the behavioral theory of the firm, emphasizing the generation and discovery of alternative strategies under bounded rationality. Finally, we distill five search heuristics from the CSR and multiple-goals literatures, forming the basis for our subsequent model.

The Challenge of Multidimensional Performance

A fundamental difficulty in managing for multiple objectives stems from the fact that performance is not a single, scalar value, but rather a vector with multiple dimensions (consistent with Miller et al.'s 2013 view of performance as a domain of “separate constructs”). Multiple objectives reflect organizational members' (and potentially external stakeholders') diverse, and sometimes conflicting, interests, requiring organizations to consider multiple performance dimensions simultaneously—financial, social, environmental, and others. This multidimensionality is central to classic and contemporary work on multiple goals in organizations (Cyert & March 1963, Gaba & Greve 2019, Greve 2008, Greve & Gaba 2020), and it is increasingly salient in strategy research emphasizing that firms must simultaneously manage multiple performance dimensions (e.g., mission vs. margin, cost–quality trade-offs, and other persistent tensions) (Levinthal & March 1993, McGahan 2023, Porter 1996, Shaked & Sutton 1982).

A straightforward numerical example illustrates the challenge of comparing alternatives when performance is multidimensional. If projects *A* and *B* have net present values of, say, \$100 and \$200 million, it is clear that project *B* is preferable. That is, finding the best alternative is straightforward: one simply picks the alternative that ranks the highest according to the performance dimension. But if performance is multidimensional, picking a superior alternative ceases to be straightforward. Consider the case where project *A* produces \$100 million and 2 units of social good versus project *B*, which produces \$200 million and 1 unit of social good. The answer to “which of the two projects should a firm choose?” is no longer self-evident. Picking among the alternatives now requires comparing two vectors— $(\$100, 2)$ and $(\$200, 1)$ —and vectors, unlike one-dimensional quantities, do not have a well-defined “greater than” operator, a concept mathematicians refer to as the absence of a “total order” (Halmos 1974, Section 14). In other words, how to compare different alternatives whose performance is multidimensional is not generally defined.

This is problematic, as comparing and choosing among alternatives whose performance is multidimensional is at the core of strategic decision-making. For instance, the search literature suggests that firms compare nearby alternatives to the status quo and adopt a new strategy if it

moves the firm “up” along the performance dimension of interest (Levinthal 1997; for a survey, see Baumann, Schmidt, and Stieglitz 2019). However, as seen with the previous numerical example, there is no “up” or “down” when comparing multidimensional performance outcomes. This is not exclusive to the behavioral understanding of firms as searching on a landscape; it also applies to the economic understanding of firms as optimizers, as optimizing also relies on comparing alternatives.

Because of this, economists have advised against even trying to optimize along multiple objectives. Jensen (2002: 238) argues that “[i]t is logically impossible to maximize in more than one dimension at the same time unless the dimensions are monotone transformations of one another” as an additional reason why “shareholder primacy” (Friedman 1970)—just focusing on improving financial performance—is the appropriate course of action for firms.

The CSR literature provides a compelling, real-world illustration of this multidimensionality through its focus on two key performance metrics: corporate financial performance (CFP) and corporate social performance (CSP) (Kaul & Luo 2018: 1655).¹ Numerous studies have examined whether social and financial performance are positively related by measuring the correlation between these metrics. This extensive body of research has been summarized by several meta-analyses (Allouche & Laroche 2005, Gallardo-Vázquez, Barroso-Méndez, Pajuelo-Moreno, & Sánchez-Meca 2019, Margolis & Walsh 2003, Orlitzky et al. 2003, Vishwanathan et al. 2020, Wang, Dou, & Jia 2016). The main finding is that, on average, there is a small positive correlation between social and financial performance. The most recent and comprehensive meta-analysis (Vishwanathan et al. 2020), covering 344 individual studies, finds a mean correlation of 0.07 (with the 95% confidence interval going from -0.14 to 0.29) and reveals substantial variation across industries (e.g., a correlation of 0.08 in manufacturing and 0.23 in financial services).

The correlation between CFP and CSP is shaped by factors—including technology, consumer preferences, regulations, and stakeholder dynamics—that vary across industries. In some settings, that correlation is negative, reflecting a trade-off between financial and social goals. For example,

¹We acknowledge that treating CSP as a single dimension is a simplification. A true multi-stakeholder perspective would imply multiple social dimensions—perhaps one for each stakeholder group. However, we adopt this two-dimensional framing for tractability while preserving generality. Moreover, in many CSR settings, social performance can be relatively well-defined (e.g., CO₂ emissions for airlines), making this simplification reasonable for our purposes.

Hart et al. (1997) argue that when “quality” (aligned with social outcomes) is hard to contract for, firms face a cost–quality trade-off: cutting costs to boost financial performance can undermine quality, leading to negative CFP–CSP correlations, especially in sectors where safety, environmental protection, or inclusiveness are costly and hard to monitor.

Conversely, CFP and CSP may be positively correlated when firms identify win-win opportunities that create shared value. Flammer (2015) finds support for a positive relationship between CSR policies and financial performance. In her study of large companies that voted on CSR policy-related shareholder proposals and narrowly adopted them, these companies achieved greater financial performance afterward compared to firms that narrowly rejected these proposals. Similarly, Dorobantu et al. (2017) also support a positive relationship between stakeholder reactions and financial performance. In their study of critical events in small gold mines, those associated with more positive stakeholder reactions in response to critical events showed greater financial performance, while more negative stakeholder reactions led to lower performance.

In other sectors, the CFP–CSP correlation is weak because the activities driving social and financial performance may operate largely independently. For example, a bank’s financial performance depends primarily on its core lending and investment activities, while its social performance often stems from separate philanthropic initiatives with minimal interaction between the two domains. Additional factors contributing to low CFP–CSP correlation include firms facing an evolving array of stakeholder demands (McGahan 2023) and significant variability in how organizations measure CSP (Kroeger & Weber 2014, Rawhouser, Cummings, & Newbert 2017).

Multi-Objective Decision-Making

While the traditional understanding of maximization does not directly apply to multiple objectives, the literature on multi-objective optimization offers a formal framework for conceptualizing what it means to “maximize” performance across multiple dimensions (Ehrgott 2005, Emmerich & Deutz 2018, Keeney & Raiffa 1976). This framework provides important distinctions that allow for a more precise and formal language to describe and analyze multidimensional performance. By introducing concepts such as domination and the efficient frontier, it gives us the tools needed

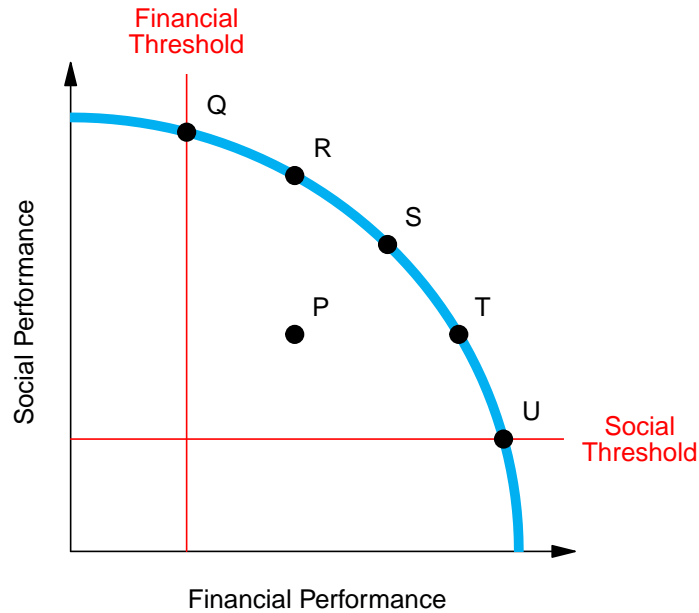


Figure 1: The social–financial frontier.

to rigorously compare alternatives across multiple dimensions. These distinctions serve as the analytical foundation for the search heuristics and modeling approach developed in the remainder of the paper.

A foundational idea in this literature is representing performance as a point in a multidimensional space. For CSR, with its dual focus on social and financial outcomes, a firm’s strategy corresponds to a point on a two-dimensional plane. The set of all feasible strategies forms a region in this space, as depicted in Figure 1.²

In the plane of possible strategies, some positions are objectively inferior to others, as they perform worse in both dimensions (a relationship called *strict domination*) or perform worse in one dimension while being at parity in the other (*weak domination*). For instance, in Figure 1, *S* strictly dominates *P*, while both *R* and *T* weakly dominate *P*. Some points are not dominated by any other points (in the example, all points except *P* are not dominated by any other point). The set of all non-dominated points is called the *efficient frontier* (Keeney & Raiffa 1976: 70).

Because firms often lack complete knowledge of the frontier, they strive for an “ideal point”

²If the space of strategies is not continuous, instead of a “region” it would be a “discrete set of points.” We use the former term for succinctness.

reflecting their preferences for social and financial outcomes. Due to technological and practical limitations, the best achievable outcome is the “best viable point”—the point on the frontier closest to the ideal. Firms’ locations on the frontier can represent a range of strategic positions. For instance, a solely financially focused firm might aim for point U , a socially focused foundation for point Q , and firms balancing diverse stakeholders somewhere in between, like point S . Firms are limited in where they locate by financial and social thresholds. Below the financial threshold, a firm will go bankrupt, and below the social threshold, the firm is not meeting its social goals.

The correlation between performance dimensions, which meta-analyses have shown to vary substantially across industries (recall the 95% confidence interval of correlation going from -0.14 to 0.29), significantly influences the shape and size of the feasible region, and consequently, the efficient frontier (Adner et al. 2014); see Figure 2 for an illustration. For instance, in industries where the meta-analyses show negative correlations between CSP and CFP, such as certain manufacturing sectors, reducing environmental impact may often come at a financial cost, resulting in a backward-sloping set of positions and a larger frontier (see panel (a)). Conversely, in industries like financial services where meta-analyses reveal stronger positive correlations, improvements in social and financial performance tend to be mutually reinforcing, leading to a relatively smaller frontier (see panel (c)). Industries with correlations close to zero, which align with the overall mean correlation of 0.07 found in the literature, might allow for a wider range of social–financial combinations (see panel (b)).

Correlation and trade-offs are inversely related: as correlation increases, the prevalence of trade-offs diminishes. Trade-offs arise when a firm cannot improve along one dimension without sacrificing performance along another. This situation occurs for firms that are *on the frontier* (the highlighted arcs in Figure 2). As correlation increases—that is, as one moves from panel (a) to panel (c)—the size of the frontier shrinks. Consequently, the number of firms facing trade-offs (those on the frontier) decreases as correlation increases.




Figure 2: The effect of correlation on the efficient frontier. Panels (a)–(c) show negative, zero, and positive correlations, respectively, between firms’ social and financial performance; black dots represent individual firms. The arcs represent stylized efficient frontiers.

Multi-Objective Search

The previous subsection established how to conceptualize “optimizing” with multiple objectives for a given set of alternatives. However, it does not address the *process* by which firms might discover high-performing (or frontier) points in the first place. The behavioral theory of the firm, grounded in the concept of bounded rationality (Simon 1955), is critical here. Firms face significant cognitive and informational limitations; they cannot possibly know all potential strategies and their associated outcomes across all relevant performance dimensions. Instead of comprehensively evaluating all possibilities, firms engage in *search* (Cyert & March 1963, Levinthal 1997): they must generate and discover alternative strategies before they can compare these to the status quo. Decision-making, by contrast, is the selection among currently available alternatives. Search is thus logically prior to decision-making, and it is path dependent: once a new strategy is adopted, it becomes the status quo that shapes which neighboring alternatives the organization is likely to discover next.

This distinction matters especially under multiple goals. With multiple performance dimensions and potential trade-offs, preferences do not only determine which of the *known* alternatives a decision maker selects; they also shape how search proceeds. Put differently, preferences must be *implemented* in some procedure for evaluating candidate moves during search, and that procedure influences which regions of the performance space the firm ever reaches. In practice, search and choice are not neatly sequential: each step of exploration creates familiarity and commitment that

shape subsequent evaluation (Gans, Stern, & Wu 2019), so the evaluation rule is not applied once to a fixed menu but repeatedly, steering the trajectory of discovery itself.

The classical NK landscape framework. The NK landscape framework (Levinthal 1997) offers a powerful way to model how firms evaluate and ultimately search for better strategies. Each point on the landscape represents a specific combination of decisions, with its “height” indicating performance.

To illustrate, consider a car manufacturer facing three key decisions: (i) engine type (electric or combustion), (ii) car type (SUV or pickup), and (iii) market (US or global). Each of the eight possible combinations of these choices yields a different outcome in terms of performance. For example, selling combustion SUVs globally might maximize profits, and selling electric pickups in the US might be less profitable.

The parameter N in the NK model represents the number of such strategic decisions the firm makes. The parameter K captures how much these decisions interact: for instance, the profitability of selling SUVs may depend on whether the target market is the US or global, but not on whether the cars are electric or not. When K is high, these decisions are highly interdependent—changing one aspect (e.g., switching to electric engines) can have ripple effects on performance depending on the other choices made. This interconnectedness creates a “rugged landscape,” where small changes can lead to unpredictable jumps in performance.

Due to bounded rationality, managers cannot foresee the consequences of all possible combinations but can experiment with small changes to observe their effects. Thus, firms explore a rugged landscape by using local search—making incremental changes to their decisions and moving to a neighboring strategy only if it improves performance according to their chosen heuristic.

The effect of K can be visualized as increasing the ruggedness of a three-dimensional landscape (see Figure 3). Although this topographical metaphor is not fully accurate, it effectively illustrates a key insight from this literature: as complexity grows, boundedly rational actors—those who can only search their immediate neighborhood—find it increasingly difficult to reach the global peak. For readers interested in a deeper understanding of NK models, we recommend Csaszar (2018),




Figure 3: Illustration of landscape complexity (K) in a unidimensional setting. Higher K leads to a more rugged landscape with multiple local peaks.

which provides an accessible introduction to this rich literature.

A key result of the literature on rugged landscapes is that firm performance depends on the interaction between the search heuristic used by the firm and the characteristics of the landscape. This interplay embodies Simon’s (1990) “scissors”: the search outcome depends on the interaction between the task environment and the decision maker.

However, the role that search plays under multiple goals has received limited attention in the search literature, and the NK literature in particular, despite strong reasons to believe it matters.³ This is surprising, given that the management of multiple goals in organizations constitutes a central pillar of the behavioral theory of the firm (Cyert & March 1963).

Empirical research suggests that firms differ in how they prioritize and manage such trade-offs—and that these differences have significant consequences for effective search. For example, Greve (2008) shows that insurance firms balance size and profitability goals, with firms below their size targets growing faster—especially when profitability goals are met. Gaba and Greve (2019) find that airlines prioritize safety and profitability based on survival needs: firms with low profitability focus more on safety, while highly profitable firms are less responsive to safety concerns. Albert and Eklund (2026) show that organizational subgroups representing distinct goals—revenues versus patient care in hospitals—may pursue the same objectives differently depending on their hierarchical position. Given such variation in how organizations manage multiple goals, understanding how they

³An important exception is Ethiraj and Levinthal (2004), who model search under multiple goals as a constraint: any change must not decrease performance on other goals. Their work demonstrates that strict improvement requirements can limit adaptability, but focuses less on how different approaches to evaluating multiple goals shape search trajectories.

search under multiple goals—and whether different approaches lead to different outcomes—becomes an important question.

Extending to two performance landscapes (dual landscape model). To address the single-goal limitation of the NK literature, we build on the dual landscape framework, which allows each strategy to map to multiple performance values.

Returning to our car manufacturer example: each combination of decisions (engine type, car type, market) not only determines a single performance value, but actually maps to *two* performance values: one financial (e.g., profit) and one social (e.g., environmental impact). For instance, producing electric SUVs for the global market may yield high social performance (due to low emissions) but moderate financial performance if costs are high or consumer demand is uncertain. Conversely, combustion pickups in the US market may yield high financial returns but poor social outcomes.

This dual mapping is captured in the *dual landscape* framework introduced by Adner et al. (2014), in which each strategy is associated with two separate performance values—one on a financial landscape and one on a social landscape (see Figure 4). While Adner et al. (2014) originally developed the dual landscape framework to examine industrial organization questions, it provides an ideal foundation for our analysis of how organizations search under multiple objectives. The degree to which these two landscapes (or objectives) are aligned or in conflict is governed by the correlation parameter (ρ): in some industries, strategies that increase financial performance also tend to improve social outcomes (positive correlation), while in others, the opposite may happen (negative correlation).

Figure 4 illustrates how a single strategy—a particular combination of decisions—maps to both financial and social performance. Importantly, there is only *one* underlying set of decisions: the same interdependent decision vector \mathbf{d} jointly determines both outcomes, so changing any component of \mathbf{d} can have coupled consequences for financial and social performance. When the two landscapes are highly correlated, strategies that do well financially also tend to do well socially. When the correlation is negative, improving one dimension often comes at the expense of the other, making it difficult to find strategies that perform well on both.

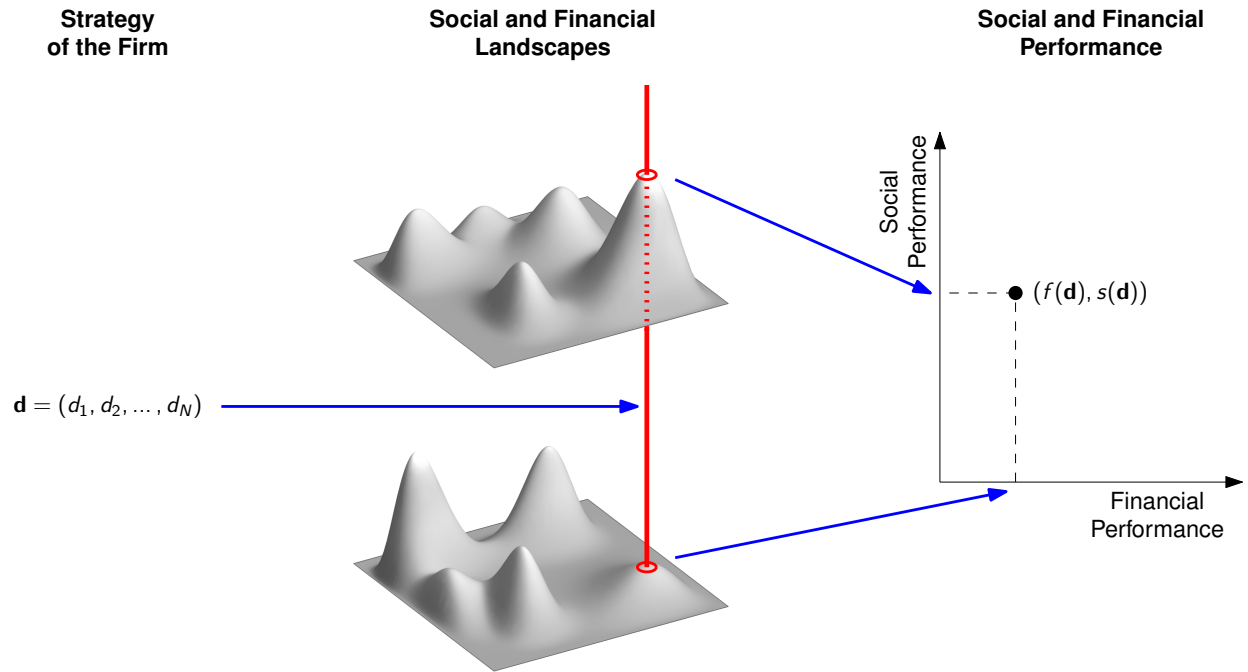


Figure 4: A dual-performance landscape. Each strategy (represented by the vector of decisions \mathbf{d}) has a corresponding performance on the financial and social performance landscapes. The point $(f(\mathbf{d}), s(\mathbf{d}))$ represents the multidimensional performance of strategy \mathbf{d} . The correlation between landscapes varies by industry and context.

Search Heuristics for Social and Financial Performance

To understand search in a multidimensional setting, we must consider how firms generate and select alternatives. The behavioral theory of the firm suggests that search is often local and iterative, informed by a firm’s existing knowledge (e.g., Fleming 2001, Katila & Ahuja 2002, Rosenkopf & Nerkar 2001), primarily promoting incremental change that is implemented if decision makers consider it superior to the status quo. That is, firms generally search for improvements to their current strategy rather than for entirely different ones. When evaluating such alternatives in a one-dimensional performance context, the heuristic applied is one of “greater than” (up or down). However, because “up” cannot be unambiguously defined in a multidimensional setting, other heuristics must take the place of the unidimensional comparison operator.

Given our focus on how organizations balance multiple stakeholder interests, and recognizing that CSR decisions (involving both social and financial performance) represent a canonical example of this challenge, we draw on the CSR and multiple-goals literatures to identify relevant heuristics.

To understand search in a multidimensional setting, it is important to distinguish between preferences and heuristics. *Preferences* represent a firm’s desired position in the performance space—for instance, aiming for point *S* in Figure 1 would reflect equal weighting of social and financial outcomes. However, preferences alone do not determine how firms move toward their goals. That process is governed by *heuristics*—the specific rules and procedures firms use to evaluate strategies that align with their preferences. Multiple heuristics can serve the same preference structure. For example, a firm equally valuing social and financial performance could employ either a Combine heuristic that evaluates both dimensions simultaneously or an Alternate heuristic that switches focus between them, as both heuristics attend to both dimensions.⁴ This distinction matters because, as we will demonstrate, different heuristics vary significantly in their effectiveness at achieving the same preferences: the same “ideal point” can be pursued through different search procedures that generate different search trajectories and, consequently, different outcomes.

Drawing from the CSR and multiple-goals literatures discussed above, we identify five distinct heuristics that capture the main ways organizations approach the dual pursuit of social and financial performance. Two of these are *compensatory* in the sense that they explicitly trade off the two dimensions on a single evaluative scale (Combine and Penalize), which typically requires commensurability of social and financial performance (Carruthers & Stinchcombe 1999, Espeland & Stevens 1998).⁵ Two are *non-compensatory* in that they avoid aggregation by focusing attention on one dimension at a time (Alternate and Satisfice), thereby reducing commensurability demands to within-dimension comparisons and/or threshold assessments. Finally, one represents a profit-first baseline that disregards social performance (Maximize). To clarify the logic of each heuristic, we first describe the rationale behind each approach—the core decision-making principle guiding search—before formally modeling them in the next section.

1. *Maximize financial performance.* Under this rationale, decision-makers ignore social performance and focus entirely on trying to maximize financial performance. Consequently, this

⁴The selection between these heuristics may also be influenced by other factors, such as a firm’s preference for stable choices or the cost of implementing a given heuristic.

⁵We thank an anonymous reviewer for encouraging us to think about the role and implications of commensurability in our model.

rationale reflects a preference structure that puts zero weight on social performance and all the weight on financial performance. This view is closest to Friedman's (1970) doctrine that "the social responsibility of business is to increase its profits." Although few firms may claim to completely disregard social performance outcomes, the profit-maximizing perspective has historically been a "dominant logic" (Crilly & Sloan 2012, Prahalad & Bettis 1986) shaping corporate behavior, as evidenced by the Business Roundtable's only recent (2019) shift toward acknowledging commitments to all stakeholders (Benoit 2019). We argue that this heuristic poses a meaningful benchmark that allows us to evaluate the opportunity costs or gains associated with using any of the multi-objective heuristics in the CSR context. For instance, some hedge funds may evaluate investments solely based on risk-adjusted returns, where social considerations matter only if they affect those returns.

2. *Combine social and financial performance.* Under this rationale, decision-makers evaluate social and financial performance simultaneously by using a composite measure that merges both dimensions into a single metric. This approach reflects a decision maker who assigns specific weights to each type of performance, rather than prioritizing one exclusively. Behavioral theory refers to such decision processes as "additive," since they sum performance across multiple dimensions (Bromiley and Harris 2014: 341; Gaba and Greve 2019: 651). Case studies in the CSR literature show that organizations implement this rationale by encouraging employees to seek compromises among different objectives (Battilana, Pache, Sengul, & Kimsey 2019: 129). Because this approach aggregates social and financial performance, it bears similarities to utilitarian ethics, which seek to maximize total welfare across all relevant individuals (Bentham 1789/1996, Mill 1848/2004). A possible example is a social enterprise operating under a "buy-one-give-one" model, where business decisions are evaluated through both commercial and social consequence criteria simultaneously (e.g., Marquis & Park 2014).

Notably, combining social and financial performance in this manner requires that these dimensions be commensurable—that is, capable of being measured and aggregated on a common scale (Carruthers & Stinchcombe 1999, Espeland & Stevens 1998). The logic of this heuristic relies on the ability to merge distinct domains into a unified metric, a process that mirrors real-world

practices such as ESG (Environmental, Social, and Governance) scoring systems and triple bottom line accounting. Both of these approaches aggregate multiple performance dimensions into a single evaluative score, using explicit or implicit weighting schemes. However, achieving such commensurability is challenging. The literature on social impact measurement highlights the lack of consensus on how to measure or compare social value creation, which is itself a multidimensional and complex construct (Kroeger & Weber 2014, Rawhouser et al. 2017). As a result, firms may adopt different methods for combining social and financial performance.

3. *Alternate between pursuing social and financial performance.* This rationale embodies the observation in the behavioral theory of the firm that organizations often focus on one goal at a time (Cyert & March 1963: 118). Like Combine, it can reflect a firm's joint preference for both dimensions. From time to time, organizations switch from one goal to another, typically when progress along one goal has plateaued and that goal therefore ceases to be salient (Greve 2008: 480, Bromiley and Harris 2014: 344). In the context of CSR, McGuire, Sundgren, and Schneeweis (1988) provide some evidence consistent with this rationale by finding that social investments are often preceded by significant improvements in financial performance. In a similar vein, Tantaló and Priem (2016: 315) point out that managers often deal with CSR pressures by rotating their attention among different trade-offs. For example, a hospital may, for an extended period, prioritize financial viability (profitability) as rising costs place increasing pressure on operations. Over time, however, management may shift its focus toward population health outcomes, that is, a mission-driven objective that includes, for instance, expanding charity care for underserved patients—a shift that could be triggered by heightened social pressures, leadership turnover, or changes in governance (e.g., takeover by a private equity owner). This “alternating” approach is closely related to the idea of vacillation in organizational choices, where firms dynamically shift their focus between alternatives over time (Nickerson & Zenger 2002). Furthermore, the alternation heuristic shares conceptual similarities with “mixed strategies” in non-cooperative game theory, where agents randomize their actions across alternatives to achieve desired outcomes (Ghemawat & Nalebuff 1985, Shinkle, Kriauciunas, & Hundley 2013). However, while mixed strategies are typically stochastic

and formalized in game-theoretic terms, the alternating heuristic described here is deterministic and rooted in behavioral theories of sequential attention to goals.⁶

4. *Penalize social underperformance.* This rationale builds on the logic outlined by McWilliams and Siegel (2001). According to this approach, firms assess different alternatives by considering a modified measure of financial performance that takes into account their desired level of social performance. For instance, a car manufacturer may strive to produce cars with zero carbon footprint. When evaluating new car designs, the company will penalize designs that fall short of this desired carbon footprint. This rationale reflects the decision-makers' aim to maximize financial performance while also ensuring a desired level of social performance. Note that penalizing social underperformance differs from maximizing financial performance because it not only accounts for financial performance (which already incorporates the costs or benefits associated with the firm's social performance) but it also takes into account the failure to meet a self-imposed social threshold. For example, a manufacturer may conclude that because its social performance goals were not met, overall performance should be discounted—even if financial and operational targets were nominally achieved. This “recalibration” can, for instance, reduce bonus payments for decision makers. Like the Combine heuristic, Penalize requires that social and financial dimensions be made commensurable so that shortfalls in social performance can be translated into financial penalties and aggregated into a single evaluative metric (Carruthers & Stinchcombe 1999, Espeland & Stevens 1998). Because of this, this approach is also consistent with utilitarianism (Bentham 1789/1996, Mill 1848/2004).

5. *Satisfice social performance.* Under this rationale, firms prioritize financial performance as long as an aspired social threshold is cleared. Similar to Penalize, this logic reflects a decision maker's preference for financial performance, conditional upon achieving their chosen level of social performance. When the social threshold is not cleared, the firm focuses solely on improving social performance. This rationale has its roots in the behavioral theory of the firm; in particular, the argument that firms engage in problemistic search when they underperform their aspiration on a

⁶Please note that the reported results are still stochastic for Alternate in the sense of Monte Carlo simulations.

specific criterion (March and Simon 1958/1993: 161–162, Cyert and March 1963: 119–120). In the context of CSR activities, this implies that attempts to enhance social performance cease once the firm has attained a satisfactory social outcome. The focus then shifts toward enhancing financial performance.

This logic is reflected in practices such as B Corp certification, regulatory compliance, or industry codes of conduct, where firms must meet minimum standards (e.g., for labor rights, environmental impact, or diversity)—or aim to achieve certain certifications (e.g., ISO standards)—before pursuing additional financial or strategic objectives. Importantly, this approach is closely related to the concept of lexicographic preferences in economics, where decision-makers prioritize one criterion absolutely over others, only considering secondary dimensions once the more relevant dimensions have been satisfied (Encarnación 1964, Moldau 1992). In this way, the satisficing heuristic reflects a lexicographic ordering of objectives, in contrast to the compensatory logic of utilitarian or additive approaches. Unlike Combine and Penalize, Satisfice does not require translating social performance into financial units; rather, it requires that the decision maker can (i) assess whether the social threshold is met and (ii) evaluate improvements within the currently focal dimension. For example, a firm that intends to satisfice on labor practices may easily assess whether the minimum standards are met (e.g., a cap on working hours per day) without measuring worker welfare in financially commensurable terms or measuring welfare beyond that threshold.

These five heuristics represent distinct approaches to navigating the multidimensional performance challenges prevalent in the CSR context. In the following sections, we formally model these heuristics, explore their performance under varying conditions, and offer insights into their relative effectiveness and the trade-offs they entail.

The focus of the analyses in the ensuing sections is to understand the consequences of using these different heuristics (i.e., their performance outcomes) rather than their antecedents (i.e., why organizations adopt one heuristic over another). This latter question, while important, falls outside the scope of our paper. However, it is worth noting that heuristic adoption is likely influenced by several factors, including incentive conflicts between principals and agents (Jensen & Meckling

1976) and the tendency to focus on outcomes that are more easily measured and rewarded (Ethiraj & Levinthal 2009, Kerr 1975). For instance, managers may gravitate toward heuristics that emphasize financial performance when social outcomes are difficult to measure or when compensation structures primarily reward financial results. These complexities in organizational decision-making further underscore the importance of understanding how different heuristics perform once adopted.

MODEL

To understand the effect of search heuristics on firms' social and financial performance, we develop a model to compare the performance of the five search heuristics described above when applied to a multi-attribute landscape (Adner et al. 2014). The remainder of this section (i) formally describes the five search heuristics and (ii) explains how we simulate and compare their social and financial performance.

Modeling Dual-Goal Search Heuristics

The heuristics we model capture the five decision-making rationales mentioned in the previous section: *Maximize* financial performance, *Combine* social and financial performance, *Alternate* between pursuing social and financial performance, *Penalize* social underperformance, and *Satisfice* social performance. From here on, we simply call these heuristics Maximize, Combine, Alternate, Penalize, and Satisfice. Our heuristics can further be categorized by the preferences they inherently reflect, that is, single dimension preference (Maximize), joint preference for social and financial performance (Combine and Alternate), and contingent preferences (Penalize and Satisfice).

Following prior research, we model firms as boundedly rational agents (Simon 1955), whose search for better strategies is inherently local and incremental, centered around modifications to their current strategy (Cyert & March 1963, Levinthal 1997). At each time period, each firm in our model performs local search, comparing its performance to the performance of the neighboring strategies. A period represents one evaluation/choice cycle (e.g., a stage-gate, KPI review, budgeting round, or managerial attention cycle), not a calendar year. For simplicity, we assume that strategy vectors only include 0s and 1s; thus, there are N neighboring strategies (e.g., if the current strategy

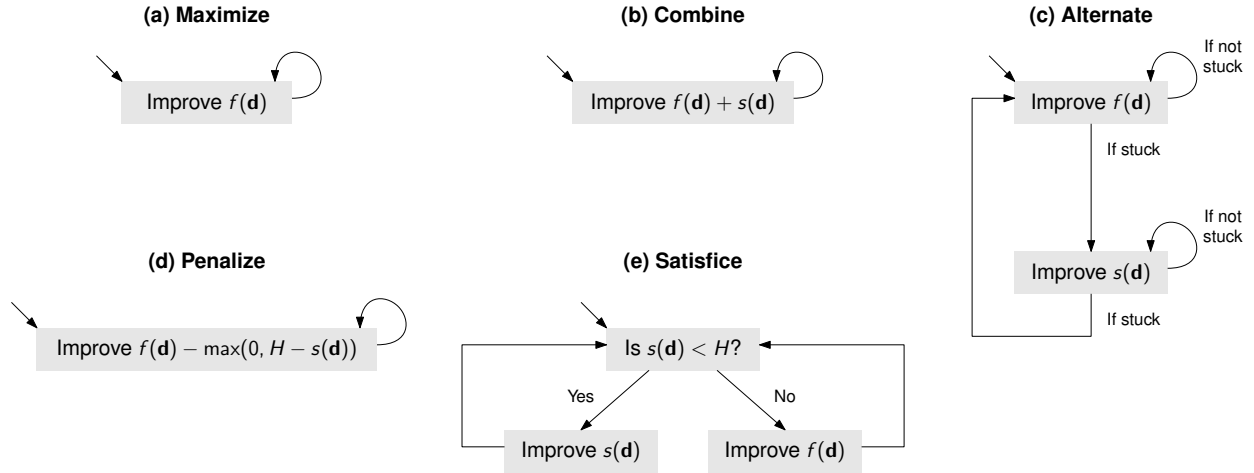


Figure 5: Flow diagrams describing the five search heuristics we analyze.

is $(0, 0, 0)$, there are three neighboring strategies: $(1, 0, 0)$, $(0, 1, 0)$, and $(0, 0, 1)$.

How performances are compared depends on the heuristic. Each strategy \mathbf{d} has a financial value $f(\mathbf{d})$ and a social value $s(\mathbf{d})$, which we normalize to $[0, 1]$ for comparability across runs. Figure 5 summarizes the five heuristics we analyze using the following notation: \mathbf{d} is the firm’s decision vector; $f(\mathbf{d})$ and $s(\mathbf{d})$ are its financial and social performance (the “height” on each landscape in Figure 4); and H is the social threshold. The “improve” boxes denote one step of local search: the firm evaluates its N one-bit neighbors and moves if the heuristic identifies a superior neighbor. Importantly, only the compensatory heuristics (Combine and Penalize) require cross-dimensional commensuration—i.e., an ability to aggregate or trade off $f(\mathbf{d})$ and $s(\mathbf{d})$ in a single evaluative criterion—whereas Alternate and Satisfice rely on within-dimension comparisons and/or a threshold check, imposing weaker commensurability demands (we return to these implementation implications in the Discussion).

Maximize (panel (a)). This heuristic only measures performance in terms of $f(\mathbf{d})$ and, thus, is indifferent to social performance. At each time period, Maximize tries to improve financial performance by performing a local search on the financial landscape. Maximize provides a baseline against which to compare the remaining heuristics.

Combine (panel (b)). This heuristic measures performance as the sum of financial and social performance. That is, at each time period, a firm using this heuristic tests all neighboring strategies

and picks the one that has the highest sum of social and financial performance (that is, it makes choices based on $f(\mathbf{d}) + s(\mathbf{d})$).⁷ In Appendix A, we explore the more general class of “Weighted Combine” heuristics, where performance is measured as the weighted sum $wf(\mathbf{d}) + (1 - w)s(\mathbf{d})$, with w in $[0, 1]$.

Alternate (panel (c)). This heuristic works on one goal at a time. It picks a first goal at random and then continues trying to improve that goal for as long as this is possible (that is, until “stuck” at a local peak). Once improvement on a given goal ceases to be possible, the heuristic switches to trying to improve the other goal, and so on.

Penalize (panel (d)). This heuristic measures performance as financial performance minus a penalty if the firm is underperforming the social threshold H . The penalty is the amount of social underperformance, $\max(0, H - s(\mathbf{d}))$.

Satisfice (panel (e)). This heuristic tries to improve social performance as long as that performance is below the social threshold H ; otherwise, it attempts to improve financial performance.

Note that it is possible to model other heuristics apart from the five just described. For instance, one could model the mirror images of Maximize and Satisfice (i.e., “maximize social performance” and “satisfice financial performance”). We focus on the five heuristics described so far to keep the presentation of the results tractable and because they best illustrate the five logics highlighted by the CSR and multiple goals literatures. In any case, once one understands the behavior of these five heuristics, it is not difficult to predict the behavior of other heuristics.

Measuring Performance via Simulation

To compare the performance of the five search heuristics, we proceed as follows. First, we create a dual-performance landscape with complexity K and correlation ρ (an overview of this method was presented in the Theoretical Background section; for details see Adner et al. 2014: 2796–2798). To make results comparable across simulations, we scale social and financial performance so that 0

⁷This rule is different from Ethiraj and Levinthal’s (2009: 14) multi-objective benchmark rule, which only moves to a new position if that position weakly dominates the current position (that is, if no performance dimension declines). A more fundamental difference with Ethiraj and Levinthal (2009) is that they only examine the effect of search heuristics on one performance measure (fitness).

corresponds to the global minimum and 1 to the global maximum of its respective landscape. This makes it easier to interpret our comparisons.

Second, we create five firms that start at the same random point on the landscape. Each firm uses exclusively one of the five heuristics for its search over the multi-attribute landscape. Third, we let the firms search for 200 time periods. We chose this number because, by then, each firm’s average performance has plateaued. We record each firm’s social and financial performance in period 200. Finally, we repeat the previous three steps 10,000 times for each combination of parameters we explore. This allows us to estimate the expected social and financial performance of each heuristic given complexity (K), correlation (ρ), and social threshold (H). In additional analyses reported in the Appendix B (and summarized in the Results section), we partially relax strict local search by allowing occasional non-local (random) moves during search.

RESULTS

We report our results using scatterplots that allow for an intuitive and precise description of the model’s behavior. The points in these scatterplots show the different heuristics measured against social and financial performance (on the vertical and horizontal axes, respectively). Unless otherwise noted, each point represents average performance at the end of the simulation ($T = 200$) over 10,000 Monte Carlo trials for a given parameter setting.

The particular values used to generate the plots were chosen because they illustrate the whole behavior of the model. We keep the number of decisions fixed at $N = 12$, as different values do not change the behavior of our model qualitatively as long as complexity (K) is scaled with N . For complexity K , we plot low, medium, and high values ($K = 1, 6, \text{ and } 11$) and for correlation ρ , we plot negative, neutral, and positive values ($\rho = -0.5, 0, \text{ and } 0.5$). For the social threshold H , our first analyses use a medium value ($H = 0.5$), which we later expand to include low and high values ($H = 0.2 \text{ and } 0.8$). Understanding our model’s behavior under this set of values is enough to understand its whole behavior.

We structure the exposition of the results as follows: We first describe general patterns that are robust to parameter changes. We then explain what drives the expected performance of each

heuristic and we examine the variability of the results. Finally, we examine the conditions under which some heuristics can match or outperform the social and financial performance of Maximize.

General Patterns: The Effect of Complexity and Correlation

Figure 6 shows nine performance plots (labeled (a) through (i)), which differ in complexity (row-wise) and correlation (column-wise). For all nine plots, the social performance threshold is fixed at $H = 0.5$. Each panel shows the average social and financial performance of each heuristic. The gray broken line in each panel represents the average position of the efficient frontier in that panel.⁸

An overarching observation from Figure 6 is that varying complexity (K) and correlation (ρ) does not change the relative ordering of the heuristics.⁹ For example, Maximize is always near the bottom-right corner and Combine is always the closest to the top-right corner. The reason is that K and ρ affect the landscapes on which search occurs, but they do not alter the directional pull built into each heuristic's acceptance rule. Maximize follows financial improvements only, leaving social performance to depend on the social value of the financial peak it reaches. Combine evaluates moves using the joint criterion $f(\mathbf{d}) + s(\mathbf{d})$, leading it toward balanced compromise positions. Alternate switches between financial and social search, moving firms toward locally attractive positions on each landscape. Penalize and Satisfice both depend on the social threshold H , but in different ways: Penalize adjusts the evaluation of financial moves by subtracting social shortfalls, whereas Satisfice switches to social search until the threshold is met. Thus, K and ρ change how difficult search is, and whether gains in one dimension tend to carry over to the other, but not the basic direction imposed by each heuristic.

This observation suggests that heuristic choice is central to explaining firms' relative positions in the social–financial space. Within the same environment, different heuristics can lead firms to reach systematically different outcomes, depending on whether they search by maximizing one objective, combining objectives, alternating attention between them, or using a social threshold.

⁸Each frontier is approximated using three points (averaged over 10,000 simulations): maximal financial performance, maximal social performance, and maximal sum of social and financial performance.

⁹Here, “relative ordering” refers to the ranking of heuristics by their average locations in the social–financial plane and to which heuristics tend to be closer to the top-right, bottom-right, or top-left regions of each panel.

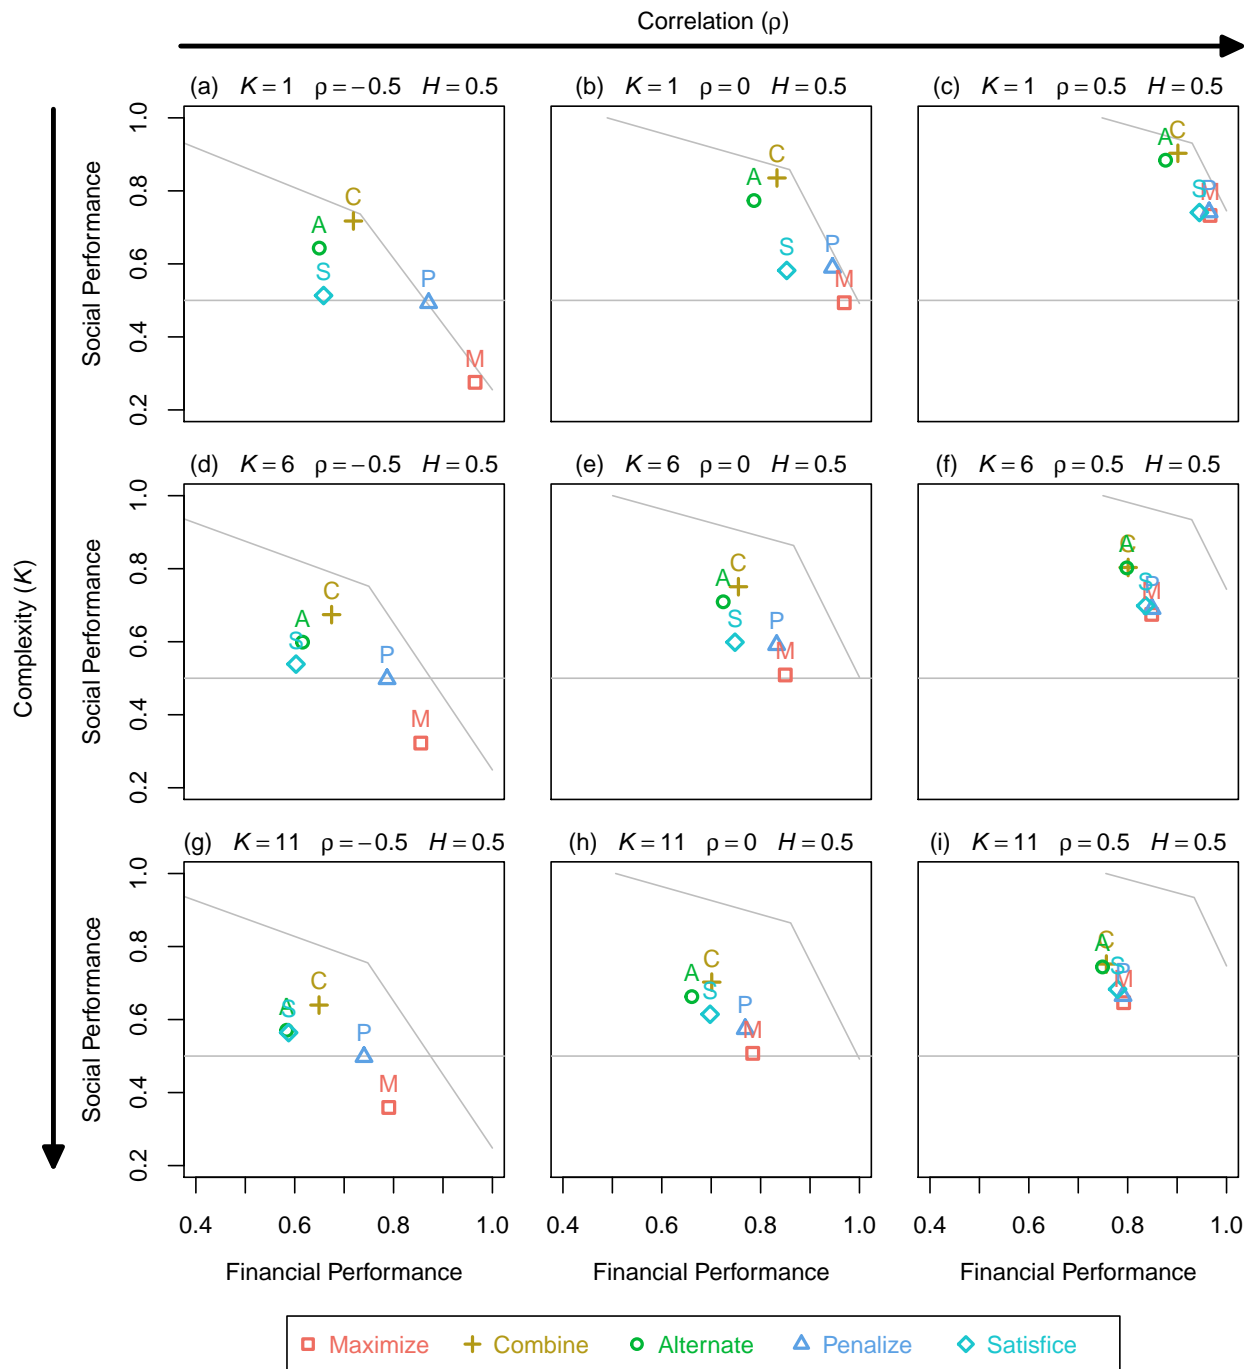


Figure 6: Performance plots for the heuristics as a function of complexity (K) and correlation (ρ).

At the same time, environmental structure affects the stakes of heuristic choice. When ρ is high or K is low, the heuristics cluster more tightly, making the choice among them less consequential within the model. When ρ is low and K is high, differences across heuristics become larger because joint improvement is harder and local search is more constrained. Thus, environmental structure shapes the difficulty of joint improvement, whereas heuristic choice shapes where search tends to land within that environment.

As correlation increases, the performance of all heuristics improves along both dimensions. This is evident in Figure 6, where all points move closer to the top-right corner of each panel as one moves from left to right across the columns. The underlying reason is that higher correlation means that improvement in one dimension is likely to lead to improvement in the other. For example, if social and financial performance are correlated, a firm that enhances its social performance is also likely to see gains in its financial performance. In other words, increased correlation strengthens the extent to which progress on one landscape facilitates progress on the other. This effect is particularly apparent for the Maximize heuristic, which ignores social performance but nevertheless sees improvements in this dimension as correlation rises. Overall, increasing the correlation between performance dimensions not only boosts performance across all heuristics but also reduces the absolute differences among them.

When complexity increases, the performance of all heuristics decreases along both dimensions. In Figure 6, this can be seen as all points moving farther away from the top-right corner of each panel as one moves row-wise from top to bottom. The reason is that as complexity increases, the performance landscape on both dimensions becomes more rugged and, hence, harder to search; the landscapes have more local peaks where search can get stuck (Levinthal 1997). Hence, all heuristics are less likely to find the global peak and instead quickly get stuck near their starting position.

The fact that the ordering of the heuristics is stable across the parameter settings we examine allows us to simplify the exposition of the results and focus on explaining the relative financial and social performance of each heuristic. Because performance differences are most visible in panel (a) in Figure 6, our descriptions below refer to that panel unless stated otherwise.

Relative Performance of the Heuristics That Do Not Depend on a Threshold

We now move on to describe the relative ordering of the different rules. For clarity, we first analyze the decision rules that do not take into account the social threshold (H); namely, Maximize, Combine, and Alternate.

Maximize achieves the highest financial performance and the lowest social performance across all panels in Figure 6. To understand this behavior, recall that this decision rule performs an unconstrained search over the financial landscape. Hence, this heuristic devotes all its attention to maximizing financial performance without having to compromise to achieve any particular social performance. As a result, it is natural that Maximize appears at the bottom-right corner of all panels in Figure 6.

The next two heuristics—Combine and Alternate—both reflect a joint preference for social and financial performance, yet they differ with respect to how they pursue these preferences. While Combine pays attention to both objectives *at the same time*, always looking at the sum of social and financial performance, Alternate pays the same attention to these goals but *over time*, switching its focus back and forth between social and financial performance. In both cases, joint attention to both performance measures is reflected in the fact that the average performance of both heuristics falls along the 45-degree diagonal of each panel.

Interestingly, Combine comes closer to performing ideally than Alternate; it reaches closer to each panel's top-right corner, where both social and financial performance are maximized. A priori, there are arguments for either of these to outperform the other. On the one hand, Combine searches over a landscape that is the sum of the social and the financial landscapes and therefore generally more rugged than either of those. The higher effective complexity faced by Combine means that it will get stuck quickly; however, it will do so at a peak that represents a compromise between the two performance dimensions. On the other hand, Alternate will face less rugged landscapes, as it searches one dimension at a time. However, Alternate is unlikely to settle on any given point: once it has found a local peak on, say, the financial landscape, it will switch to searching on the social landscape, and so on. That is, Alternate never settles at a point that is a good compromise between

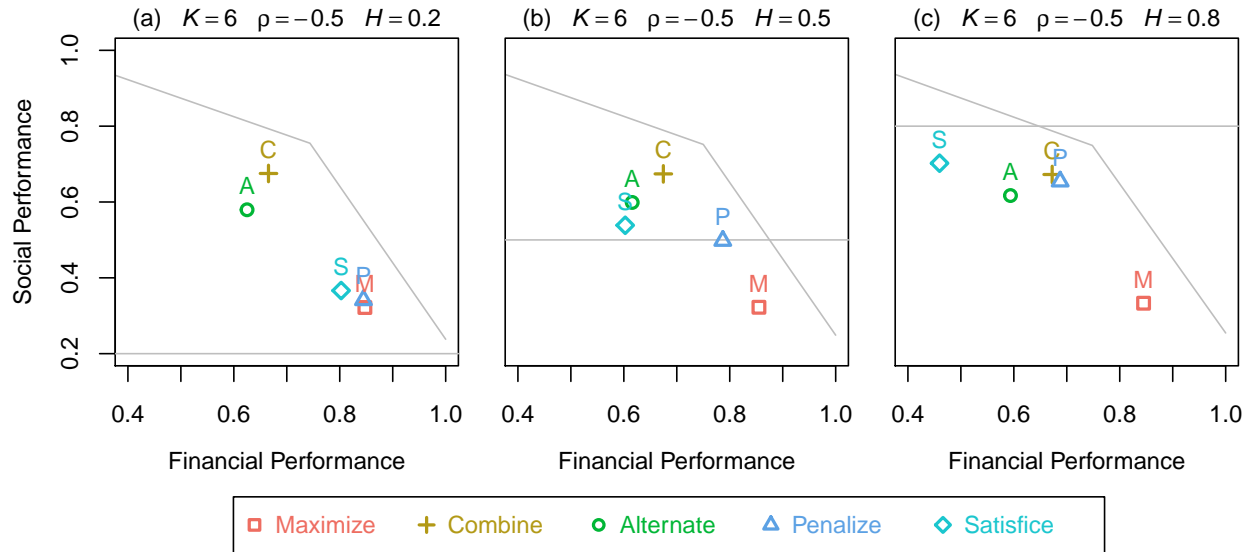


Figure 7: Performance plots for different social thresholds (H).

the two performance dimensions (unless that point happens to be a local peak in both landscapes). The result is that Combine finds a good position and stays there, while Alternate continues to explore and the average of that continuous exploration is farther from the ideal than Combine’s solution. We will add nuance to this conclusion later when we discuss the ability of the different heuristics to explore.

Relative Performance of the Heuristics That Depend on a Threshold

We now analyze the remaining heuristics, Penalize and Satisfice, both of which reflect a preference for financial performance that is contingent on meeting a minimum social performance threshold (H). Here, we will be looking at Figure 7, which uses the same K and ρ values as panel (d) in Figure 6 while varying the social threshold to take low, medium, and high values ($H = 0.2, 0.5,$ and 0.8).

The Penalize heuristic moves from being close to Maximize to being close to Combine as the social threshold increases (i.e., as one moves from panel (a) to panel (c) in Figure 7). This happens because the threshold affects the topography of the penalty-adjusted landscape that this heuristic searches. When $H = 0$, there is never a penalty associated with social performance, as it is impossible to underperform this threshold. Hence, the penalty-adjusted landscape is simply

the financial landscape and, thus, the Penalize heuristic behaves identically to Maximize (that is, if $H = 0$, then $f(\mathbf{d}) - \max(0, H - s(\mathbf{d}))$ reduces to $f(\mathbf{d})$). Conversely, when $H = 1$, social performance is always underperforming the threshold, making the penalty-adjusted landscape an even mix of social and financial performance. Penalize then behaves identically to Combine (that is, if $H = 1$, then $f(\mathbf{d}) - \max(0, H - s(\mathbf{d})) = f(\mathbf{d}) + s(\mathbf{d}) - 1$, which makes the same choices as Combine's $f(\mathbf{d}) + s(\mathbf{d})$ formula). For intermediate values of the threshold ($0 < H < 1$), the average behavior of Penalize interpolates between these two extremes.

The Satisfice heuristic moves from being close to Maximize to achieving high social performance and low financial performance as the social threshold increases (i.e., as one moves from panel (a) to panel (c) in Figure 7, Satisfice approaches the top-left corner). This happens because the social threshold affects which performance dimension this heuristic will search. When $H = 0$, Satisfice will never switch to searching for social performance improvements, as the social threshold is obviously cleared. Satisfice is then equivalent to Maximize. Conversely, when $H = 1$, the firm will only search on the social landscape, effectively turning Satisfice into a rule that could be called “Maximize Social” and which is the mirror image of Maximize but with respect to social performance (which is why Satisfice is close to the top-left corner in panel (c)). For intermediate values of the threshold ($0 < H < 1$), the average behavior of Satisfice interpolates between the two extremes. Note that unless the social threshold is very high (as in panel (c) in Figure 7), Satisfice is usually dominated or closely matched by one of the other rules, as we see in all of Figure 6 and in panels (a) and (b) in Figure 7.

Variability of Performance

So far, we have looked at the average performance of the five heuristics. We now examine the variation around these averages; that is, we look at the whole range of performance for firms using a given heuristic. This gives us an idea of the degree to which a heuristic determines firm performance. Each scatterplot in Figure 8 represents the social and financial performance at the end of the simulation ($T = 200$) for 10,000 firms, each starting at a random position in a random landscape. Each panel represents one heuristic. The environmental parameters are kept fixed at

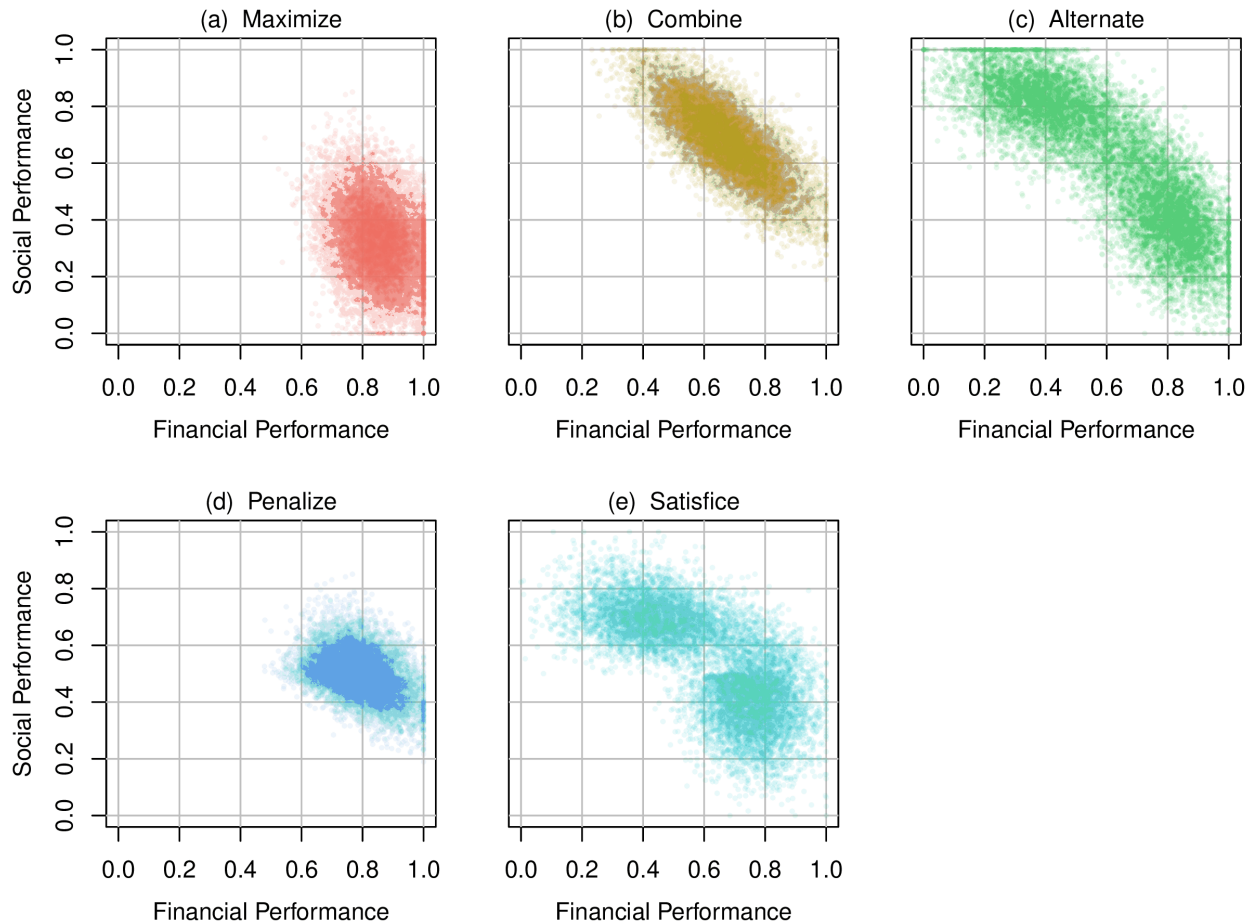


Figure 8: Social and financial performance of 10,000 simulations of each heuristic.

$K = 6$, $\rho = -0.5$, and $H = 0.5$ (the same as in panel (d) in Figure 6).

A general observation from Figure 8 is that although within each panel there is much variability in the positions achieved (i.e., each cloud of points covers a large area), there are marked differences across panels (i.e., each cloud of points has a distinctive shape and location)—that is, heuristics have an important effect on which outcomes are more or less likely to occur. Below, we examine what drives each panel’s most salient characteristics.

The cloud of points representing Maximize is taller than it is wide (panel (a)). This happens because Maximize reliably gets firms to a local peak in the financial landscape while not paying any attention to social performance. Hence, the width of the cloud of points in panel (a) represents the range of financial local peaks, while the height of the cloud of points represents whatever social performance happens to be associated with those financial peaks.

The clouds of points for Combine and Alternate are diagonally oriented (panels (b) and (c)). Even though the shape of the clouds is similar, the underlying reason is quite different. Combine tries to find points that rank highly on the sum of both performance dimensions. Because correlation is negative in Figure 8, high performance on one attribute will usually be associated with low performance on the other; hence the downward-sloping cloud.

The shape of the cloud for Alternate (panel (c)) is due to the sequential nature of the search. Recall that while this heuristic is searching on one dimension, the other dimension is ignored. Hence, the cloud represents a snapshot of firms at different moments in their oscillation between goals: some are near the peak that Maximize would find and some are near the peak a “Maximize Social” heuristic would find. Note that Alternate is more spread out than Combine. This is because Alternate typically does not settle on a given point (unless that point happens to be a local peak in both landscapes).

Penalize (panel (d)) produces the least variation in outcomes. The reason is that Penalize “sees” a very rugged landscape, one that combines the ruggedness of the financial landscape and the ruggedness of the landscape representing the penalties (that is, Penalize sees the landscape generated by adding $f(\mathbf{d})$ and $-\max(0, H - s(\mathbf{d}))$). Hence, firms using Penalize get stuck at a local peak more rapidly than with the other heuristics.

Satisfice (panel (e)) leads to a bimodal distribution, as firms using this heuristic behave differently depending on whether or not they can clear the social threshold. If they cannot clear it, they try to improve social performance until they do, at which point they try to improve their financial performance, and so on. Thus, the bimodal distribution represents a search that oscillates between the points at which the two search modes are triggered.

One last observation from Figure 8 is that the size of the clouds of points depends on the extent to which each heuristic explores the landscape. That is, the heuristics that explore the least, like Maximize (which searches until getting stuck on a financial local peak) or like Penalize and Combine (which search on more rugged landscapes and, hence, get stuck earlier) exhibit less variability of outcomes than the heuristics that explore more, like Alternate and Satisfice, which switch between

exploring the social and financial landscapes. Appendix B shows that allowing occasional non-local (random) jumps does not alter these central qualitative conclusions, but as random jumps become more frequent, differences among heuristics compress, and outcomes drift toward chance-level performance.

Outperforming Maximize

Proponents of Maximize (e.g., Friedman 1970, Jensen 2002) suggest that any multi-objective decision-making will fare worse financially than attempts to solely maximize financial performance.¹⁰ We investigate this statement by tracking the probability that our search heuristics perform at least as well as Maximize on *both* performance dimensions. When we say heuristic *A* outperforms heuristic *B*, we mean that *A weakly dominates B*: *A* achieves at least as high performance on both dimensions ($f_A \geq f_B$ and $s_A \geq s_B$) and may be strictly higher on one or both dimensions. We use “match” to refer to the special case in which performance is equal on both dimensions.

Figure 9 plots the same nine scenarios as Figure 6; however, the *y*-axis now shows the cumulative probability of *matching or outperforming* Maximize (that is, weakly dominating Maximize at some point during the search trajectory). Specifically, at each point in time, we check whether each of the other four heuristics has, at any point up to that time, reached a position that weakly dominates Maximize’s performance.¹¹ This measure counts the fraction of cases in which a heuristic that does not solely pursue financial gain can do at least as well as one that does.

Note that this comparison differs from Figure 6 in an important way. Figure 6 reports where each heuristic is located at the end of the search process, on average. Figure 9, by contrast, asks whether a heuristic has discovered, at any point along its search trajectory, a strategy that weakly dominates the position reached by Maximize. The figure therefore captures discovery rather than final location. In practice, this assumes that firms can retain or revisit previously discovered alternatives, as when a firm explores multiple possibilities before settling on one encountered earlier.

¹⁰A closely related critique in the CSR debate is that pursuing social objectives may distract from profitability and thereby reduce financial performance (see, e.g., Karnani 2011).

¹¹More formally, the *y*-axes in Figure 9 measure performance for heuristic *H* at time *T* as $\frac{1}{\#\text{simulations}} \sum_{\text{simulations}} \mathbb{1} \left[\bigvee_{t=1}^T ((f_t^H \geq f_t^{\text{Maximize}}) \wedge (s_t^H \geq s_t^{\text{Maximize}})) \right]$.

For example, Microsoft's decision to rekindle its efforts to develop Windows after terminating its deep involvement with IBM's OS/2 (Richards 1991) illustrates how firms may return to a previously explored trajectory after searching elsewhere.

This distinction explains why Combine can appear closer to the frontier in Figure 6, while Alternate stands out in Figure 9. Combine searches on the composite landscape $f(\mathbf{d}) + s(\mathbf{d})$ and therefore tends to settle near balanced positions, but these need not match the financial local peak reached by Maximize. Alternate, by contrast, may end farther from the frontier because it keeps shifting between objectives. Yet that shifting also helps it discover new regions: after social search moves the firm away from Maximize's financial peak, later financial search can climb to a different financial peak. Some of these peaks also have high social performance, allowing Alternate to weakly dominate Maximize in both dimensions.

The main result stemming from Figure 9 is a challenge to the proponents of Maximize: some heuristics match or outperform Maximize more than half of the time. That is, some heuristics can often discover strategies that are at least as good as Maximize's current financial peak while also delivering at least as high social performance. The explanation for this result is that bounded rationality dictates that firms cannot optimize, but only search. This means that Maximize cannot really find the maximum financial performance but simply a local peak in the financial landscape. In contrast, heuristics that explore more broadly can, given enough time, discover positions that dominate the particular local peak found by Maximize. In other words, using heuristics that attend to multiple dimensions of performance can lead to *both* better social performance and, in many cases, better financial performance than focusing narrowly on financial performance.

To get a better intuition for the underlying mechanisms driving this result, let us delve into some of the patterns observable in Figure 9. In all scenarios, Alternate is able to surpass the 50% line (i.e., given enough time, Alternate matches or outperforms Maximize). Reaching the 50% line happens because half of the time Alternate starts its search on the financial landscape and therefore behaves exactly like Maximize until it reaches the financial local peak, after which it switches to searching on the social landscape. Surpassing the 50% line is due to Alternate's greater exploration:




Figure 9: Probability of matching or outperforming Maximize.

by alternating attention between the two landscapes, it can leave the financial local peak found by Maximize, move to a different region while searching socially, and then climb financially from that new starting point. This process can reveal financial peaks that Maximize, because of its narrower path-dependent search, never visits.

Another observation is that when correlation is positive (i.e., panels (c), (f), and (i)), Penalize has the greatest probability of matching or outperforming Maximize. This is because Penalize searches what can be termed a “stretched financial landscape,” where all points below H have their financial performance reduced. In these particular panels, financial outcomes exhibit positive correlations with social outcomes, meaning that this reduction of the landscape does not significantly alter the ranking of neighboring values. Therefore, it is quite common for Penalize to locate the same position as that found by Maximize.

The finding that certain heuristics can identify strategies that match or surpass those of the Maximize approach suggests that these heuristics are effective at pursuing what Eno and Schmidt (2001) call an “oblique strategy.” By alternating attention between social and financial goals, heuristics like Alternate are able to zigzag across the performance landscape, exploring regions that Maximize, with its singular financial focus, typically overlooks. This broader search allows Alternate to escape local financial optima and sometimes reach positions that deliver higher performance on both dimensions. Thus, Alternate is not necessarily the best heuristic if performance is evaluated only by the average final position. Its distinctive advantage is that it expands the opportunity set generated by search, making it more likely that the firm discovers a strategy that does as well as Maximize financially while also doing at least as well socially.

DISCUSSION

A central challenge for boundedly rational organizations is how to search for strategies when performance is multidimensional and goals potentially conflict. The behavioral theory of the firm has long recognized that organizations pursue multiple goals, yet the search literature has largely focused on single-objective landscapes. Our study addresses this gap by examining how different search heuristics shape the discovery of alternative strategies, when firms must navigate

dual performance dimensions. Below, we elaborate on the broader contributions of our work.

Contributions to Behavioral Strategy Research

Our work contributes to understanding the relationship between multi-objective decision-making and organizational search. The multi-objective decision-making literature has traditionally focused on how decision-makers evaluate and select among a given set of alternatives (Keeney & Raiffa 1976). This framing treats the set of alternatives as fixed and asks which should be chosen. In contrast, the behavioral theory of the firm emphasizes that alternatives are not given but must be discovered through search (Cyert & March 1963, Levinthal 1997). Our work bridges these perspectives by examining how evaluation heuristics shape the trajectory of organizational search—that is, how the criteria used to compare currently available alternatives influence which new alternatives the firm will subsequently discover. In a multi-goal setting, this distinction is especially consequential because preferences are not only expressed in one-shot choice; they are implemented repeatedly, step by step, through a local evaluation rule that generates path dependence in what the organization ever learns is feasible.

This connection between evaluation and search helps explain why the Alternate heuristic can match or even surpass the performance of the single-goal focused Maximize approach. Alternate does not outperform Maximize because it is superior at any single evaluation. Rather, the sequential nature of its evaluation criteria produces a distinctive search trajectory: by oscillating between objectives, Alternate generates alternatives in regions of the landscape that Maximize, with its singular focus, is less likely to traverse. The idea of oblique strategies captures this mechanism. As Kay (2010, Chapter 3) argues, firms may accomplish their objectives by not approaching them directly—that is, by avoiding narrowly focused and linear pursuit. In our setting, the obliqueness stems from shifting the evaluative criterion over time: the wider pattern of alternative generation that emerges from Alternate’s oscillating priorities enables it to uncover high-performing solutions that remain hidden to heuristics with more narrowly focused evaluation criteria.

Our results also add nuance to classic behavioral arguments about sequential attention to goals. The behavioral theory of the firm has emphasized that firms may address conflicting goals by

rotating attention among them (Cyert and March 1963: 118, Fang and Kim 2018: 2553, Gaba and Greve 2019: 649). In our simulations, however, simultaneous evaluation can, on average, outperform sequential evaluation (Combine outperforms Alternate throughout Figure 6). Combine has in its favor the ability to continuously make trade-offs between dimensions, whereas Alternate’s one-dimensional focus can drive it “up” one landscape at the expense of the other. This is disadvantageous if the primary aim is to approach a balanced ideal point, but it can be advantageous when the aim is broader exploration—a distinction that becomes central when evaluating the probability of matching or outperforming Maximize.

Finally, we contribute by extending the NK landscape framework to study multiple performance dimensions. While the NK framework is a staple in the search literature, it has customarily been used to examine search with respect to *one* overall measure of performance, often called “fitness.”¹² Adner et al. (2014) developed a way to create landscapes with multiple performance dimensions, which they used to study spatial competition but not search. Building on their work, we show how to model search on multi-attribute landscapes and how heuristic choice interacts with landscape structure (complexity K , correlation ρ , and threshold H) to shape *joint* performance outcomes. Our “grammar” for describing multi-objective heuristics (Figure 5) provides a compact template that can be extended to additional rules in future work.

Contribution to the CSR literature

Jensen’s (2002: 238) famous statement that “[i]t is logically impossible to maximize in more than one dimension at the same time [. . .]” relies on a basic understanding of optimization. We have suggested that the literature on multi-objective maximization provides a better theoretical foundation for reasoning about multiple performance dimensions. This shift calls attention to the role of the search heuristics that take the place of the “greater than” operator used in unidimensional settings. It also brings to the fore the distinction between preferences and heuristics and the fact that even

¹²Although there is research in the NK literature that has modeled multiple sources of performance (e.g., Csaszar and Levinthal 2016 assign performance to different attributes of a product and Rivkin and Siggelkow 2003 and Ethiraj and Levinthal 2009 assign performance to different divisions of a firm), such research has only studied how search heuristics affect just *one* measure of fitness.

heuristics that reflect the same preferences can generate different outcomes because they guide search differently.

Our work helps develop the role of decision processes in CSR research. Prior research has often overlooked this dimension, in part because of the methodological challenges involved in studying how organizations actually make multi-objective decisions (Aguinis & Glavas 2012: 953). The methodology we introduce allows us to open this black box by analyzing how alternative heuristic implementations translate into systematically different search trajectories and, consequently, different social and financial outcomes.

Because search heuristics are so consequential in the environments we model, choices should be made carefully. In particular, some heuristics are rarely attractive: Satisfice is dominated (or closely matched) in most conditions, and Maximize systematically yields the lowest social performance, often leaving the firm outside the viable region implied by nontrivial social aspirations. The practical implication is therefore not that one heuristic is universally best, but that heuristic choice can materially shape which compromises are reached and which opportunities are discovered under bounded search.

Choosing among heuristics is nuanced because heuristics can encode similar preferences while differing in implementation. Penalize and Satisfice, for instance, both express contingent preferences—high financial performance subject to clearing a social threshold—yet Penalize is typically more effective at attaining that combination in our simulations. Likewise, Combine and Alternate can both reflect joint concern for social and financial outcomes, but they operationalize that concern differently: Combine tends to move toward a balanced compromise and then settle, whereas Alternate tends to continue exploring by shifting attention across dimensions. A key consequence of this exploration advantage is that Alternate can sometimes discover strategies that match or even outperform Maximize—including on financial performance—which is consistent with the view that pursuing social objectives can, under some conditions, open paths to superior financial outcomes (Dorobantu et al. 2017, Flammer 2015).

Beyond explaining the workings of specific heuristics, developing the role of heuristics in CSR

is valuable because it helps address the call by Margolis, Elfenbein, and Walsh (2009: 31) to identify organizational practices that facilitate the joint achievement of social and financial objectives. Our results also provide formal backing to research suggesting that valuable CSR initiatives can emerge from idiosyncratic search processes (Haanaes, Michael, Jurgens, & Rangan 2013), and they align with a central idea in stakeholder theory: that the role of management is to “search for means to overcome trade-offs” (Hörisch, Freeman, & Schaltegger 2014: 331).

We also clarify how environment structure affects the *stakes* of heuristic choice. Across the parameter ranges we examine, the relative ordering of heuristics is robust to changes in landscape complexity (K) and correlation between objectives (ρ), even though these contingencies shift absolute performance levels and can compress differences among heuristics (e.g., when ρ is high, improvement on one dimension tends to carry over to the other). Thus, while K and ρ shape how difficult joint improvement is, our results suggest that variation in how firms *implement* preferences as search heuristics can be an important contributor to heterogeneity in social and financial outcomes—a possibility that complements the dispersion documented in the CSP–CFP literature (Allouche & Laroche 2005, Gallardo-Vázquez et al. 2019, Margolis & Walsh 2003, Orlitzky et al. 2003, Vishwanathan et al. 2020, Wang et al. 2016).

Our work also contributes to the literature on commensurability. Prior research has emphasized that aggregating different performance dimensions—such as social and financial outcomes—often requires that these dimensions be made commensurable, or comparable on a single scale (e.g., Carruthers & Stinchcombe 1999, Espeland & Stevens 1998). This requirement underpins the logic of heuristics like Combine and Penalize, which necessitate explicit weighting or translation of social and financial performance into a common metric. However, achieving commensurability in practice is notoriously challenging, as the measurement and valuation of social outcomes are often contested and context-dependent (Kroeger & Weber 2014, Rawhouser et al. 2017).

Our findings suggest that strict cross-dimensional commensurability may be less essential for achieving balanced outcomes than previously assumed. For instance, as shown in Figure 6, the expected performance of Alternate—a heuristic that does not require aggregating social and financial

performance into a single metric—is remarkably similar to that of Combine, which does. This does not mean that measurement of social performance is unnecessary. Alternate still requires the organization to assess whether a move improves social performance when social performance is the focal dimension. The point is narrower: Alternate does not require the organization to translate social performance into financial units, assign an exchange rate between the two dimensions, or evaluate marginal trade-offs between them on a common scale.

Put differently, organizations facing barriers to constructing a unified metric can sometimes approximate the outcomes of aggregation-based approaches by using heuristics grounded in sequential attention rather than simultaneous integration. This resemblance is notable precisely because the heuristics impose different informational demands. Combine and Penalize require cross-dimensional integration, such as weights or penalties that translate social performance into the same evaluative currency as financial performance. Alternate requires within-dimension comparison: the organization must be able to assess whether a move improves social performance when social performance is the focal dimension, but it need not compare that improvement against a financial gain on a common scale. A threshold-based heuristic like Satisfice can further reduce measurement demands by shifting the question from “how much better?” to “have we cleared the threshold?” This is particularly useful when it is difficult to assess marginal improvements along a social dimension (Kroeger & Weber 2014, Rawhouser et al. 2017).

More broadly, our approach shifts attention from whether performance dimensions can be perfectly commensurated to how organizations should select among heuristics with different informational and commensurability requirements. Whereas prior work debates whether and how different types of value can be compared or combined, our results show that heuristics requiring cross-dimensional aggregation and heuristics relying on sequential attention or threshold checks can sometimes yield comparable outcomes. This opens a wider practical menu for firms managing multiple objectives, especially when social performance can be assessed within its own domain but cannot be easily monetized or integrated into a unified metric.

Contribution to Other Multi-Objective Contexts

Our work provides an answer to Greve and Gaba's (2020: 323) call to develop a better understanding of search under multiple goals. In particular, we analyze the trade-offs among a set of heuristics for pursuing multiple goals and show how heuristics that encode similar preferences can nonetheless generate systematically different search paths and outcomes. Understanding such trade-offs may shed light on problems not only involving social and financial performance but also other (conflicting) goals such as willingness-to-pay and cost (Porter 1996), short- and long-term performance (Levinthal & March 1993), and quality and quantity (Shaked & Sutton 1982).

Such multi-objective challenges are pervasive across individual and organizational contexts. For instance, employee performance evaluations commonly require individuals to navigate incentive systems that combine multiple metrics—such as sales, customer satisfaction, and teamwork (Ethiraj & Levinthal 2009, Kerr 1975). Firms also face multidimensional objectives when responding to industrial policy incentives that simultaneously target domestic investment, employment, and innovation outcomes (Li & Csaszar 2019). In product development, design teams must iteratively search for configurations that balance competing criteria—performance, cost, time-to-market, and reliability—where each design decision (e.g., choice of materials, component architecture, or manufacturing process) jointly affects multiple outcomes (Vinokurova & Kapoor 2020). A team that evaluates prototypes solely on projected profit margin may converge on a local optimum that a team alternating attention between margin and user experience would have avoided, mirroring the dynamics our model captures. In the healthcare sector, hospitals must balance service pricing decisions to ensure financial viability while simultaneously fulfilling population-health mission commitments (Albert & Eklund 2026). In all these settings, the problem of searching for satisfactory solutions under multiple goals is structurally similar to the one we have modeled, and our framework offers a generalizable approach for analyzing such decision processes.

Our framework also speaks to debates in the New Stakeholder Theory (NST), which examines how stakeholder engagement enables firms to create and appropriate value (Barney 2018, Kaplan 2023, McGahan 2023). NST has been mostly silent, however, on how firms *search* for the strategies

that yield such value creation. Its core normative question—how organizations should aggregate and prioritize heterogeneous stakeholder claims—can be recast in terms of the heuristics we formalize: some approaches rely on cross-dimensional aggregation and trade-offs, while others refuse such trade-offs through priority rules or aspiration-based thresholds, or shift attention sequentially across demands. Yet as Lazzarini (2025) argues, NST has not sufficiently specified the normative assumptions underpinning its theoretical mechanisms, including how “value” is operationalized across stakeholder groups. Our results add a behavioral dimension to this challenge: even when normative commitments are clear, the search heuristic used to implement them steers organizations toward different regions of the performance space. Value creation is therefore not only about what firms *aim* to create or how stakeholders define value, but also about which search procedures enable firms to *discover* strategies that jointly improve multiple dimensions.

Practical Implications

Our research highlights two critical decisions firms must make. First, they need to decide what position they want to occupy on the social–financial frontier. The desired position will depend on managers’ preferences as well as pressures stemming from competitors and stakeholders such as regulators, investors, and customers (Battilana, Obloj, Pache, & Sengul 2022). Hence, different firms may aim for different targets. For instance, a nonprofit may aim for point *Q* in Figure 1 while a firm facing strong competitive and shareholder pressures may aim for point *U*.

Second, once the desired location on the frontier has been determined, the firm needs to pick a heuristic likely to get it there. Our research offers some initial guidance. Maximize may be useful when financial performance is critical for survival and social performance requirements are undeveloped, as in nascent industries and markets that lack regulation and oversight. Penalize is appropriate when financial performance is most important but clearing a self-imposed social threshold is desired. Combine is most useful for firms that aim to jointly pursue social and financial goals and can sustain the commensurability infrastructure required to aggregate them. Alternate is helpful when, apart from jointly pursuing both goals, the firm also wants to foster exploration—including exploring strategies that can ultimately exceed a purely financial searcher’s local peak.

Satisfice could be chosen by firms aspiring to high social performance with only limited financial aims, as could be the case for nongovernmental organizations and certified benefit corporations (Gehman, Grimes, & Cao 2019), and it may be especially practical when fine-grained measurement of social performance is difficult.

We have also found that Figure 1—the social–financial frontier—is a valuable framework in the MBA classroom, as it allows students to think more clearly about several issues related to CSR and multiple stakeholders: (i) that “performance” is not just financial performance, since all strategies also have a social performance; (ii) that trade-offs between social and financial goals are not always unavoidable—they are not when firms are inside the frontier or when new technologies push the frontier; (iii) that firms should aim to be on the frontier; (iv) that not all positions on the frontier are viable (only those that clear the social and financial thresholds); and (v) that picking a target position on the viable section of the frontier is an important choice that depends on the values of the firm’s decision-makers.

Future Work

Our research can be extended in several ways. In terms of empirical work, ethnographic or experimental research methods could be used to (i) surface additional heuristics that firms may use when dealing with multiple goals and (ii) identify the situations in which the different heuristics are more likely to be used. Future empirical work could also test the predictions of our model using a large sample. One way would be to use the composition of the board of directors or the top management team (TMT) as a proxy for the type of heuristics used by the firm. For instance, one could test whether firms with finance-heavy representation in the TMT tend to perform like the Maximize heuristic or whether firms with high-turnover TMTs tend to perform like the Alternate heuristic. The social threshold could be evaluated by measuring the extent to which a company’s ESG goals exceed regulatory standards. Future work could examine the implementation costs of multi-objective heuristics, including their measurement demands.

In terms of theoretical work, future research could (i) study heuristics that deal with more than two performance dimensions, as in designing a product to satisfy many different stakeholders, (ii)

study heuristics that include multiple actors, some of whom may have opposing ideas regarding the firm's proper goals, and (iii) incorporate processes by which firms may actively shape how performance is incentivized and measured. These extensions would bring the model closer to multi-stakeholder settings by decomposing "social performance" into multiple stakeholder-specific dimensions and allowing goal conflict to shape search.

Related to this last point, a promising direction for future research is to connect our findings to the emerging literature on market shaping (Helfat 2021, Vinokurova 2019) by considering how the adoption of particular search heuristics can itself serve as a lever for shaping markets. In multidimensional contexts, the way firms "see" the landscape is mediated by their chosen heuristic—for example, the Combine heuristic aggregates social and financial dimensions, while the Alternate heuristic focuses on one dimension at a time. This implies that shaping is not limited to directly changing the social and financial payoff structures, but can also involve influencing which heuristics firms adopt, whether through regulation, industry standards, or stakeholder pressure. For instance, policy interventions might mandate that firms apply specific weights to different objectives, effectively institutionalizing a preferred heuristic and thereby guiding organizational search and aggregate market outcomes. Exploring how such mechanisms affect the distribution of firms along the social–financial frontier, and the collective achievement of goals like sustainability, offers a fruitful avenue for extending both shaping theory and the study of multi-objective decision processes.

Conclusion

This paper sheds light on the fundamental behavioral challenge of how organizations search for satisficing strategies when multiple objectives matter. We show that some multi-objective search heuristics are much more likely than others to achieve particular stakeholder goals. Choosing a search heuristic is therefore a critical matter. Interestingly, this choice depends mostly on what the firm wants to achieve rather than on characteristics of its environment. To develop these insights, we have extended the search literature to formally study search under multiple goals. We have thus contributed to opening the black box of search strategies for stakeholder value while—in the manner of some of the heuristics we model—pursuing the dual goals of rigor and relevance.

REFERENCES

- Adner, R., Csaszar, F. A., & Zemsky, P. B. 2014. Positioning on a multi-attribute landscape. *Management Science*, 60(11): 2794–2815.
- Aguinis, H., & Glavas, A. 2012. What we know and don't know about corporate social responsibility: A review and research agenda. *Journal of Management*, 38(4): 932–968.
- Albert, D., & Eklund, J. 2026. Margin or mission? Understanding the relationship between multilevel functional subgroups and pricing. *Organization Science*, 37(1): 323–352.
- Allouche, J., & Laroche, P. 2005. A meta-analytical investigation of the relationship between corporate social and financial performance. *Revue de Gestion des Ressources Humaines*, 57: 18–41.
- Barney, J. B. 2018. Why resource-based theory's model of profit appropriation must incorporate a stakeholder perspective. *Strategic Management Journal*, 39(13): 3305–3325.
- Battilana, J., Obloj, T., Pache, A.-C., & Sengul, M. 2022. Beyond shareholder value maximization: Accounting for financial/social tradeoffs in dual-purpose companies. *Academy of Management Review*, 47(2): 237–258.
- Battilana, J., Pache, A.-C., Sengul, M., & Kimsey, M. 2019. The dual-purpose playbook. *Harvard Business Review*, 97(2): 124–133.
- Baumann, O., Schmidt, J., & Stieglitz, N. 2019. Effective search on rugged performance landscapes: A review and outlook. *Journal of Management*, 45(1): 285–318.
- Benoit, D. 2019. Move over, shareholders: Top CEOs say companies have obligations to society. *The Wall Street Journal* (Aug 19).
- Bentham, J. 1789/1996. An introduction to the principles of morals and legislation. Oxford, UK: Oxford University Press.
- Bromiley, P., & Harris, J. D. 2014. A comparison of alternative measures of organizational aspirations. *Strategic Management Journal*, 35(3): 338–357.
- Carruthers, B. G., & Stinchcombe, A. L. 1999. The social structure of liquidity: Flexibility, markets, and states. *Theory and Society*, 28(3): 353–382.
- Crilly, D., & Sloan, P. 2012. Enterprise logic: Explaining corporate attention to stakeholders from the 'inside-out'. *Strategic Management Journal*, 33(10): 1174–1193.
- Csaszar, F. A. 2018. A note on how NK landscapes work. *Journal of Organization Design*, 7(15): 1–6.
- Csaszar, F. A., & Levinthal, D. A. 2016. Mental representation and the discovery of new strategies. *Strategic Management Journal*, 37(10): 2031–2049.
- Cyert, R. M., & March, J. G. 1963. A behavioral theory of the firm. Englewood Cliffs, NJ: Prentice-Hall.
- Dorobantu, S., Hennisz, W. J., & Nartey, L. 2017. Not all sparks light a fire: Stakeholder and shareholder reactions to critical events in contested markets. *Administrative Science Quarterly*, 62(3): 561–597.
- Ehrgott, M. 2005. Multicriteria optimization (2nd ed.). Berlin, Germany: Springer.
- Emmerich, M. T. M., & Deutz, A. H. 2018. A tutorial on multiobjective optimization: Fundamentals and evolutionary methods. *Natural Computing*, 17(3): 585–609.
- Encarnación, J. 1964. Constraints and the firm's utility function. *The Review of Economic Studies*, 31(2): 113–120.
- Eno, B., & Schmidt, P. 2001. Oblique strategies. Boxed set of cards. Available at <https://enoshop.co.uk>.
- Espeland, W. N., & Stevens, M. L. 1998. Commensuration as a social process (Vol. 24).
- Ethiraj, S. K., & Levinthal, D. A. 2004. Modularity and innovation in complex systems. *Management Science*, 50: 159–173.

- Ethiraj, S. K., & Levinthal, D. A. 2009. Hoping for A to Z while rewarding only A: Complex organizations and multiple goals. *Organization Science*, 20(1): 4–21.
- Fang, C., & Kim, J.-H. 2018. The power and limits of modularity: A replication and reconciliation. *Strategic Management Journal*, 39(9): 2547–2565.
- Flammer, C. 2015. Does corporate social responsibility lead to superior financial performance? A regression discontinuity approach. *Management Science*, 61(11): 2549–2568.
- Fleming, L. 2001. Recombinant uncertainty in technological search. *Management Science*, 47(1): 117–132.
- Friedman, M. 1970. A Friedman doctrine: The social responsibility of business is to increase its profits. *The New York Times*(Sep 13).
- Gaba, V., & Greve, H. R. 2019. Safe or profitable? The pursuit of conflicting goals. *Organization Science*, 30(4): 647–667.
- Gallardo-Vázquez, D., Barroso-Méndez, M., Pajuelo-Moreno, M., & Sánchez-Meca, J. 2019. Corporate social responsibility disclosure and performance: A meta-analytic approach. *Sustainability*, 11(4): 1115.
- Gans, J. S., Stern, S., & Wu, J. 2019. Foundations of entrepreneurial strategy. *Strategic Management Journal*, 40(5): 736–756.
- Gehman, J., Grimes, M. G., & Cao, K. 2019. Why we care about certified B corporations: From valuing growth to certifying values practices. *Academy of Management Discoveries*, 5(1): 97–101.
- Ghemawat, P., & Nalebuff, B. 1985. Exit. *RAND Journal of Economics*, 16(2): 184–194.
- Greve, H. R. 2008. A behavioral theory of firm growth: Sequential attention to size and performance goals. *Academy of Management Journal*, 51(3): 476–494.
- Greve, H. R., & Gaba, V. 2020. Performance feedback in organizations and groups. In L. Argote & J. M. Levine (Eds.), *The Oxford handbook of group and organizational learning* (pp. 314–336). New York: Oxford University Press.
- Haanaes, K., Michael, D., Jurgens, J., & Rangan, S. 2013. Making sustainability profitable. *Harvard Business Review*, 91(3): 110–115.
- Halmos, P. R. 1974. *Naive set theory*. New York: Springer-Verlag.
- Hart, O., Shleifer, A., & Vishny, R. W. 1997. The proper scope of government: Theory and an application to prisons. *The Quarterly Journal of Economics*, 112(4): 1127–1161.
- Helfat, C. E. 2021. What does firm shaping of markets really mean? *Strategy Science*, 6(4): 360–370.
- Hörisch, J., Freeman, R. E., & Schaltegger, S. 2014. Applying stakeholder theory in sustainability management. *Organization & Environment*, 27(4): 328–346.
- Jensen, M. C. 2002. Value maximization, stakeholder theory, and the corporate objective function. *Business Ethics Quarterly*, 12(2): 235–256.
- Jensen, M. C., & Meckling, W. H. 1976. Theory of firm: Managerial behavior, agency costs and ownership structure. *Journal of Financial Economics*, 3(4): 305–360.
- Kaplan, S. 2023. The promises and perils of corporate purpose. *Strategy Science*, 8(2): 288–301.
- Karnani, A. 2011. Doing well by doing good: The grand illusion. *California Management Review*, 53(2): 69–86.
- Katila, R., & Ahuja, G. 2002. Something old, something new: A longitudinal study of search behavior and new product introduction. *Academy of Management Journal*, 45(6): 1183–1194.
- Kaul, A., & Luo, J. 2018. An economic case for CSR: The comparative efficiency of for-profit firms in meeting consumer demand for social goods. *Strategic Management Journal*, 39(6): 1650–1677.

- Kay, J. A. 2010. *Obliquity: Why our goals are best achieved indirectly*. London, UK: Profile Books.
- Keeney, R. L., & Raiffa, H. 1976. *Decisions with multiple objectives: Preferences and value tradeoffs*. Cambridge, UK: Wiley.
- Kerr, S. 1975. On the folly of rewarding A, while hoping for B. *Academy of Management*, 18(4): 769–783.
- Kroeger, A., & Weber, C. 2014. Developing a conceptual framework for comparing social value creation. *Academy of Management Review*, 39(4): 513–540.
- Lazzarini, S. G. 2025. The strategic imperative: Do we need normative considerations in strategic theories of stakeholder engagement? *Strategy Science*, 10(4): 300–315.
- Levinthal, D. A. 1997. Adaptation on rugged landscapes. *Management Science*, 43(7): 934–950.
- Levinthal, D. A., & March, J. G. 1993. The myopia of learning. *Strategic Management Journal*, 14: 95–112. Special Issue.
- Li, C., & Csaszar, F. A. 2019. Government as landscape designer: A behavioral view of industrial policy. *Strategy Science*, 4(3): 175–192.
- March, J. G. 1991. Exploration and exploitation in organizational learning. *Organization Science*, 2(1): 71–87.
- March, J. G., & Simon, H. A. 1958/1993. *Organizations*. New York: John Wiley & Co.
- Margolis, J. D., Elfenbein, H. A., & Walsh, J. P. 2009. Does it pay to be good . . . and does it matter? A meta-analysis of the relationship between corporate social and financial performance. Available at SSRN: <https://ssrn.com/abstract=1866371>.
- Margolis, J. D., & Walsh, J. P. 2003. Misery loves companies: Rethinking social initiatives by business. *Administrative Science Quarterly*, 48(2): 268.
- Marquis, C., & Park. 2014. Inside the buy-one give-one model. *Stanford Social Innovation Review*, 12(1): 28–33.
- McGahan, A. M. 2023. The new stakeholder theory on organizational purpose. *Strategy Science*, 8(2): 245–255.
- McGuire, J. B., Sundgren, A., & Schneeweis, T. 1988. Corporate social responsibility and firm financial performance. *Academy of Management Journal*, 31(4): 854–872.
- McWilliams, A., & Siegel, D. 2001. Corporate social responsibility: A theory of the firm perspective. *Academy of Management Review*, 26(1): 117–127.
- Mill, J. S. 1848/2004. *Principles of political economy*. Amherst, NY: Prometheus Books.
- Miller, C. C., Washburn, N. T., & Glick, W. H. 2013. The myth of firm performance. *Organization Science*, 24(3): 948–964.
- Moldau, J. H. 1992. On the lexical ordering of social states according to Rawls' principles of justice. *Economics and Philosophy*, 8(1): 141–148.
- Nickerson, J. A., & Zenger, T. R. 2002. Being efficiently fickle: A dynamic theory of organizational choice. *Organization Science*, 13(5): 547–566.
- Orlitzky, M., Schmidt, F. L., & Rynes, S. L. 2003. Corporate social and financial performance: A meta-analysis. *Organization Studies*, 24(3): 403–441.
- Porter, M. E. 1996. What is strategy? *Harvard Business Review*, 74(6): 61–78.
- Prahalad, C. K., & Bettis, R. A. 1986. The dominant logic: A new linkage between diversity and performance. *Strategic Management Journal*, 7(6): 485–501.
- Rawhouser, H., Cummings, M., & Newbert, S. L. 2017. Social impact measurement: Current approaches and future directions for social entrepreneurship research. *Entrepreneurship Theory and Practice*, 43(1): 82–115.

- Richards, E. 1991. Airing the 'Computer Papers'. The Washington Post, June 23.
- Rivkin, J. W., & Siggelkow, N. 2003. Balancing search and stability: Interdependencies among elements of organizational design. *Management Science*, 49(3): 290–311.
- Rosenkopf, L., & Nerkar, A. 2001. Beyond local search: Boundary-spanning, exploration, and impact in the optical disk industry. *Strategic Management Journal*, 22(4): 287–306.
- Shaked, A., & Sutton, J. 1982. Relaxing price competition through product differentiation. *Review of Economic Studies*, 49: 3–13.
- Shinkle, G. A., Kriauciunas, A. P., & Hundley, G. 2013. Why pure strategies may be wrong for transition economy firms. *Strategic Management Journal*, 34(10): 1244–1254.
- 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.
- Tantalo, C., & Priem, R. L. 2016. Value creation through stakeholder synergy. *Strategic Management Journal*, 37(2): 314–329.
- Vinokurova, N. 2019. Reshaping demand landscapes: How firms change customer preferences to better fit their products. *Strategic Management Journal*, 40(13): 2107–2137.
- Vinokurova, N., & Kapoor, R. 2020. Converting inventions into innovations in large firms: How inventors at Xerox navigated the innovation process to commercialize their ideas. *Strategic Management Journal*, 41(13): 2372–2399.
- Vishwanathan, P., van Oosterhout, H. J., Heugens, P. P., Duran, P., & van Essen, M. 2020. Strategic CSR: A concept building meta-analysis. *Journal of Management Studies*, 57(2): 314–350.
- Wang, Q., Dou, J., & Jia, S. 2016. A meta-analytic review of corporate social responsibility and corporate financial performance. *Business & Society*, 55(8): 1083–1121.

APPENDIX A: THE WEIGHTED COMBINE HEURISTIC

The Weighted Combine heuristic, denoted $\text{Combine}(w)$, searches on the combined landscape $wf(\mathbf{d}) + (1 - w)s(\mathbf{d})$. For example, $\text{Combine}(0.8)$ assigns 80% weight to financial performance and 20% to social performance. $\text{Combine}(0.5)$ is equivalent to the Combine heuristic analyzed in the body of the paper.

Figure A.1 simulates 10,000 runs of $\text{Combine}(w)$ for $w = 0.2, 0.5,$ and 0.8 under the same assumptions used to generate Figure 8. The main effect of w is to tilt outcomes toward the weighted dimension: as w increases from 0 to 1, search shifts from high-social/low-financial performance toward low-social/high-financial performance.

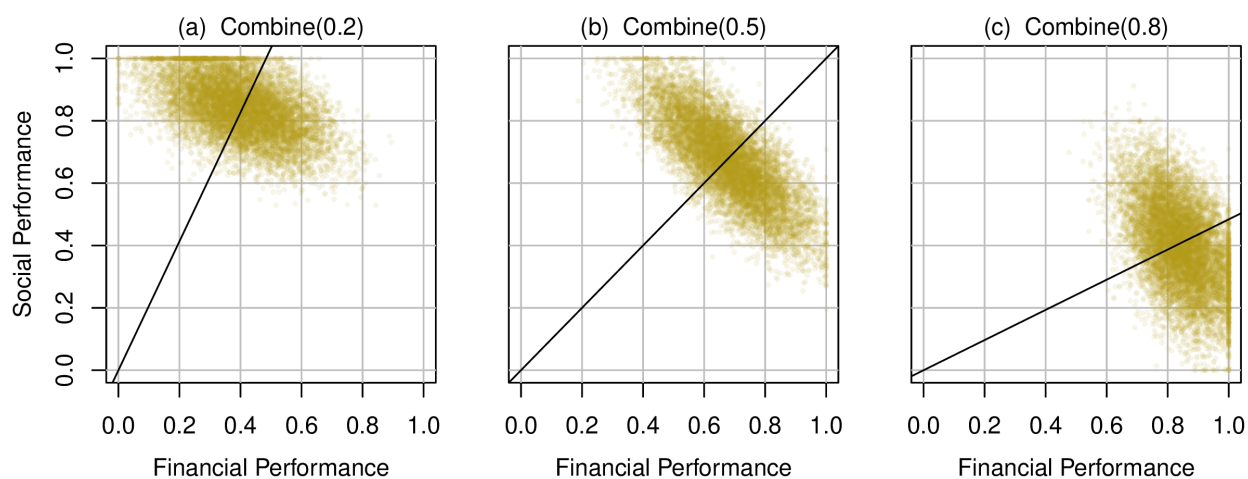


Figure A.1: Illustration of the Weighted Combine heuristic under the same conditions as Figure 8.

APPENDIX B: RELAXING THE LOCAL SEARCH ASSUMPTION

A key assumption in our baseline model is strict local search: firms only evaluate strategies that differ by a one-bit change. To examine robustness to this assumption, we allow firms to follow their focal heuristic most of the time but, with a small probability, make a non-local (random) move.¹³ Occasional non-local moves reduce the extent to which any one heuristic becomes “trapped” at a particular local peak, thereby dampening (though not eliminating) differences that arise purely from

¹³Conceptually, such non-local moves capture episodes of more exploratory variation (e.g., experimentation with distant projects, acquisitions, or major strategic reorientations) that can arise under conditions such as crisis, ownership change, or active investor intervention. This is consistent with classic arguments that organizational adaptation reflects a shifting balance between exploration and exploitation (see, e.g., March 1991).




Figure B.1: Performance plots for different exploration probabilities (R).

locality constraints.

We model exploration via a parameter R , which denotes the probability that the firm performs a “random jump” rather than a step of local search in a given time period. Thus, $R = 0$ corresponds to the analyses reported in the body of the paper (i.e., full local search), while, for example, $R = 0.2$ corresponds to a 20% probability of performing a random jump per period. A random jump is modeled as drawing a new decision vector by assigning each of the N binary decisions in \mathbf{d} at random (consistent with Levinthal 1997: 938). Figure B.1 uses the same K , ρ , and H values as panel (d) in Figure 6 while varying the probability of performing a random jump ($R = 0, 0.2$, and 0.4).

Several patterns emerge from Figure B.1. First, the qualitative ranking of heuristics is preserved across all three panels: Maximize attains the highest financial performance but the lowest social performance, while Combine and Alternate remain the strongest on the social dimension. Second, Maximize falls below the social threshold $H = 0.5$ in every panel, indicating that random jumps alone do not remedy the social-performance deficit of a purely financial focus. Third, the spread among heuristics shrinks as R grows: random jumps dilute each heuristic’s local search logic and compress outcomes toward one another (i.e., closer to random performance $(0.5, 0.5)$). At $R = 0.4$, Combine, Alternate, and Satisfice are nearly indistinguishable, since their distinct local search patterns are increasingly overridden by random exploration. Overall, relaxing strict local search dampens but does not eliminate the differences among heuristics.