



This postprint was originally published by the American Economic Association (AEA) as:

Artinger, F. M., Gigerenzer, G., & Jacobs, P. (2022). **Satisficing: Integrating two traditions.** *Journal of Economic Literature*, 60(2), 598–635. <https://doi.org/10.1257/jel.20201396>.

Nutzungsbedingungen:

Dieser Text wird unter einer Deposit-Lizenz (Keine Weiterverbreitung - keine Bearbeitung) zur Verfügung gestellt. Gewährt wird ein nicht exklusives, nicht übertragbares, persönliches und beschränktes Recht auf Nutzung dieses Dokuments. Dieses Dokument ist ausschließlich für den persönlichen, nicht-kommerziellen Gebrauch bestimmt. Auf sämtlichen Kopien dieses Dokuments müssen alle Urheberrechtshinweise und sonstigen Hinweise auf gesetzlichen Schutz beibehalten werden. Sie dürfen dieses Dokument nicht in irgendeiner Weise abändern, noch dürfen Sie dieses Dokument für öffentliche oder kommerzielle Zwecke vervielfältigen, öffentlich ausstellen, aufführen, vertreiben oder anderweitig nutzen. Mit der Verwendung dieses Dokuments erkennen Sie die Nutzungsbedingungen an.

Terms of use:

This document is made available under Deposit Licence (No Redistribution - no modifications). We grant a non-exclusive, nontransferable, individual and limited right to using this document. This document is solely intended for your personal, non-commercial use. All of the copies of this documents must retain all copyright information and other information regarding legal protection. You are not allowed to alter this document in any way, to copy it for public or commercial purposes, to exhibit the document in public, to perform, distribute or otherwise use the document in public. By using this particular document, you accept the above-stated conditions of use.

Provided by:

Max Planck Institute for Human Development
Library and Research Information
library@mpib-berlin.mpg.de

Satisficing: Integrating Two Traditions[†]

Florian M. Artinger, Gerd Gigerenzer, and Perke Jacobs*

In 1955, Herbert Simon introduced the notion of satisficing: an agent satisfices by searching for an alternative that meets an aspiration level but does not optimize. We survey more than 60 years of advances in understanding satisficing in economics, psychology, and management, identifying two research traditions that address two classes of situations: under risk, satisficing is typically inferior to optimization strategies and modeled according to the neoclassical framework; under uncertainty, satisficing strategies are often derived empirically and can be highly effective. We integrate the two research traditions and show the conditions under which satisficing can be rational. (JEL D11, D80, D90)

1. Introduction

Ever since Simon (1955) initiated the behavioral revolution in economics, its poster child has been satisficing. Satisficing refers to the observation that agents make choices with the help of aspiration levels that do not necessarily coincide with utility maximization. The normative appeal of utility maximization has led many to dismiss satisficing uniformly as an undesirable quirk of human behavior. In this article, we distinguish two separate research traditions that can be traced back to Simon's (1955) original visions of satisficing but are largely disconnected today. Reviewing both traditions, we show how they can be integrated within a framework for understanding decision-making beyond utility maximization.

Specifically, we argue that the rationality of satisficing strategies depends on the class of decision environment. Broadly, environments can be divided into two classes. Under risk, optimization sets the rational benchmark and satisficing can yield suboptimal decisions. Under uncertainty or intractability, where the optimal action cannot be determined, satisficing can outperform complex strategies, including rational choice models. Satisficing has been examined in both types of decision environments, resulting in two distinct and largely unconnected literatures.

In section 2, we trace the historical trajectory of the notion of satisficing and provide a conceptual overview of the meaning of the term. In section 3, we provide a review of the two traditions in research on satisficing

*Artinger: Max Planck Institute for Human Development, Simply Rational—The Decision Institute, and Berlin International University of Applied Sciences. Gigerenzer: Max Planck Institute for Human Development, Simply Rational—The Decision Institute, and Harding Center for Risk Literacy. Jacobs: Max Planck Institute for Human Development. The authors thank Konstantinos Katsikopoulos, Reinhard Selten, Shenghua Luan, and the ABC Research Group for comments on an earlier version of this text, and Rona Unrau for copyediting the manuscript.

[†] Go to <https://doi.org/10.1257/jel.20201396> to visit the article page and view author disclosure statement(s).

that have evolved. In section 4, we propose a unifying framework to better understand when and why satisficing can be rational and what that means. Section 5 closes with four methodological conclusions, advocating competitive out-of-sample tests to evaluate decision strategies under uncertainty and intractability.

2. Satisficing

2.1 Historical Context

In his seminal contribution to economics, Simon advocated and developed a model of bounded rationality, positing that “the task is to replace the global rationality of economic man with a kind of rational behavior that is compatible with the access to information and computational capacities that are actually possessed by organisms, including man, in the kinds of environments in which such organisms exist” (Simon 1955, p. 99). The period in which this paper was published was characterized by the popularization of neoclassical rational choice theory, or global rationality, as Simon referred to it. This body of theories includes von Neumann and Morgenstern’s (1944) and Savage’s (1954) work on expected utility theory and the work by Nash (1950) on equilibrium in noncooperative games. Common to these theories is the assumption of an agent who has complete information about the available alternatives, including perfect foresight about all possible consequences and sufficient knowledge of the probabilities with which they occur. Agents are then predicted to act as if they were solving an optimization problem to maximize expected utility. Friedman (1953, p. 15) maintained that such an assumption is justified, irrespective of whether it is deemed realistic, because “the relevant question to ask about the assumptions of a theory is not whether they are descriptively realistic, for they never are, but whether they are sufficiently good approximations for the purpose in hand. And this question can be answered only by seeing whether the theory works, which means whether it yields sufficiently accurate predictions.” Compared with this approach, Simon’s differed in two respects: First, his interest was in models of agents’ actual decision processes, not only of their outcomes. Second, he was interested in situations “where the conditions for rationality postulated by the model of neoclassical economics are not met” (Simon 1989, p. 377).

The early responses in economics to Simon’s writings were twofold. On the one hand, the proposition of bounded rationality evoked vigorous defenses of rational choice theory.¹ At the same time, many economists were somewhat open to the notion of bounded rationality. This group included Robert Solow, who—reviewing Simon’s (1957) book *Models of Man: Social and Rational*, which expands on his 1955 paper—found himself “torn between an impulse to display the interdisciplinary scope of my ignorance by commenting on every essay and a more rational disposition to fight it out along the main line” (Solow 1958, p. 81). Although the book received very positive reviews by some economists (e.g., Shubik 1958), Simon’s departure from the neoclassical economic canon presumably made it difficult for many economists to incorporate his ideas into their theorizing.

For the years up to 1969, we found only 35 citations of Simon’s 1955 article on Web of Science. Only during the 1970s did his early contributions start to gain recognition in the economic literature. Inspired by Simon being awarded the 1978 Nobel Memorial Prize in Economic Sciences, a community of economists and psychologists dedicated their

¹ For an overview of the arguments put forth in favor of rational choice theory over the years as well as the counterarguments, see Conlisk (1996).

work to studying the behavioral foundations of economic theory, which developed into behavioral economics. Today, Simon's work is often cited as the predecessor of Tversky and Kahneman's (1974) work on heuristics and biases. However, not until 1981 did Tversky and Kahneman begin to relate their work to Simon's study of bounded rationality (Gigerenzer 2004). Yet bounded rationality does not mean the same in both programs. To Simon, it meant the study of behavior in situations where the conditions assumed in neoclassical economics are not met, whereas Kahneman and Tversky assumed that these conditions are met and that deviating behavior implies a lack of rationality. Simon (1985, p. 297) made this difference between them clear: "Bounded rationality is not irrationality."

At the same time, both approaches to behavioral economics can be characterized as empirically falsifying the assumptions underlying neoclassical economic theory (for complementary reviews on the topic, see Harstad and Selten 2013, Crawford 2013, Rabin 2013). Unlike Simon, however, the heuristics-and-biases program attributed behavioral deviations from neoclassical theory to flaws in people's minds rather than to potential flaws in the application of the theory. This allowed contemporary behavioral economics to retain the underlying norm of an agent who integrates all information and maximizes utility. Simon's writings, in contrast, were followed by research that studied decisions beyond the domain of rational choice theory, including the work by Cyert and March (1963) on the behavioral theory of the firm, Winter's (1971) work on evolutionary economics, and the work by Gigerenzer and colleagues on fast and frugal heuristics (Gigerenzer, Hertwig, and Pachur 2011; Gigerenzer and Selten 2001). These analyses are based on a satisficing agent and study the decision processes, routines, and rules of thumb that agents and organizations actually use when facing complex and dynamic environments that provide only limited information.

Winter (1971) provides a quote of Simon as a central source for his own inspiration:

The equilibrium behavior of a perfectly adapting organism depends only on its goals and its environment; it is otherwise completely independent of the internal properties of the organism (...). [T]o predict the short-run behavior of an adaptive organism, or its behavior in a complex and rapidly changing environment, it is not enough to know its goals. We must also know a great deal about its internal structure and particularly its mechanisms of adaptation. (Simon 1959, p. 255)

That is, equilibrium strategies derived from a stylized representation of the world, specifically its incentive structure, can substantially differ from the strategies that agents actually use to navigate an uncertain and complex world. Going back to Smith (1962), there is a substantial literature in economics demonstrating that equilibrium also obtains with naive, merely privately informed agents (for a review, see Smith 2008). Gode and Sunder (1993) even find that zero-intelligence traders, who randomize within their budget constraints, produce allocative efficiency. Much of the work in behavioral economics does not make a clear distinction between the individual and the aggregate levels of analysis.

In order to account for differences between an equilibrium perspective and the actual behavior of an agent, Smith (2008) proposes a distinction between two types of analyses.² The first, *constructivist rationality*, applies deductive reasoning from first principles: it identifies the incentive structure

² The principal distinction between two such rational orders can already be found in the writings of Adam Smith (1776 [1976], 1759 [1981]), Hume (1739), and later Hayek (1937, 1945), as well as Savage (1954) and Simon (1955, 1956).

and deduces the equilibrium by sufficiently abstracting and simplifying. In contrast, an analysis of *ecological rationality* proceeds empirically by determining the decision strategies used by agents and then evaluating the performance of that strategy competitively against other relevant strategies in the given context. The term *ecological rationality* thereby refers to the degree to which a strategy is adapted to the environment, evaluated in terms of a fitness measure such as profit or accuracy of predictions (Gigerenzer, Todd, and ABC Research Group 1999).

In the present article, we examine satisficing through the lens of ecological rationality. This perspective, yet uncommon in economics, offers a framework for thinking about decision strategies in a broader way. As we will argue, it explains how Simon's early writings inspired two largely distinct research traditions. Because this perspective examines strategies relative to the environment, we include a short primer on different degrees of uncertainty that lead to fundamentally different classes of decision environments.

2.2 Risk, Ambiguity, Intractability, and Uncertainty

Keynes (1921) and Knight (1921) use a dichotomy of two broad categories of environments, both of which are characterized by the absence of certainty: Risk and uncertainty. Whereas risk is commonly understood, uncertainty has been assigned different meanings. In order to define those relevant for this article, we begin with the terminology of Savage (1954) developed in *The Foundations of Statistics*, in which he axiomatized subjective expected utility theory. Building on this terminology allows us to offer a more detailed definition of different kinds of uncertainty (see also table 1).

Savage defines a decision problem as a pair $\{S, C\}$, where S is the exhaustive and mutually exclusive set of all future states of the world and C the exhaustive set of their consequences associated with each alternative. The alternatives or actions are defined on $\{S, C\}$, and each state s in S has an assigned probability. Choice under certainty means that for each alternative, there is only one state with probability 1; all others have probabilities 0. Choice under risk means that more than one state has nonzero probability and that the probabilities attached to each state are known (table 1, top); the expected utility of an alternative is the sum of the consequences multiplied by their respective probabilities over all possible states. A situation of ambiguity is identical to this, apart from the probability distribution not being known, as in the gambles underlying the Ellsberg paradox (table 1, middle; Ellsberg 1961, see also Anscombe and Aumann 1963). What these three situations—certainty, risk, and ambiguity—have in common is that the complete set of alternatives, future states of the world, and consequences is known. Such problems are said to be *well-defined*.

A well-defined problem can be tractable or not. Any decision problem under certainty, risk, and ambiguity is considered computationally intractable (with subdivisions into NP-hard, NP-complete, etc.) if the set of alternatives or states is so large that the best one cannot be identified by mind or machine. This means that no efficient (i.e., polynomial-time) algorithm exists to solve it (e.g., Garey and Johnson 1979). Examples include games such as chess and Go. To understand the order of magnitude of this limitation, note that chess has approximately 10^{120} unique sequences of moves or games, a number greater than the estimated number of atoms in the universe (Shannon 1950). Many important tasks are intractable, including scheduling, capital budgeting, and itinerary problems, among others (Papadimitriou and Steiglitz 1998, Markose 2005). Savage (1954, p. 16) is explicit that intractable problems are outside of the

TABLE 1
TYPES OF ENVIRONMENTS

Decisions Under Risk					
Alternative	State 1 $p = p_1$	State 2 $p = p_2$	State 3 $p = p_3$...	State M $p = p_M$
a_1	$c_{1,1}$	$c_{1,2}$	$c_{1,3}$...	$c_{1,M}$
a_2	$c_{2,1}$	$c_{2,2}$	$c_{2,3}$...	$c_{2,M}$
a_3	$c_{3,1}$	$c_{3,2}$	$c_{3,3}$...	$c_{3,M}$
\vdots	\vdots	\vdots	\vdots	\ddots	\vdots
a_N	$c_{N,1}$	$c_{N,2}$	$c_{N,3}$...	$c_{N,M}$

Decisions Under Ambiguity					
Alternative	State 1 $p = ??$	State 2 $p = ??$	State 3 $p = ??$...	State M $p = ??$
a_1	$c_{1,1}$	$c_{1,2}$	$c_{1,3}$...	$c_{1,M}$
a_2	$c_{2,1}$	$c_{2,2}$	$c_{2,3}$...	$c_{2,M}$
a_3	$c_{3,1}$	$c_{3,2}$	$c_{3,3}$...	$c_{3,M}$
\vdots	\vdots	\vdots	\vdots	\ddots	\vdots
a_N	$c_{N,1}$	$c_{N,2}$	$c_{N,3}$...	$c_{N,M}$

Decisions Under Uncertainty					
Alternative	State 1 $p = ??$	State 2 $p = ??$	State 3 $p = ??$...	$??$ $p = ??$
a_1	$c_{1,1}$	$c_{1,2}$	$c_{1,3}$...	$??$
a_2	$c_{2,1}$	$c_{2,2}$	$c_{2,3}$...	$??$
a_3	$c_{3,1}$	$c_{3,2}$	$c_{3,3}$...	$??$
\vdots	\vdots	\vdots	\vdots	\ddots	\vdots
$??$	$??$	$??$	$??$...	$??$

domain of expected utility theory. Solving intractable problems requires a different kind of decision theory, in both descriptive and normative terms, that includes heuristic strategies (Gigerenzer and Selten 2001).

The final class of problems involves degrees of what we call uncertainty, which has elsewhere been termed radical or fundamental uncertainty (Kay and King 2020). In economics, the terms ambiguity and uncertainty are commonly used interchangeably (Etner, Jeleva, and Tallon 2012). However, the distinction between them is fundamental. Ambiguity means that a problem is well-defined, that is, the exhaustive set of alternatives, possible states, and their consequences is known. Uncertainty, in contrast, means that the problem is ill-defined, that the exhaustive set of states of the world and their consequences is not knowable or foreseeable at the point of decision-making (table 1, bottom). Savage (1954, p. 16) lists as an example planning a picnic, where events can occur that one cannot know ahead of

time. He points out that expected utility theory cannot and should not be applied under uncertainty.

One of the contributions of this article is to relate these classes of situations to satisficing. In section 3, we will show that there are two different traditions of satisficing, one assuming well-defined situations such as risk and the other addressing situations of uncertainty and intractability. First, however, we define the basic concepts of satisficing.

2.3 Satisficing: Definition

To illustrate his vision of bounded rationality, Simon spends a good portion of his landmark 1955 article describing a satisficing decision strategy. Such a strategy, he posits, is more descriptive of human decision processes than the traditional model of rational choice, for which he sees “a complete lack evidence that, in actual human choice situations of any complexity, these computations can be, or are in fact, performed” (Simon 1955, p. 104). His alternative model consists of two elements that are characteristic for choice processes: (i) the aspiration level, which is a simplified value function that can be adapted over time, and (ii) search.

The first element, a direct simplification of neoclassical theory, is perhaps the most controversial element of his proposal. Consistent with our earlier notation, Simon (1955, p. 104) suggests that “[o]ne route to simplification is to assume that [the value function] $V(s)$ necessarily assumes one of two values, $(1; 0)$, or one of three values, $(1; 0; -1)$, for all s in S . Depending on the circumstance, we might want to interpret these values, as (a) (satisfactory or unsatisfactory), or (b) (win, draw, or lose).” A binary value function implies the use of an aspiration level, which refers to the minimum level of a given scale of interest that is acceptable to the agent in order to choose an alternative. The term aspiration level has a long tradition in psychological theory (Gardner 1940), appearing first in German (as *Anspruchsniveau*) in work by Dembo (1931) and then in English (as level of aspiration) in work by Lewin (1935). Although Simon first defines the aspiration level in terms of the value or utility of an alternative, he later maintains that it is more realistically set separately for each attribute under consideration (see Simon, 1955, pp. 109–10). When choosing among different houses, for example, agents may have multiple aspiration levels, one for each attribute such as price, floor size, location, or number of rooms, etc., instead of one aspiration level for the overall utility of each alternative. In this view, agents are not assumed to integrate different attributes on a cardinal scale but instead to evaluate each attribute separately. This explicitly allows for incommensurability, the proverbial comparison of apples and oranges where one cannot readily compare two attributes on the same scale.

Aspiration levels do not necessarily remain constant over time. After one or more alternatives are encountered that do not meet the aspiration level, the agent may adjust the level and then proceed to examine the next alternative until one is encountered that fulfills the most recent aspiration level. For an illustration, consider figure 1, which extends figure 2 of Simon (1955). Aspiration levels for alternatives a_1, a_2, a_3 with two attributes, x_1 and x_2 , are initially set at A_1 and A_1 , which defines the set of acceptable alternatives (shaded area). At $t = t_1$, A_1 is lowered to A_1' , while A_2 remains. This adjustment redefines the set of acceptable options from the shaded area to the shaded and dotted areas. Alternative a_1 lies below both aspiration levels and is never chosen, whereas a_3 is chosen at all times and a_2 is chosen only after t_1 .

The division into acceptable and unacceptable alternatives could be interpreted as a choice of a less-than-optimal alternative. Simon (1955) addresses such concerns by pointing out that the dichotomy of the value function can reflect preferences (e.g., as an




Figure 1. Aspiration-Level Adaptation

Notes: Aspiration levels for alternatives a_1 , a_2 , a_3 with two attributes, x_1 and x_2 , initially set at A_1 and A_2 , which defines the set of acceptable options (shaded area). At $t = t_1$, A_1 is lowered to A_1' while A_2 remains. The change redefines the set of acceptable options (shaded and dotted areas). Alternative a_1 lies below both aspiration levels and is never chosen, whereas a_3 is always chosen and a_2 is chosen only after t_1 .

approximation to a value function with sufficiently strong decreasing marginal returns), but can also be considered an element to navigate the environment, for example, in situations where alternatives are decided upon sequentially. Note that this implies that the alternative chosen is the best so far, a phrase that Simon later says is a better summary of the satisficing idea than the common notion of an alternative that is good enough (Gigerenzer 2004).

The second element of Simon's satisficing assumes that the agent has to search for information. Simon points to two possibilities for exploration. First, the agent is unaware of the complete set of alternatives and discovers alternatives sequentially. Second, the agent is aware of all alternatives but is agnostic about the complete set of states of the world that each alternative entails. The agent can search either externally by acquiring additional pieces of information or internally by picturing the consequences of different alternatives. In both cases, information gathering leads agents to sequentially explore alternative-state combinations.

Both elements of satisficing, aspiration level and search, jointly define the decision process in situations where the agent cannot fully explore the space spanned by alternatives and states.³ The aspiration level governs both the quality of the alternative chosen and the duration of the search process—even in situations where the cost of

³ Simon provides the example of a chess player pondering the next move by simulating internally how the game would continue under different alternatives until a move is found that clearly leads to a winning position. He points out that this particular task is only manageable by using a satisficing strategy, as it reduces the space of alternatives from an estimated 10^{24} to fewer than 100 moves to be considered.

gathering information is not known to the agent. Setting the aspiration level is therefore at the heart of a satisficing model.

Generally, models of aspiration level setting and adaptation can take many forms, ranging from computationally intensive methods, including Bayesian approaches, to simple ones. Although Simon himself derived a computationally intensive method in the appendix of his 1955 paper (see also Wall 1993, Gilboa and Schmeidler 2001b), he notes that such a method is psychologically implausible given agents' typical lack of necessary information and the complexity of the models. Instead, he posits that an agent "will set his acceptance [level] quite high, watch the distribution of offers he receives, and gradually and approximately adjust his acceptance [level] downward or upward until he receives an offer he accepts—without ever making probability calculations" (Simon 1955, p. 117). Simon's basic model of aspiration-level adaptation can be summarized as follows:

Step 1: Set an aspiration level.

Step 2: Continue search until finding the first alternative that meets or exceeds the aspiration level.

Step 3: If no alternative meets the aspiration level within a fixed period, adapt it by a particular value and return to Step 2.

In his article of 1955, Simon leaves this class of models unnamed and only later introduces the term satisficing in a follow-up article in *Psychological Review* (Simon 1956). In that article, he shows how an organism with very limited cognitive capabilities can successfully apply a well-adapted satisficing strategy to maintain its subsistence. Although the organism in the article of 1956 relies solely on a satisficing strategy, Simon highlights that the satisficing model is but one strategy for many realistic decision situations; other situations may call for different strategies (Simon 1955, p. 104).

3. Two Traditions of Satisficing

Ever since Simon's inception of the term satisficing, it has been used widely and diversely in the literature, which has produced a fragmented understanding of it. In its most general sense, satisficing is defined simply as the antithesis of optimization, without imposing additional constraints on the decision model. By this definition, any decision model that does not rely on optimization techniques satisfies. We could not find this interpretation of satisficing in Simon's early writings, but it emerged later (e.g., Simon 1979), after the term had become emblematic of his more general critique of rational choice theory.

We argue that satisficing has been understood in two different ways, each of which originates in the writings of Simon. Both research traditions make use of aspiration levels in their decision models, albeit in different ways. The key difference between these two traditions lies in the different decision environments they assume: whereas one tradition is primarily concerned with decisions under risk, with and without search, the other tradition examines search-based decisions under uncertainty and intractability. In this section we give an overview of both these traditions and their relevant literature, and selectively highlight some of the central contributions.

3.1 Satisficing under Risk

The first of these research traditions examines satisficing under risk, where all information is available. On the basis of Simon's conception of satisficing above, we make a distinction between models that focus on aspiration levels only and those that model both aspiration levels and search. Although

TABLE 2
MODELS OF SATISFICING WITHOUT SEARCH

Approach	Reference	Keyword
Utility Function	Charnes and Cooper (1963)	binary utility
	Borch (1968)	binary utility
	Fishburn (1977)	risk preferences
	Payne, Laughhunn, and Crum (1980, 1981)	risk preferences
	Brown and Sim (2009)	risk preferences
	Brown, De Giorgi, and Sim (2012)	risk preferences
	Diecidue and van de Ven (2008)	augmented concave utility
	Kahneman and Tversky (1979)	utility with reference point
Preference Orders	Kőszegi and Rabin (2006)	utility with reference point
	Jamison and Lau (1973)	semi-orders
	Krishnan (1977)	semi-orders
	Lioukas (1984)	semi-orders
	Rubinstein and Salant (2006)	semi-orders
	van Rooij (2011)	semi-orders
Choice Rules	Aleskerov, Bouyssou, and Monjardet (2007)	interval orders
	Manzini, Mariotti, and Tyson (2013)	lexicographic semi-orders
Strategic Interaction	Wierzbicki (1982)	principal-agent problem
	Haller (1985)	principal-agent problem
	Pazgal (1997)	cooperation
	Oechssler (2002)	cooperation
	Güth, Levati, and Ploner 2010; Güth 2010	bargaining games
	Papi (2012b, 2013, 2018)	consumer choice in monopoly

models in the former group lack one of our defining elements, they represent a widespread interpretation of the term.

3.1.1 Models without Search

We begin with the group of static models, which is rooted in Simon's proposition of a simplified binary value function. However, over time more diverse implementations have emerged; table 2 classifies this diverse set of models into broad categories.

One straightforward means of introducing aspiration levels into the neoclassical framework has been to modify the utility function. In place of a standard Bernoulli function (Bernoulli 1954) of concave curvature, utility functions are modified to accommodate aspiration levels. Borch (1968) was one of the first to put forth a model of a decision maker seeking to minimize the probability of bankruptcy. A decision maker in this model seeks to have positive wealth and chooses courses of action that jointly maximize the probability of achieving this outcome. Such a maximization is computationally equivalent to maximizing Simon's (1955) step-utility function, which assumes a value of zero for all outcomes below the aspiration level and a value of one for all outcomes at or above the aspiration level.

The problem with models of successful probability maximization lies in their coarseness. Assuming that each alternative exceeding the aspiration level yields the same amount of utility appears somewhat counter-intuitive. One attempt to overcome this issue

is presented by Diecidue and van de Ven (2008). Rather than introducing a step utility function, their model maintains and augments a concave utility function: The value of an alternative is described by the sum of the expected utility of the alternative, its probability of success, and its probability of failure. Formally, the value of alternative a is given by

$$(1) \quad V(a) = \sum_{e=1}^E p_e u(c_e) + \mu P(c^+) - \lambda P(c^-),$$

where c_e denotes the payoff conditional on event e that occurs with probability p_e , c^+ and c^- denote the set of payoffs above and below a specified aspiration level, respectively, and μ and λ are constant behavioral quantities. Here, c^+ and c^- are defined by an aspiration level.

Whereas the models by Diecidue and van der Ven use expected utility theory as a starting point, prospect theory (Kahneman and Tversky 1979, Tversky and Kahneman 1992) represents a more radical departure from traditional economic theory. Importantly, prospect theory uses a reference point that separates the domain of gains from that of losses. Gains are valued along a conventional concave function, and losses are valued along a convex function that is steeper in slope than the gain function. The reference point has been interpreted as an aspiration level (e.g., Heath, Larrick, and Wu 1999); exceeding it yields returns in a conventional concave fashion, whereas falling short of it yields disproportionately negative returns.

An alternative approach to modeling satisficing modifies the preference relation underlying the utility function. In this literature, satisficing is most commonly understood in terms of semi-orders (Luce 1956), a class of preference orderings that allow for intransitive indifference relations of the following type: $a_1 \sim a_2$; $a_2 \sim a_3$; $a_1 \succ a_3$. That is, an agent is indifferent between a_1 and a_2 and between a_2 and a_3 in paired comparisons, albeit preferring a_1 over a_3 . This intransitive relationship is implied by the differential threshold of vision or touch formulated in Weber's law, $jnd = \frac{\delta v}{v}$, where jnd is the just-noticeable difference, v the value of the stimuli, and δv the change in the stimuli. Similarly, Luce (1956) contends that such indifference relations occur when agents are only able to distinguish alternatives that are sufficiently distinct and shows that semi-orders are consistent with maximizing a generalized form of utility function.

Van Rooij (2011) argues that satisficing gives rise to a preference semi-order, resulting from agents' inability to distinguish alternatives above the aspiration level. Like a binary utility function, this interpretation of satisficing emphasizes the perceived equivalence of different alternatives. However, as Tversky (1969) points out, a semi-order can result from a lexicographic choice rule. By this choice rule, attributes are examined in a fixed order: Initially, alternatives are ranked according to the first attribute; only if they are too similar is the second attribute considered. Similarity is usually assessed using a threshold that can be interpreted as an aspiration level. For instance, Manzini and Mariotti (2007, 2012) develop axiomatic characterizations of choice data that are consistent with the use of lexicographic choice rules. Using this framework, Manzini, Mariotti and Tyson (2013) characterize a specific lexicographic procedure in which a satisficing strategy is applied at the first stage, followed by a maximization procedure on the selected subset if no unique solution is found beforehand.

Overall, static models of satisficing under risk use expected utility theory as a starting point and modify it to incorporate an aspiration level and make the theory more consistent with observed behavior. Their similarity to expected utility theory enables satisficing

to be contrasted with utility maximization, where satisficing is often understood as perceived equivalence of two alternatives that objectively differ in quality. In this modeling approach, satisficing is considered a deviation from rational choice.

3.1.2 Models of Search

The second branch of satisficing under risk is characterized by the use of aspiration levels in the context of search. Here, an agent is initially unaware of the full set of available alternatives and needs to explore them sequentially. Depending on the search problem, the agent can either recall alternatives discovered earlier or only select the alternative that was last examined. Even if earlier alternatives can be chosen, search may be costly, resulting in a trade-off between investing in further exploration and exploiting current knowledge. Given limited resources such as time, it can be advantageous to terminate search before exploring all available alternatives (Sims 2003, Reis 2006, Gabaix 2014, Caplin and Dean 2015).

Risky search implies that the agent may be unaware of the available alternatives but has meta-information, for example, regarding their distribution or the cost of search. The first model of search under risk was developed by Simon. In the appendix of his 1955 paper, firmly within the rational choice tradition, he develops an optimal search model that relies on an aspiration level for selling a house. Each day the agent receives a price offer from a known distribution. The agent sets the reservation price, or aspiration level, such that it maximizes the expected value. Following this example, Stigler (1961) popularized the topic of search in economics, emphasizing optimal search. At the same time the topic of search rose to prominence in statistics with the theory of optimal stopping (DeGroot 1970). A common finding in this literature is that optimally behaving agents should continue search until finding an alternative that meets a fixed utility threshold, or aspiration level, similar to Wald's (1950) approach to statistical inference. During the 1970s, this sequential paradigm was also adopted by economists who then used reservation prices as aspiration levels to characterize optimal search (e.g., Telser 1973, Rothschild 1974). Since then, models of optimal search under a range of assumptions made aspiration-based stopping rules a tradition in economic and statistical theory, staying within the tradition of expected utility theory (e.g., Gilboa and Schmeidler 1995, 2001a, Rubinstein and Salant 2006).

In search problems without recall, agents can observe the value of an alternative directly but alternatives are only available sequentially. A classic problem in this literature is the secretary problem, where the goal is to choose the best alternative from a random sequence of which only the most recently seen alternative is available for choice. As Gilbert and Mosteller (1966) demonstrate analytically, the optimal strategy is a satisficing strategy. When each alternative is described solely by its rank within the observed sample, the chances of choosing the best alternative are maximized when the agent examines the first $1/e \approx 37$ percent of the sequence, uses the best alternative encountered so far as an (implicit) aspiration level, and then selects the first alternative exceeding this aspiration level. Gilbert and Mosteller (1966) describe a computationally more intensive satisficing strategy for maximizing the probability of choosing the best alternative. Abstracting from the classical secretary problem, Todd and Miller (1999) introduce additional goals beyond the probability of finding the best alternative, such as maximizing the expected value of the chosen secretary. According to their results, achieving these goals requires shorter search than would be necessary to maximize the probability of finding the best alternative. This divergence in goals, they argue, may

explain the finding that experimental participants search less than necessary to find the best alternative. Most of the search models presented here are derived deductively. They are obtained from the properties of the decision problem by determining the optimal decision strategy that achieves a given goal. Notably, these optimal responses often take the form of a satisficing model.

When studying whether people use an aspiration level, the empirical challenge is that it is not sufficient to rely merely on observed outcomes, as Friedman (1953) postulates. This challenge provides the motivation for studying the decision process in terms of (i) the search process identifying the information sequentially inspected by the agent, (ii) the stopping rule specifying the aspiration level that terminates search, (iii) and the decision strategy specifying how the agent derives the decision from the information inspected (Gigerenzer and Goldstein 1996; Handel and Schwartzstein 2018).

Caplin, Dean, and Martin (2011) are among the first to empirically demonstrate the use of aspiration levels akin to Simon (1955) in an incentivized experiment. Participants need to infer the values of the alternatives based on attributes represented by positive or negative numerical values, facilitating commensurability, and to indicate their preferred alternative at any given moment. The authors show that a satisficing model best describes behavior: Participants switch from lower- to higher-value alternatives, indicating that information is being absorbed on an item-by-item basis. Search stops when participants encounter an alternative that exceeds their aspiration level.

The empirical evidence consistently shows that participants do not choose the best alternative that would be possible in the event of omniscience. However, if people have the opportunity to learn through experience, they are able to approximate an optimal stopping rule (Hey, Permana, and Rochanahastin 2017; Manski 2017; Goldstein et al. 2020), which includes the response time when searching internally (Navarro-Martinez et al. 2018). Taking the sequential nature of search into account, participants generally choose the best alternative among those observed (Bearden and Connolly 2007; Caplin, Dean, and Martin 2011; Caplin and Dean 2011; Reutskaja et al. 2011), thereby meeting the requirement of rational choice theory for sequential search (Caplin and Dean 2011). That is, people do indeed choose the “best so far.”

3.2 Satisficing under Uncertainty and Intractability

As noted in the previous section, Simon (1955) was the first to develop an optimal search model that uses an aspiration level. Yet he suggests that this is inadequate in many settings:

It is interesting to observe what additional information the seller needs in order to determine the rational acceptance price, over and above the information he needs once the acceptance price is set. He needs, in fact, virtually complete information as to the probability distribution of offers for all relevant subsequent time periods. Now the seller who does not have this information, and who will be satisfied with a more humbling kind of rationality, will make approximations to avoid using the information he doesn't have (Simon 1955, p. 117).

One way to address such a situation is by applying heuristics. Since the 1970s, the term heuristics has acquired a negative connotation in economics, psychology, and management, referring to the shortcomings of human reasoning (Tversky and Kahneman 1974). In computer science, however, it is used in line with its original Greek meaning—“to find out or discover”—to describe

comparatively simple algorithms for making intelligent inferences with incomplete information in situations of uncertainty or intractability.⁴ We follow this tradition and use the term to describe simple decision processes that use limited information.

Heuristics are typically derived from observation of expert decision-making in natural environments that are often fraught with uncertainty. To assess the performance of strategies under uncertainty, one cannot rely on the axioms of rational choice theory, which apply solely to situations of risk. Instead, one can assess quality by comparing performance among a set of strategies, for example, that of rational choice strategies with satisficing heuristics. In these comparisons, heuristics often perform surprisingly well or even outperform highly complex strategies, vindicating their use by experts. Determining the conditions under which strategies work well under uncertainty forms the subject of the study of ecological rationality. In other words, heuristics provide the answer to the question of how decisions are made, while ecological rationality is the answer to the question of why a given strategy works well.

Heuristics are often specified in terms of the three elements highlighted before: a search rule, a stopping rule, and a decision rule (Gigerenzer and Goldstein 1996). Using this taxonomy, we can characterize satisficing models more precisely as decision models that (i) search through alternatives or attributes and (ii) use an aspiration level in their stopping rules. Learning and adaptation can lead agents to rely on specific classes of strategies tailored to classes of decision problems. The resulting assemblage of strategies represents an “adaptive toolbox” (Gigerenzer and Selten 2001). We will review several classes of decision problems along with classes of satisficing heuristics (see table 3 for an overview; heuristics are ordered by their appearance in the text).

3.2.1 *Aspiration-Level Adaptation*

Smith (1962) observed a paradox: markets quickly converge to equilibrium even though agents operate under information conditions that are much weaker than specified by the theory that characterizes the aggregate market. But what are the decision strategies that agents actually use in such a context to solve the problem? Addressing this puzzle, Artinger and Gigerenzer (2016) conduct an analysis investigating both constructivist and ecological rationality by studying pricing in the online used car market, which is characterized by a large degree of uncertainty. They find that the market is well fitted to the aggregate data by an equilibrium model by Varian (1980) that captures both observed price dispersion and average price in the tradition of a constructivist analysis. Unlike the equilibrium model, the aspiration-level adaptation heuristic, originally proposed by Simon (1955), correctly predicts the actual dynamic setting of the price. This heuristic is virtually used by all dealers, who initially start with a high aspired price A_1 and lower it at fixed time intervals t , usually by a predetermined margin δ , until a car sells. The aspiration-level adaptation heuristic systematically captures observed phenomena such as high initial price, price stickiness, and the “cheap twin paradox,” whereby two highly similar cars at the same dealership have a different price tag due to the simple fact that the price of the car that has been on offer for longer has been reduced after a fixed time interval.

Artinger and Gigerenzer (2016) show that the parameters the dealers use—the initial price A_1 , the duration t that the price

⁴ The first textbook on heuristics in computer science was written by Pearl (1984), who, like Simon, received the Turing Award, the highest honor in the field of computer science and often compared to the Nobel Prize; see Lucci and Kopec (2016) for an up-to-date treatment.

TABLE 3
CLASSES OF SATISFICING HEURISTICS

Heuristic	Aspiration Level	Search Rule	Stopping Rule	Decision Rule
Aspiration-level adaptation	A_1 : minimum value on single attribute (e.g., profit or price)	by alternatives i	$a_{i1} > A_1$; if after time t search could not be stopped, lower A_1 by Δ	Choose a_i
Multi-attribute aspiration-level adaptation	A_j : minimum value on multiple attributes	by alternatives i	$a_{ij} > A_j$ for all j ; if after time t search could not be stopped, lower A_j by Δ_j	Choose a_i
Hiatus (recency)	A_t : time passed since last event	by alternatives i	$a_{it} > A_t$	Classify a_i into one of two categories
Tallying	A_j : minimum values on multiple attributes	by attributes j	$a_{ij} > A_j$ for k out of n	Classify a_i into one of two categories
Fast-and-frugal trees	A_j : minimum values on ordered attributes	by attributes j	$a_{ij} > A_j$ for first j that permits exiting the tree	Classify a_i into one of two categories
Take-the-best/ Δ -inference	A_j : minimum value of difference in j between a_1 and a_2	by attribute j in order of validity	first attribute for which $ a_{1j} - a_{2j} > A_j$	Choose alternative with higher value on j
Elimination-by-aspects	A_j : minimum level on attributes	by attributes j in random order	for each j , eliminate all a_i with $a_{ij} < A_j$, until only one a_i remains	Choose remaining a_i

Notes: a_i denotes alternative 1, 2, ..., i , ..., N , x_j denotes attribute 1, 2, ..., j , ..., n , a_{ij} denotes the value of alternative a_i on attribute j , A_j denotes the aspiration level for x_j ; for ease of presentation, we assume that the aspiration level is a minimum rather than a maximum value and is lowered when adapted.

is held constant, and price reduction δ — vary systematically with the local market conditions, an indication of the ecological rationality of aspiration-level adaptation pricing. Specifically, the higher the population density in the local market and number of dealerships, the shorter the duration t that the price is held constant. In a more densely populated area with more competition, a dealer can more quickly infer that a car is unlikely to sell for a given price, whereas in less densely populated areas with less competition the price needs to be held constant for a longer time.

Pricing offers an illustration of a simple heuristic that formulates an aspiration level on one attribute, the price, and adapts it if necessary. Camerer et al. (1997) hypothesize that taxi drivers terminate their shifts after earning a daily income target. However, if there is an increase of demand on a given day, and taxi drivers could predict it, this would imply that drivers stop their shifts too early. Subsequent research has tested this hypothesis by comparing the descriptive powers of neoclassical and reference-dependent versions of utility theory (e.g., Farber 2008, 2015; Crawford and Meng

2011). However, Artinger, Gigerenzer, and Jacobs (2020) find that hourly earnings per driver are barely predictable and therefore hypothesize that drivers use some form of heuristic, rather than strategies that rely on rational expectations, such as utility theory. To test their hypothesis, the authors compare two utility models and four satisficing heuristics in predicting taxi drivers' shift ends. They find that the behavior of the vast majority of drivers is best predicted by satisficing heuristics that terminate shifts after working for a fixed number of hours or after earning a fixed amount of income. The authors speculate that the choice of the aspiration variable reflects which attribute drivers prioritize, whereas the aspiration level is chosen based on experience such that a reasonable balance of income and leisure is reached. Both heuristic models use fixed aspiration levels that are not adapted over time. Similarly, Berg (2014) reports that developers of high-rise office building and malls decide in favor of an investment if they can get at least "X percent return" in two or three years, that is, a return that exceeds an aspiration level A for a time t .

3.2.2 Multi-attribute Aspiration-Level Adaptation

In other situations such as academic job search, several attributes are relevant, for instance, prestige of the institution, salary, location, quality of local schools, and spouses' and family preferences. In addition, some of these attributes cannot be traded for others but may be perceived as incommensurable. For these situations, Selten (1998) and Sauermann and Selten (1962) provide a solution that closely follows Simon's (1955) aspiration-level adaptation heuristic (see table 3). Agents have an aspiration grid for a set of n incommensurable attributes. After each alternative a_i is examined, aspiration levels A_j are adapted for some attributes j according to a ranking that is affected by preferences and may change as search progresses. When the aspiration for a specific attribute is not met by the examined alternative, it is lowered; otherwise it is increased. When no further increases are feasible, the alternative that meets all aspiration levels is chosen.

In an experimental setup, Stüttgen, Boatwright, and Monroe (2012) use an eye-tracking device that enables monitoring of the search, stop, and decision process for choices among brands of instant noodles with which participants are familiar. Such a naturalistic task makes it possible to investigate a situation with multiple, potentially incommensurable attributes. They test the predictive accuracy of two different models: one based on expected utility and the other on Simon's (1955) aspiration-level heuristic, albeit without adaptation, where an agent forms aspiration levels for each attribute separately. The satisficing model predicts the observed data much more closely than the utility model does, suggesting that even in such a mundane task, incommensurability is at work. In particular, when an alternative is found that meets all aspiration levels, search is concluded after a verification phase in which the agents acquire additional information.

3.2.3 Hiatus Heuristic

Aspiration levels can also be used for predicting whether an event observed in the past will occur again in the future. A case in point is that marketing practitioners often rely on the hiatus heuristic, predicting that customers will make further purchases if they have made a purchase within the past t days, otherwise not (table 3). Here, the aspiration level A_t refers to time passed. This practice contrasts sharply with the Pareto–negative binomial distribution (NBD) model (Schmittlein, Morrison, and Colombo 1987), a stochastic model that also predicts future purchases. Using purchasing data from three industries, Wübben and von Wangenheim (2008) set out to demonstrate the superiority of the stochastic model but find that

the hiatus heuristic yielded the same or better out-of-sample predictions than did the Pareto–NBD model (see also Persson and Ryals 2014).

Theirs is not an isolated finding. An exclusive reliance on recency, that is, the time since the last event occurred, which ignores all other variables, has long been considered irrational yet has been observed in many different domains (e.g., Kunreuther 1976, Gallagher 2014, Malmendier and Nagel 2011). Using 60 different datasets across many different domains, Artinger et al. (2018) show that the hiatus heuristic is the single best strategy for predicting future events such as purchases, outperforming logistic regressions and highly sophisticated machine learning methods such as random forests that make use of recency, frequency, and any other available information.

3.2.4 Tallying

Predicting the next president of the United States is a problem entailing high uncertainty about voters' behavior. In November 2016, when prediction markets, polls, and big data analytics predicted Hillary Clinton's victory by a large margin, Lichtman (2020) instead predicted that Donald Trump would win. Using a tallying heuristic, Lichtman's model has correctly predicted all presidential elections since 1984. The heuristic considers 13 attributes that Lichtman calls "keys," which comprise yes–no questions such as whether the incumbent party holds more seats in the US House of Representatives after midterm elections than it did after the previous midterm election, whether the incumbent-party candidate is the sitting president, and whether real annual per capita economic growth during the term equals or exceeds mean growth during the two previous terms. The tallying rule is:

Keys to the White House: If six or more attributes are negative, then the challenger will win.

Note that there is no attempt to estimate the weights of the attributes or their covariances; all are given equal weight and simply counted. Unlike in aspiration-level adaptation, where the attributes are given and search takes place by evaluating alternatives, here the alternatives are given and search takes place by evaluating attributes. Search is then stopped when 6 out of 13 attributes are negative. As an aside, the model correctly predicted the election of Joe Biden in 2020 (Lichtman 2020).

Åstebro and Elhedhli (2006) report evidence that a tallying heuristic for classifying early-stage ventures performs at least as well as computationally intensive models, while being faster and requiring less information. The heuristic first tallies the positive and negative attributes of a specific venture. If the tally of the positive attributes exceeds the aspiration level and the tally of the negative attributes falls short of the aspiration level, a venture is classified as promising. Testing the strategy competitively against a linear model, the authors found that the satisficing strategy reached a predictive accuracy of 83 percent, compared with 79 percent for the linear model. Similarly, Jung et al. (2020) show that a system of bail decisions based on tallying leads to recommendations that are as accurate as those of complex machine learning systems, but less costly and more transparent.

In general, consider a choice between alternatives a_1 and a_2 , and n binary or continuous attributes x_j with aspiration levels A_j . The prediction or choice is made by a simple rule:

Tallying: Choose a_1 if $a_{1,j} \geq A_j$, for at least k out of n attributes. Otherwise choose a_2 .

Note that aspiration-level adaptation requires that all A_j are met, whereas tallying requires that only some k with ($k \leq n$) of A_j met. That makes it possible to consider a larger number of attributes and determine

whether a critical subset is met. Tallying is compensatory, treating the attributes equal and exchangeable. In contrast, the next class of heuristics treats these in a non-compensatory way: if the first aspiration level is met, search is stopped and a decision is made.

3.2.5 Fast-and-Frugal Trees

In classification problems, an agent needs to assign an alternative to one of several classes according to its attributes. One class of heuristics that addresses these problems are fast-and-frugal trees (Martignon, Katsikopoulos, and Woike 2008). Fast-and-frugal trees order attributes sequentially and define aspiration levels for each of them. Unlike other classification trees, however, fast-and-frugal trees can reach a decision after one or a few attributes are examined. That is, for attributes that are split at their aspiration levels, they have $n + 1$ exits for n attributes, of which the first $n - 1$ attributes have one exit each and the final one has two. Thus, a fast-and-frugal tree has $n + 1$ exits, while a full tree has $2n$ exits. These two features, order and aspiration levels, reduce estimation error and prevent intractability with large numbers of attributes.

For illustration, consider figure 2, which displays a decision tree by Aikman et al. (2021) developed for the Bank of England to identify banks at default risk. This tree uses $n = 3$ attributes, leading to $n + 1 = 4$ exits overall. The first attribute that is examined is the bank's leverage ratio. If this ratio falls short of the aspiration level of 4.1 percent, the bank is immediately classified as vulnerable, without inspecting subsequent attributes. Only when the leverage ratio is at least 4.1 percent is the next attribute examined, the market-based capital ratio, which can also lead to an immediate decision. The sequential nature of the fast-and-frugal tree models a form of incommensurability: if a bank has an extremely poor leverage ratio, an excellent loan-to-deposit ratio cannot compensate for that. This is analogous to many other real-life systems: a healthy liver cannot compensate for a failing heart. The exit structure of the fast-and-frugal tree reflects the cost structure of misclassifications (Luan, Schooler, and Gigerenzer 2011). The bank classification tree has a "red flag exit" after each attribute and thereby minimizes misses at the cost of false alarms, compared to the three other possible trees with the same order of attributes but different exit structures. In principle, there are two means of parameterizing the heuristic: by estimating the parameters statistically or by a combined method in which an expert determines the attributes, their order, and the exits and then determines the thresholds for each variable using statistical methods. Aikman et al. (2021) show that the statistically estimated fast-and-frugal tree is better than a logit model at predicting vulnerable banks, and that the combined method outperformed both. The construction of fast-and-frugal trees, as well as of tallying models, and their predictive accuracy relative to regression and machine learning models are explained by Katsikopoulos et al. (2020).

3.2.6 Take-the-Best, Δ -Inference, and Elimination-by-Aspects

Consider a choice problem where the agent needs to select one alternative from a set based on the attributes of each alternative but does not know the overall value or utility of each alternative. For example, an agent wants to identify the house with the highest quality based on attributes such as price, floor size, location, or number of rooms. Here, the agent knows the available options but does not know the outcome (quality) and needs to infer it from the given attributes, where the relation between attributes and quality is not or is only partially known. This meets the definition of uncertainty with regard to outcomes; the alternatives are known. When choice is between two alternatives,

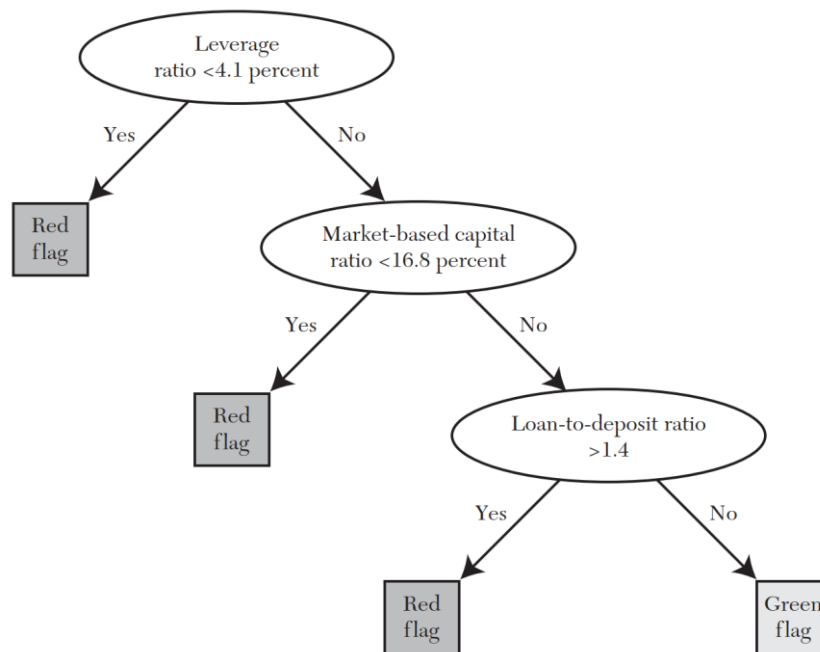


Figure 2. Fast-and-Frugal Tree for Bank Classification (Aikman et al. 2021)

Note: Note that the order of attributes affects the decision.

the take-the-best heuristic can be used (Gigerenzer and Goldstein 1996). This strategy examines attributes lexicographically, that is, they are ordered by their validity, defined as the percentage of correct inferences made by that attribute alone. If the two options have identical values for the attribute, it is ignored and the next attribute is examined. Search through attributes stops as soon as one is found that discriminates, that is, when the difference in value exceeds the aspiration level of zero. Once such an attribute is found, no further attributes are examined, and the decision is based solely on the one discriminating attribute. Whereas take-the-best is limited to decisions with binary attributes, Δ -inference (Luan, Schooler, and Gigerenzer 2014) can be used when attributes are continuous.

Testing the performance of take-the-best and Δ -inference shows that these heuristics are surprisingly powerful in prediction. Across 20 different datasets, ranging from school dropout rates to property prices, take-the-best was less accurate than multiple regression in choice from seen alternatives (data fitting) but more accurate in choice from unseen alternatives, that is, in out-of-sample testing (Czerlinski, Gigerenzer, and Goldstein 1999). Similarly, comparing take-the-best with machine learning models such as classification-and-regression trees (CART) and support vector machines shows that take-the-best can match or even outperform these in an out-of-sample setting, while using less information (Brighton and Gigerenzer 2008, 2012). Luan, Schooler, and Gigerenzer (2014) and Luan, Reb,

and Gigerenzer (2019) find that Δ -inference yields better out-of-sample performance than linear regression and machine learning models such as random forest across 20 additional datasets, irrespective of whether the aspiration level was set optimally (based on past data) or to its most robust value.

When alternatives abound, direct comparisons may no longer be feasible. With more than three alternatives available, the number of possible direct comparisons exceeds the number of alternatives. In these cases, efficiency would not require agents to compare alternatives with one another, but instead to assess them individually. This is done, for instance, by the elimination-by-aspects heuristic, which forms aspiration levels for each attribute and examines them in the order of their importance (Tversky 1972). For each attribute, it eliminates all alternatives that do not meet its aspiration level until a single alternative remains, which is then chosen. The structure of elimination-by-aspects resembles the class of consideration-set heuristics that have been proposed as a strategy for dealing with large sets of alternatives in research on consumer decisions: Rather than examining all available alternatives in detail, consumers heuristically exclude options from detailed analysis. The resulting consideration set is then submitted to detailed examination at a second stage (e.g., Hauser and Wernerfelt 1990, Hauser et al. 2010, Marewski et al. 2010). Hauser (2014) reviews various heuristics that have been proposed for the formation of consideration sets, including satisficing approaches. These heuristic models are highly predictive of how consumers form small consideration sets (Yee et al. 2007, Dzyabura and Hauser 2011). Consideration sets bear a resemblance to the choice rules discussed earlier (Manzini and Mariotti 2007; Manzini, Mariotti, and Tyson 2013), with the difference that the work on choice rules makes specific assumptions about preference orders, whereas consideration set heuristics do not assume a specific preference ordering. This difference reflects diverging assumptions about the decision environment: The notion of preference orders assumes a situation of risk where alternatives can be at least partially ranked because all necessary information is available to the agent. In contrast, consideration set heuristics have been devised for situations where the set of alternatives is simply so large that one cannot meaningfully deduce an order.

Note that all heuristics listed in table 3 can effectively deal with incommensurability between attributes and do not need to integrate them onto a single cardinal scale as in expected utility theory. Search proceeds on an attribute-by-attribute basis where the value of an attribute is evaluated with respect to an aspiration level. This is also the case for the aspiration-level heuristic, which searches by alternatives and evaluates these attribute by attribute. Also note that not all heuristics that deal with decisions under uncertainty are satisficing heuristics. For instance, $1/N$, which allocates an investment equally over N assets, relies neither on an aspiration level nor on search. Its rationale is to reduce error from estimating asset weights and the covariance matrix (DeMiguel, Garlappi, and Uppal 2009).

3.2.7 Heuristics for Well-Defined but Intractable Problems

The above examples address uncertainty. However, even if the complete set of alternatives, possible future states of the world, and consequences is known, intractability surprisingly quickly makes rational choice infeasible and a heuristic strategy becomes an effective solution. Gabaix et al. (2006) show this by comparing two different conditions that resemble such a mundane task as selecting a television characterized by a few attributes. The first condition is simple: Facing three alternatives with only one

attribute each and a stochastic payoff, the decision maker sequentially explores the attributes and the values that realize, and in turn the value of an alternative. In this condition, information acquisition is costly and the agent can stop acquiring information at any time. The Gittins–Weitzman index solves the problem optimally by establishing a complete sequence with which to explore the products and their attributes (Gittins 1979, Weitzman 1979). The more complex condition is characterized not by three but by eight alternatives with not one but nine attributes. In the experiment, the attributes are positive or negative numerical values, which facilitates commensurability and which participants get to know during a sequential search process. The authors report that the complex task is computationally intractable for the rational choice model because the problem suffers from the curse of dimensionality (see also Salant 2011, González-Valdés and de Dios Ortuzar 2018). However, they show that participants in a laboratory experiment employed the directed cognition heuristic, a simple, myopic strategy that looks just one step ahead instead of solving the complete sequence. It can solve not only the simple problem but also the complex problem that rational choice cannot address. The heuristic proceeds as follows: First, it compares the value of stopping search immediately and the expected value of stopping immediately after the next attribute. This is a myopic calculation because it does not incorporate the possible consequences of continuing search beyond the realization of the next attribute. Second, the heuristic inspects the attribute with the highest expected value. Third, the first two steps are repeatedly executed until the costs of searching and inspecting another attribute outweigh the expected value of the next most attractive attribute, which represents the aspiration level.

Myopic search is a general tool for finding good solutions to intractable problems. Examples are scheduling problems in transportation, where the task is to find the shortest route through cities, beginning and ending at the same city. For 50 cities, finding the shortest route would require searching through more than 10^{62} possible routes. Heuristics such as the nearest-neighbor algorithm—move to the nearest unvisited city— can provide excellent solutions where the best one cannot be found (Lucci and Kopec 2016).

3.2.8 Behavioral Theory of the Firm

The use of an aspiration level as a heuristic decision strategy was already at the core of Simon's dissertation, published in 1947 under the title "Administrative Behavior." Here, Simon was primarily concerned with firms; a more general decision context is found in his seminal paper of 1955. Nonetheless, it was "for his pioneering research into the decision making process within economic organizations" that Simon was awarded the 1978 Nobel Memorial Prize in Economic Sciences. His work was developed further (March and Simon 1958) and ultimately inspired the "behavioral theory of the firm" (Cyert and March 1963). In its analysis of the fundamental decisions of the firm, such as price, output, and resource allocations, it lays "an explicit emphasis on the actual process of decision making as its basic research commitment" (Cyert and March 1963, p. 19). This stands in sharp contrast to traditional economic theory that focuses on market level outcomes and classically models firms as rational actors.

The theoretical foundations on which behavioral theory of the firm builds in order to understand the actual decision processes are a) satisficing instead of maximization, where the first alternative that is satisfactory with respect to an aspiration level is chosen; b) search for information when not all possible outcomes of any choice alternative can be anticipated; and c) the use of robust rules in the face of uncertainty,

circumventing predictions about the distant future. Its quest is to align models as closely as possible with the empirical observations of both the output that organizations produce and the process they use (for a review, see Gavetti et al. 2012). The theory has risen to prominence particularly in the domains of organization and strategy (March 1981, 1991; Winter 2000).

The behavioral theory of the firm predicts that an aspiration level is a function of recent performance, past historical aspiration levels, and recent performance of other firms. To understand organizational learning processes, models of aspiration formation have been developed that are clearly rooted in the tradition of bounded rationality. In economics, models of this form have a long history as adaptive expectations (e.g., Sterman 1987, Chow 1989) or adaptive learning (e.g., Jacobs and Jones 1980). The main difference is that in expectation models, aspiration levels are not explicitly modeled, but subsumed in an overall function. An ecology of learning organizations and competition yields a continuous adaptation process. Search is initially local; if no solution is found that yields satisfactory performance, a broader search ensues (Levinthal and March 1981). If this does not yield a satisfactory outcome, the aspiration level is adapted.

The behavioral theory of the firm has also inspired research in economics, giving rise to the field of evolutionary economics (Nelson and Winter 1973, 1982), which examines industrial evolution (e.g., Dosi, Nelson, and Winter 2000). Nelson and Winter (1982) model the evolution of an industry with firms using particular rules or routines in situations where the world is characterized by complexity and uncertainty. Routines emerge as a response to an adaptive process where no optimal solution can be found *ex ante*. A firm uses a given routine as long as its output remains above an aspiration level; only when the output falls below this does it engage in exploration for an alternative routine (Winter 1971). Such behavior necessarily results in companies failing to survive, as shown for instance by Witt (1986). Comparing three algorithms in a multi-period market competition where the first maximizes expected profits, the second uses an aspiration level and satisfices when setting prices, and the third algorithm is based on simple reinforcement learning, Witt shows that the survival of an algorithm strongly depends on the initial conditions and that optimization does not dominate the other algorithms. Given uncertainty, Heiner (1983, 1989) formally shows that firms will adapt only a limited number of simple decision rules. Rules are added solely if they exceed an aspiration level, which refers to the reliability with which the rule will generate profitable future actions. The larger the degree of uncertainty, the fewer the rules.

The rich body of empirical research on firms operating under uncertainty shows that they rely on various forms of satisficing (Artinger et al. 2015). Much of this research concerns the actions of firms and the evaluation of performance with regard to an aspiration level. Aspirations determine whether past performance is framed as a success or failure, which influences subsequent strategic decisions (Lant 1992). The first to provide empirical evidence for the use of an aspiration level in the formation of organizational goals is Lant (1992). In a laboratory experiment, she uses a management game employed by companies for in-house training that captures the complexity and dynamics that managers frequently face. Teams of MBA students and managers from an executive program compete over multiple rounds of producing and selling two types of consumer products. After teams set their own goals, they make strategic and resource allocation decisions on a range of different variables. The underlying software uses a complex nonlinear algorithm that simulates

a competitive market in a multidimensional, interdependent world. Lant (1992) observes the goals, or aspiration levels, that the teams set in terms of their targeted sales as well as their actual performance. Her findings indicate that the teams are best described as satisficing, whereas the rational expectation model receives relatively little support (see also Lant and Shapira 2008, Audia and Greve 2006). The results by Lant (1992) have been replicated in a study by Mezias, Chen, and Murphy (2002), who use field rather than laboratory data of decision makers in a financial services company. Following Lant (1992), considerable evidence has accumulated on the use of aspiration levels in evaluating and guiding firm performance in such diverse contexts as research and development expenditures, outsourcing, and firm growth (e.g., Bolton 1993; Miller and Chen 1994; Greve 1998; Audia, Locke, and Smith 2000; Baum et al. 2005; De Boer, Gaytan, and Arroyo 2006; Greve 2008; Berg 2014). Blettner et al. (2015) study the process of aspiration adaptation in a longitudinal dataset from the news magazine industry. They find that when constructing an aspiration level, organizations largely focus on their own past performance and whether they achieve the goals they set for themselves. When they are close to bankruptcy, however, they start focusing on competitors' performance.

In summary, considering the search process is an important element in evaluating the quality of individuals' and firms' decisions. Given search, empirical evidence shows that agents do choose the alternative that is the best so far. Given alternatives with two or more potentially incommensurable attributes, heuristics dispense with the need to integrate these by applying aspiration levels. Heuristics can be powerful tools to make good decisions given uncertainty or intractability. But what exactly are the conditions under which heuristics perform well?

4. A Framework for Integration

As we have argued in the previous section, satisficing strategies are investigated within two distinct research traditions. These traditions differ in the assumed decision environment (risk versus uncertainty) and arrive at different conclusions about the rationality of satisficing: Although models of satisficing under risk without search conclude that satisficing is not commendable because it produces systematic mistakes, satisficing strategies given uncertainty or intractability are often found to yield better decisions than those by alternative models. There are two possible explanations for such divergent findings that allow for integration of these two different traditions: (i) cost–benefit trade-off in information acquisition and in information processing, and (ii) the bias–variance trade-off, which provides a framework to understand what type of strategy is best given the information available, beyond the costs of search and computation.

4.1 Cost–Benefit Trade-Off

Operating on the basis of limited information can be a result of the costs of information search. Given such costs, it can be beneficial to stop search early and make a decision based on a limited sample such that people deliberately decide not to know (for a review in an applied context, see Hertwig and Engel 2016). Since at least the work of Hayek (1945), economists have studied the trade-off between learning costs and decision quality. Agents in general do indeed choose the better alternatives or, as Simon would put it, the best so far (Bearden and Connolly 2007; Caplin, Dean, and Martin 2011; Reutskaja et al. 2011). Costs of information acquisition can also rationalize some apparent mistakes in the field. For instance, Chetty, Looney, and Kroft (2009) find that agents buy unnecessarily expensive products due to search costs. Similarly, costs

for information acquisition also play a central role when agents visit only a relatively small number of websites while buying over the internet (De Los Santos, Hortaçsu, and Wildenbeest 2012). Search costs feature in several models in the form of information cost functions (e.g., Verrecchia 1982, Sims 2003, Caplin and Dean 2015). What these models have in common is that they make it possible to determine the optimal trade-off between benefits and costs of search. At the same time, however, they make some strong assumptions about agents' knowledge and the degree of complexity of the environment.

Similarly, operating with a simple decision strategy can result from the costs of computation. As shown, for instance, by Gabaix et al. (2006) and Salant (2011) and other work, mainly in computer science (for a review, see Lucci and Kopec 2016), computing an optimal solution can entail substantial costs or even quickly prove to be infeasible. People are inherently sensitive to such costs of computation. Payne, Bettman, and Johnson (1988) show that when the cost of computation is increased by introducing time pressure, agents switch to a simpler strategy. Experimental work concludes that humans tend to trade off benefits and costs of cognitive effort rationally (Kool and Botvinick 2014; Lieder et al. 2014; Lieder, Griffiths, and Hsu 2018). This suggests that apparent mistakes can be rationalized in terms of computational costs. Formal models that trade off computational benefits and costs have been developed, particularly in computer science and specifically in research on artificial intelligence (for a review, see Gershman, Horvitz, and Tenenbaum 2015). Such models start with a meta-level analysis of the ideal balance between computational effort and the quality of alternatives. Yet, as Gershman, Horvitz, and Tenenbaum (2015) stress, calculating an optimal solution in terms of costs and benefits of computation is frequently challenging. Few decision problems admit an analytic solution, and many problems are computationally intractable. In addition, Conlisk (1996) points to the problem of infinite regress: if computation is costly, then optimizing computation is costly ad infinitum.

Like search costs, computational costs provide a reason why satisficing and simple strategies that require little information can perform so well, albeit trading a reduction in costs for a reduction in accuracy. At the same time, section 3.2 reported a number of findings where heuristics perform on par or even outperform complex alternative models from, among others, operations research and machine learning, while reducing costs of search and computation at the same time. The performance criterion in many cases was predictive accuracy, without accounting for search and computational costs, which would have boosted the performance advantage of the heuristic even more. This calls for an alternative explanation.

4.2 Bias–Variance Trade-Off

An alternative reason why heuristics can perform so well concerns the trade-off between bias and variance. This trade-off explains why a model with fewer parameters can yield more accurate predictions than does a more generalized model with more parameters. The trade-off provides the explanation of why a simple model such as a heuristic can outperform a complex model simultaneously in terms of accuracy and other performance metrics.

A model's predictive accuracy is grounded in a basic statistical relation between bias and variance. Suppose the task is to predict y , the value of an unseen item, based on its observables, denoted by vector \mathbf{x} . Value y was generated by an unknown function $y(\mathbf{x}, \mathbf{Q})$, which combines the observables using a fixed but unknown parameter set \mathbf{Q} . A model $m(\mathbf{x}, \mathbf{q})$

is used to predict the value y . To this end, a learning sample is randomly drawn from the population of items generated by $y(\mathbf{x}, \mathbf{Q})$ and is used to calibrate \mathbf{q} , the parameter set of the model $m(\mathbf{x}, \mathbf{q})$. Based on this parameter set, a prediction about y is made, say $m(\mathbf{x}, \hat{\mathbf{q}})$. The error in prediction can be assessed in many ways, among which the mean squared error is a common choice (Geman, Bienenstock, and Doursat 1992). The mean squared error in predicting an item can be decomposed as follows:

$$(2) \quad \text{error} = \text{bias}^2 + \text{variance} + \epsilon$$

where ϵ denotes the irreducible error that cannot be eliminated even if the generating function $y(\mathbf{x}, \mathbf{Q})$ were known. An example of ϵ is unsystematic measurement error. In contrast to ϵ , bias is a systematic error of model $m(\mathbf{x}, \mathbf{q})$. To better understand this, suppose there are L possible independent learning samples of a given size, and each is used to fit the parameters of model $m(\mathbf{x}, \mathbf{q})$ and calculate a prediction about y , yielding a set of predictions $m(\mathbf{x}, \hat{\mathbf{q}}_1), m(\mathbf{x}, \hat{\mathbf{q}}_2) \dots m(\mathbf{x}, \hat{\mathbf{q}}_L)$. Bias is the difference between the average of these predictions and the expectation of the generating function:

$$(3) \quad \text{bias}^2 = \{E[m(\mathbf{x}, \hat{\mathbf{q}}_l)] - E[y(\mathbf{x}, \mathbf{Q})]\}^2,$$

where $E[m(\mathbf{x}, \hat{\mathbf{q}}_l)]$ denotes the expectation with respect to different learning samples and $E[y(\mathbf{x}, \mathbf{Q})]$ denotes the expectation with respect to unsystematic error. Error from bias reflects a systematic misspecification of model $m(\mathbf{x}, \mathbf{q})$ relative to the generating function $y(\mathbf{x}, \mathbf{Q})$. Variance, in contrast, arises from a lack of knowledge of the model's parameter values and the need to infer them from a limited sample of data. Formally, variance refers to the average variation of an individual prediction $m(\mathbf{x}, \mathbf{q}_l)$ around the average prediction:

$$(4) \quad \text{variance} = E[\{m(\mathbf{x}, \hat{\mathbf{q}}_l) - E[m(\mathbf{x}, \hat{\mathbf{q}}_l)]\}^2].$$

Error from variance reflects the sensitivity of the model to idiosyncrasies of the sample used for calibrating its parameters. An illustration is provided in the upper panel of figure 3. In it, the center of the target is the expectation of the unknown generating function, the best possible prediction. The black dots represent different predictions of $m(\mathbf{x}, \mathbf{q})$ based on different learning samples that yield different estimates of \mathbf{q} , and the gray dot is the average of these predictions. The model on the left has low bias and accurate average predictions but high variance, and its predictions fluctuate strongly with parameter estimates; the model on the right has higher bias but lower variance in predictions.

In general, a model that aims to maximize predictive accuracy needs to control both the bias and the variance components of the prediction error. Bias is reduced to the degree that model $m(\mathbf{x}, \mathbf{q})$ resembles the generating function $y(\mathbf{x}, \mathbf{Q})$. One means of achieving lower bias is to add free parameters to the model. However, each additional parameter increases (or at best keeps constant) the variance component of the prediction error, provided that these models are nested. Thus, there is a trade-off between bias and variance. The lower panel of figure 3 shows the point of the minimum prediction error for two different sample sizes given a set of free parameters estimated from the data. To the left of the point, including fewer parameters implies higher total error; to the right of it, more parameters imply higher total error. The exact location of the point is determined by factors affecting variance, such as the mathematical nature of the parameters (additive, multiplicative, etc.) and the size of the sample used to calibrate the parameters. With more data available, the model

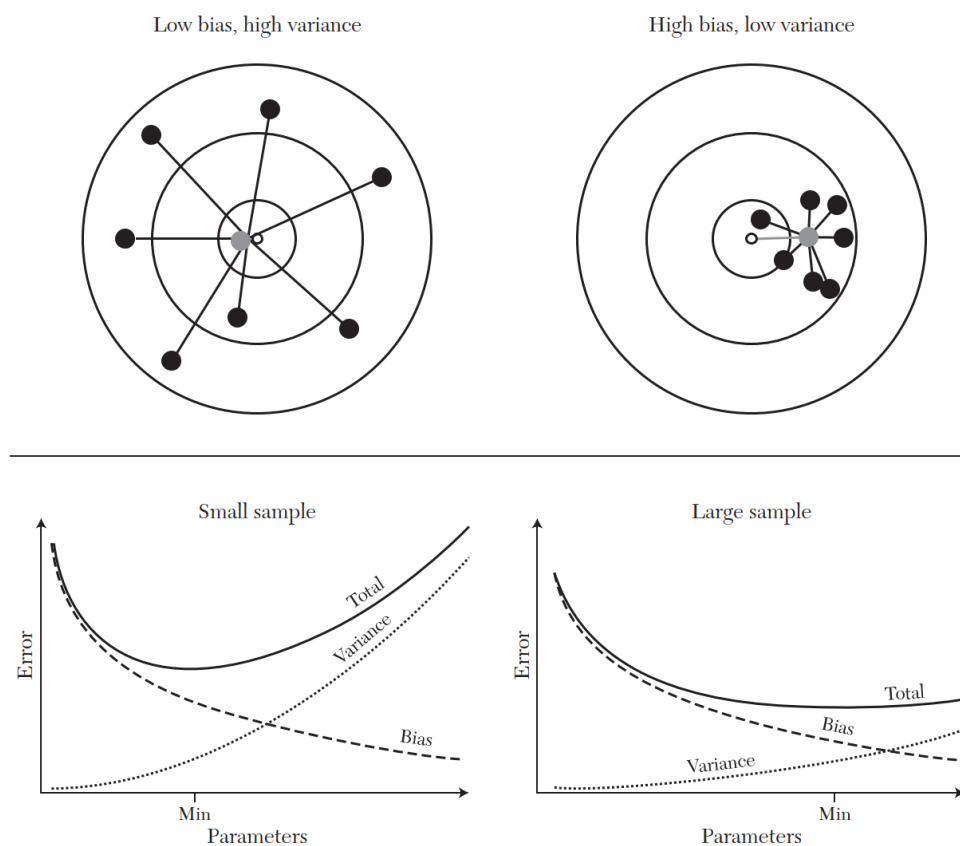


Figure 3. Bias–Variance Trade-off

Notes: Upper panel: The center of the target is the expectation of the unknown generating function, the best possible prediction. The black dots represent different predictions of $m(\mathbf{x}, \mathbf{q})$ based on different learning samples that yield different estimates of \mathbf{q} . The gray dot is the average of these predictions. The model on the left has a low bias and accurate average predictions but a high variance, and predictions fluctuate strongly with parameter estimates. The model on the right has higher bias but low variance in predictions. Lower panel: Bias diminishes but variance increases in the number of parameters. There exist a number of parameters that minimize total error, the position of which depends on many factors, including the size of the learning sample.

can afford more parameters while keeping its bias and variance in balance.

The trade-off between bias and variance is well-known in statistics and machine learning. However, empirical economics is divided in its view (e.g., Yatchew 1998, Varian 2014, Mullainathan and Spiess 2017). On the one hand, large parts of empirical economics are concerned with accurately estimating model parameters. Here, the goal is to obtain unbiased parameter estimates to evaluate, for instance, policy interventions and only rarely to predict (but see also Kleinberg et al. 2015). Overfitting is usually addressed by using tools such as the Bayesian information criterion or regularization. Yet controlling for overfitting does not seem to be a perfect safeguard. For instance, Artinger et al. (2018) and Luan, Reb, and Gigerenzer (2019) competitively test heuristics against regularized regression and random forests, models designed to be robust to overfitting. They find that the heuristics predict better than the more complex models, particularly in smaller samples. On the other hand, parts of microeconomics and most of

behavioral economics focus on comparing theories. This line of work rarely conducts out-of-sample tests of these theories, but instead typically compares the in-sample fit (e.g., Starmer 2005, Friedman et al. 2014). As the bias–variance trade-off illustrates, this practice can be misleading. Instead, theory comparison should follow two principles set out by Friedman in 1953:

(i) Theories need to be tested on prediction, not on how well they fit data:

The ultimate goal of a positive science is the development of a “theory” or “hypothesis” that yields valid and meaningful (i.e., not truistic) predictions about phenomena not yet observed (Friedman 1953, p. 7).

(ii) Theories should be tested competitively by comparing their predictions:

The question whether a theory is realistic “enough” can be settled only by seeing whether it yields predictions that are good enough for the purpose in hand or that are better than predictions from alternative theories (Friedman 1953, p. 41).

Given an uncertain environment, agents—like scholars who seek high accuracy—require strategies that balance bias and variance (Gigerenzer and Brighton 2009, Brighton and Gigerenzer 2015). We examine now what this conclusion implies for satisficing in different informational environments.

4.3 Satisficing and the Bias–Variance Trade-Off

Environments of risk allow for constructivist analysis. Under risk, the only error one can make is due to bias, whereas error due to variance is by definition not possible. With full information about the decision environment, scholars and agents are able to derive the optimal path of action. No inference about the decision environment is required and variance from estimation becomes irrelevant. Consequently, the prediction error, and hence quality of choice, depends solely on the bias of the decision strategy. In these situations, satisficing is suboptimal compared with unbiased strategies such as mathematical optimization.

Environments of uncertainty require a different approach; the analysis of the ecological rationality of strategies (Gigerenzer and Selten 2001, Smith 2008). In an environment of uncertainty, where the problem is ill-defined and the exhaustive set of states of the world and their consequences is not knowable or foreseeable, the best strategy cannot be foreseen. Thus, the study of ecological rationality under uncertainty compares strategies to determine which one likely performs better in what environment and, importantly, why. To do that, it analyzes both sources of error, bias and variance. The best-performing decision strategy does not reduce bias to zero, but finds a balance between the two kinds of errors. Here is the place for heuristics, whose simplicity can reduce error due to variance, such as by needing fewer free parameters. Given uncertainty, the optimal trade-off between bias and variance is hard to estimate. Instead, it requires two methodological approaches: competitive testing of strategies and testing on out-of-sample prediction.

Consider first the ecological rationality analysis of bias. So far, we have assumed that lower exposure to variance is paid for with increased bias. However, the amount of bias depends crucially on the statistical structure of the decision environment. When the structure of the heuristic matches the structure of the decision environment, the bias of a heuristic can be surprisingly low. To illustrate, consider the hiatus heuristic defined earlier, which predicts whether a customer will make future purchases or not (table 3). Can we identify conditions under which the bias of this heuristic is the same as that of linear models? Consider a simple example.

Assume a linear strategy with n binary attributes x_1, x_2, \dots, x_n with values of either $+1$ or -1 , where the positive value indicates future purchases. All of the weights of the attributes are positive and, like beta weights, reflect the additional contribution to the higher-ranked attributes. The linear rule infers that the customer will make future purchases if $y > 0$, otherwise not. If the following condition holds, the bias of the hiatus heuristic (or similar one-reason heuristics) is the same as that of a linear model:

Dominant Attribute Condition: The weights w_1, w_2, \dots, w_n form a dominant structure if they satisfy the inequality constraint:

$$(5) \quad w_1 > \sum_{k=2}^n w_k.$$

If there is a dominant attribute, the heuristic performs as well as and even better than the linear model because the latter incurs further error from variance (Artinger et al. 2018, Gigerenzer 2021b). Similarly, it can be shown analytically that the take-the-best heuristic has the same bias as a linear model when the weights of the linear model are non-compensatory (Martignon and Hoffrage 2002).

Non-compensatory attributes condition: Each weight, w_j , exceeds the sum of lesser weights:

$$(6) \quad w_j > \sum_{k=j+1}^n w_k.$$

In these situations, the choice made by a linear decision rule is determined exclusively by the attribute with the highest weight (see also Hogarth and Karelaia 2007). The take-the-best heuristic exploits such an environmental property by employing a non-compensatory, lexicographic decision strategy. This property leads take-the-best and a linear decision rule to have identical bias in non-compensatory environments. In addition to identical bias, lower variance leads heuristics to outperform linear models in prediction.

How often do these conditions hold in the real world? Şimşek (2013) examines 51 natural environments, ranging from car and house prices to the salaries of college professors. For 90 percent of the paired comparisons in half of the datasets, she finds that linear models yielded the same decisions as lexicographic models that decided on the basis of the first discriminating attribute (Gigerenzer 2021a).

As to variance, the larger the sample size and the smaller the number of free parameters, the lower the error by variance in general. Heuristic strategies, including satisficing heuristics, tend to have fewer free parameters than do estimation-and-optimization strategies. Heuristics and optimization strategies are rarely nested, which complicates a priori comparisons of their exposure to variance. The results of competitive performance tests, however, suggest that satisficing heuristics can be less susceptible to differences in training data and offer more robust predictions. It is important to realize that these conclusions are relative to the amount of training data used to compare these strategies. Şimşek and Buckmann (2015) show for 63 datasets that heuristics can be effective when training data is limited, but their advantage vanishes as the size of the learning sample increases. The work by DeMiguel, Garlappi, and Uppal (2009) illustrates that the necessary sample size can be very large. For portfolio selection, they compare the $1/N$ -heuristic, which allocates an investment equally over N assets, with the mean-variance model (Markowitz 1952) and a range of modern variants (see also Wang, Wu, and Yang 2015). They show for training data of 10 years that none of the sophisticated models was able to consistently outperform $1/N$. Indeed, for $N = 25$ assets, the mean-variance model requires about 250 years of stock data to outperform the simple $1/N$ -heuristic, and for $N = 50$ assets, 500 years are required, assuming that the same stocks, and the stock market itself, still exist.

The informational requirements of parameter-rich models are an important factor for understanding their performance under uncertainty. The derivation of these models under assumptions of sufficient information often ignores the important role of variance. With remarkable prescience, Simon (1981, p. 44) alludes to this issue, writing: “Although uncertainty does not [...] make intelligent choice impossible, it places a premium on robust adaptive procedures instead of strategies that work well only when finely tuned to precisely known environments.” In many situations of uncertainty, neither agents nor scholars can obtain an exhaustive representation of the decision environment at hand and therefore they cannot determine the optimal course of action. To find the best possible strategy therefore involves a competitive test of different strategies, as already highlighted by Friedman (1953).

Experts, such as car dealers or marketing managers, can generate high-performing, simple heuristics. With sufficiently large datasets, there are also statistical techniques available to generate such simple heuristics. Jung et al. (2020), for instance, develop a technique based on regularization that generates transparent and easy to understand heuristics that perform on par with black-box machine learning models such as random forest. Regularization, for instance using the Stein estimator (James and Stein 1961), the Lasso estimator (Tibshirani 1996), or ridge regression (Hoerl and Kennard 1970), reduces the exposure of the models to error due to variance. Specifically the Stein estimator, which is a biased estimator of the mean, can be shown to dominate the ordinary least squares approach in terms of a strictly better mean squared error.

There is an important caveat. The bias–variance dilemma assumes a stable population from which repeated samples are drawn. An example would be an agent, playing a lottery with unknown probabilities, who is given the opportunity to sample each alternative before choosing one. Here, agents may be able to estimate the relevant parameters of their decision environment before applying an optimization strategy, such as any form of utility maximization. Yet many situations contain more radical forms of uncertainty. The decision environment may be unstable and leave the agent without reliable learning samples for parameter estimation or the causal structure of the environment may be unknown. The problem may also suffer from computational intractability. Under these circumstances, agents are forced to employ some form of simplification in order to obtain a mathematical representation of the decision problem from which they can optimize. Because optimization strategies are only optimal relative to their assumptions or the sample they were estimated in, there is no guarantee that the decisions they yield are indeed optimal. Simon (1979) succinctly observes in his Nobel Prize speech that “decision makers can satisfice either by finding optimal solutions for a simplified world, or by finding satisfactory solutions for a more realistic world. Neither approach, in general, dominates the other, and both have continued to co-exist in the world of management science.” These alternative methodologies correspond to the two traditions of satisficing: optimal solutions assuming a situation of risk, and heuristic solutions assuming an uncertain world.

4.4 Integrating Two Research Traditions Using the Bias–Variance Trade-Off

This article set out to review advances in our understanding of satisficing and identified two research traditions. We refer to these as satisficing under risk and satisficing under uncertainty. Satisficing has been studied under both conditions, albeit in two largely unconnected literatures. Both have their roots in Simon’s 1955 article, the first

in the appendix, the second in the main text. For both classes of models the motivation is the same, namely to introduce aspiration levels and search in order to reflect what actual decision makers do. Models of satisficing under risk use rational choice theory as a starting point and then proceed to make modifications in line with a satisficing strategy to account for additional behavioral variance. Typically that requires strong assumptions about what decision makers know or can know about the future in order to use the optimization calculus. Moreover, in many of these models, only aspiration levels are considered, search is ignored, and satisficing is viewed as a failure to act rationally. Models of risk with aspiration levels and optimal search need to make additional strong assumptions about what can be known, such as complete information about the relevant probability distributions (e.g., the probability distribution of offers for all future time periods; see Simon, 1955) and the future cost of search, in order to calculate an optimal stopping point.

Satisficing under uncertainty refers to situations in which these assumptions are not met or cannot be met. Uncertainty includes situations where the exhaustive and mutually exclusive set of alternatives, or future states, cannot be known, meaning that the relevant probability distributions are also unknown. To model decision-making under uncertainty, we have used the bias–variance trade-off, whose key insights are that good models need to have not only low bias but also low variance and that there is a trade-off between the two sources of error. Variance arises in situations where unknown parameters need to be estimated; it does not arise in situations of risk, where the parameters (including probability distributions) are assumed to be known. Heuristics can reduce variance by using few free parameters, or even none (as in the hiatus heuristic with a fixed hiatus or in $1/N$), and thus can lead to more accurate predictions or decisions than more complex models. At the same time, they are fast, transparent, and reduce search costs. In situations of uncertainty, satisficing heuristics are no longer suboptimal; they can be the best one can do. But deciding which heuristic to choose in which situation requires careful study of the match between heuristic and environment, that is, by an analysis of their ecological rationality.

Thus, both traditions have their relevance, but in two quite different classes of problems. Against this dualism, one might argue that in situations of uncertainty, one could use the bias–variance dilemma to calculate the optimal trade-off between bias and variance in the same way as when using satisficing under risk, where the optimal stopping point is a trade-off between expected costs and benefits of further search. That would indeed reduce uncertainty to risk. To calculate the bias, however, one would need to know the true function that generates the data, which cannot be known in situations of uncertainty. Similarly, one could argue that all situations should be treated as ones of uncertainty. For instance, even in apparently certain conditions, unforeseeable events may happen or rules might be gamed. Yet that attempt toward reduction would be equally mistaken: Models are not unrealistic per se; they can be realistic for one class of problems and not for others. Moreover, both traditions have a different yet complementary approach to rationality, corresponding to Smith's (2008) distinction between constructivist rationality and ecological rationality.

Although Simon's original article planted the seeds for two different research traditions, their diverging methodologies and conclusions are not contradictory. Instead, the review has shown how these differences parsimoniously follow from the classes of decision environments they address. Under risk, where all the relevant information is known, optimization strategies are superior

to satisficing strategies and rational choice theory constitutes an obvious starting point. Given intractability when the problem is well-defined, or uncertainty when the problem is ill-defined, no single class of decision strategies is generally superior to another and models are derived from observation rather than function.

4.5 A Revised Understanding of Heuristics and Biases

One common view characterizes heuristics solely in terms of their bias (e.g., Tversky and Kahneman 1974). This view does not distinguish between risk and uncertainty and routinely concludes that agents' lack of computational capabilities leads them to make decisions that are not in their own interest. Generalizing the results of such studies to decision-making under uncertainty when the problem is ill-defined is more complicated than commonly assumed. The bias–variance trade-off implies that some degree of bias can be appropriate when the relative scarcity of data leads to potentially unbiased but unreliable parameter estimates. Under such conditions, biased decision strategies can yield better decisions than do unbiased ones. For example, Harry Markowitz used the $1/N$ heuristic for his retirement investments in place of calculating a mean-variance portfolio. Without an empirical test of the predictive accuracy of the $1/N$ heuristic, one likely would conclude that Markowitz relied on a suboptimal strategy. Would one also attribute this to cognitive limitations, as is usually the case for other agents' apparently biased decisions (see also Friedman 1998; Loomes, Starmer, and Sugden 2003; Friedman et al. 2007)?

Several phenomena often interpreted in the literature as biases (in the sense of systematic errors of judgment) have been shown to correspond to correct judgments in situations of uncertainty. One group of phenomena comprises judgments of randomness, such as coaches' alleged hot-hand fallacy (Miller and Sanjurjo 2018) and people's alleged erroneous intuitions about chance, including the belief in the law of small numbers (Hahn and Warren 2009). In earlier studies, these judgments are compared to known population probabilities rather than, correctly, to sample probabilities. A second group of phenomena include so-called errors in judgments of low versus high risk, such as overestimation of small risks and underestimation of large risks (Hertwig, Pachur, and Kurzenhäuser 2005) and overconfidence (Pfeifer 1994), which again look like systematic errors (bias) but are in fact largely due to unsystematic error (variance). In general, people appear to be quite sensitive to the difference between risk and uncertainty in both small and larger samples (Hertwig and Pleskac 2010, Gigerenzer 2018).

Expanding on Simon's observation that agents use simple strategies because of the mind's computational limitations, we propose that humans use heuristics because they can yield good decisions under uncertainty or intractability. Under uncertainty, their performance needs to be judged according to their ecological rationality, that is, by their success in achieving a defined criterion, not by principles of logic or consistency. Consistency and success are two different criteria, which are sometimes uncorrelated (Berg, Biele, and Gigerenzer 2016). Moreover, a survey of violations of consistency found little to no evidence that in an uncertain world, coherence violations incur material costs, or if they do, that people would fail to learn (Arkes, Gigerenzer, and Hertwig 2016). Whereas risky environments allow for general verdicts on the rationality of specific strategies, such generalizations are misguided in uncertain environments.

5. Discussion

The term satisficing has been used to mean many things. Sometimes, it is seen as

suboptimal, while at other times it seems optimal, or at least better than other strategies. In this article, we addressed this situation and have argued that there are two different research traditions that contend with two different types of problems, risk and uncertainty. In situations of risk, satisficing is suboptimal if search is ignored. However, when agents need to search for information, as considered by Simon (1955), satisficing can help to solve the cost–benefit trade-off. In situations of uncertainty, satisficing can also help to solve the bias–variance trade-off, which can lead to more accurate decisions than with more complex strategies. Although both research traditions examine distinct classes of environments, they can learn from each other. Specifically, we identify four methodological principles and related areas for fruitful future research on satisficing.

First, we encourage future research in both traditions to study how well decision theories make predictions, for example out of sample. We found very few studies that tested their decision models beyond data fitting, despite Friedman's (1953) endorsement of predictive accuracy. The bias–variance decomposition demonstrates how in-sample tests give undue advantage to parameter-rich models that prove inexpedient in predicting behavior.

Second, models of satisficing under uncertainty are currently limited in scope. To date, few of these models specify how aspiration levels are formed and how they are adjusted, although this process appears to be of theoretical importance, particularly in decisions guided by preferences (but see Cyert and March 1963; Blettner et al. 2015). The satisficing heuristics presented here were tested on inference problems for methodological convenience, although most of these strategies can be applied to decisions involving preferences. Smith (1962) already highlighted the disconnect between rational choice models that correctly predict market outcomes in a context where agents do not have the information available that the models assume. Yet little is known how the aggregate level of analysis interacts with the individual level when agents face uncertainty (but see Gode and Sunder 1993, Jamal and Sunder 2001, Artinger and Gigerenzer 2016).

Third, models of satisficing under risk remain largely within the realm of as-if theories, where additional parameters are added to a rational choice model to better account for behavior (e.g., Kahneman and Tversky 1979; Köszegi and Rabin 2006). The rational benchmark under risk also depends on the assumed structure of preferences and is not synonymous with rational choice theory. Models of the decision process rather than of decision outcomes could help uncover agents' preference structures, which in turn helps build a coherent theory of choice under risk.

Finally, we encourage further study of decisions given uncertainty, particularly when the problem is ill-defined and the exhaustive set of states of the world and their consequences is not knowable or foreseeable. Over the past years, technological advances have contributed considerably to the rise of machine learning, a branch of computer science distinctly concerned with building algorithms for decisions without full access to information. The rise of this field testifies to the fact that the problem of decisions under uncertainty poses larger problems than commonly acknowledged in normative economics. We therefore encourage competitive tests of specific decision strategies in their natural decision environments. A systematic use of competitive tests can in turn lead to formal analyses of the structural characteristics of environments that regulate the relative advantages of one class of strategies over another. We believe that this research strategy presents a feasible path to a normative theory for decisions under uncertainty.

References

- Aikman, David, et al. 2021. "Taking Uncertainty Seriously: Simplicity versus Complexity in Financial Regulation." *Industrial and Corporate Change* 30 (2): 317–45.
- Aleskerov, Fuad, Denis Bouyssou, and Bernard Monjardet. 2007. *Utility Maximization, Choice and Preference*. 2 ed. Berlin: Springer.
- Anscombe, F.J., and R.J. Aumann. 1963. "A Definition of Subjective Probability." *Annals of Mathematical Statistics* 34 (1): 199–205.
- Arkes, H.R., Gerd Gigerenzer, and Ralph Hertwig. 2016. "How Bad Is Incoherence?" *Decision* 3 (1): 20–39.
- Artinger, Florian M., and Gerd Gigerenzer. 2016. "Heuristic Pricing in an Uncertain Market: Ecological and Constructivist Rationality." <https://www.chapman.edu/research/institutes-and-centers/economic-science-institute/files/ifree-papers-and-photos/Artinger-Gigerenzer-Heuristic-Pricing-0160606.pdf>.
- Artinger, Florian M., Gerd Gigerenzer, and Perke Jacobs. 2020. "How Taxi Drivers Terminate Their Shifts When Earnings Are Hard to Predict." https://papers.ssrn.com/sol3/papers.cfm?abstract_id=3728515.
- Artinger, Florian M., Nikita Kozodi, Florian Wangenheim, and Gerd Gigerenzer. 2018. "Recency: Prediction with Smart Data." In *American Marketing Association Winter Conference Proceedings: Integrating Paradigms in a World Where Marketing Is Everywhere*, Vol. 29, ed. Jacob Goldenberg, Juliano Laran, and Andrew Stephen, L2–6. Chicago, IL: American Marketing Association.
- Artinger, Florian, Malte Petersen, Gerd Gigerenzer, and Jürgen Weibler. 2015. "Heuristics as Adaptive Decision Strategies in Management." *Journal of Organizational Behavior* 36 (S1): S33–52.
- Åstebro, Thomas, and Samir Elhedhli. 2006. "The Effectiveness of Simple Decision Heuristics: Forecasting Commercial Success for Early-Stage Ventures." *Management Science* 52 (3): 395–409.
- Audia, Pino G., and Henrich R. Greve. 2006. "Less Likely to Fail: Low Performance, Firm Size, and Factory Expansion in the Shipbuilding Industry." *Management Science* 52 (1): 83–94.
- Audia, Pino G., Edwin A. Locke, and Ken G. Smith. 2000. "The Paradox of Success: An Archival and a Laboratory Study of the Strategic Persistence following Radical Environmental Change." *Academy of Management Journal* 43 (5): 837–53.
- Baum, Joel A.C., Timothy J. Rowley, Andrew V. Shipilov, and You-Ta Chuang. 2005. "Dancing with Strangers: Aspiration Performance and the Search for Underwriting Syndicate Partners." *Administrative Science Quarterly* 50 (4): 536–75.
- Bearden, J. Neil, and Terry Connolly. 2007. "Multi-attribute Sequential Search." *Organizational Behavior and Human Decision Processes* 103 (1): 147–58.
- Berg, Nathan. 2014. "Success from Satisficing and Imitation: Entrepreneurs' Location Choice and Implications of Heuristics for Local Economic Development." *Journal of Business Research* 67 (8): 1700–1709.
- Berg, Nathan, Guido Biele, and Gerd Gigerenzer. 2016. "Consistent Bayesians Are No More Accurate Than Non-Bayesians: Economists Surveyed About PSA." *Review of Behavioral Economics* 3 (2): 189–219.
- Bernoulli, Daniel. 1954. "Exposition of a New Theory on the Measurement of Risk." *Econometrica* 22 (1): 23–36. (Orig. pub. 1738).
- Blettner, Daniela P., Zi-Lin He, Songcui Hu, and Richard A. Bettis. 2015. "Adaptive Aspirations and Performance Heterogeneity: Attention Allocation among Multiple Reference Points." *Strategic Management Journal* 36 (7): 987–1005.
- Bolton, Michele Kremen. 1993. "Organizational Innovation and Substandard Performance: When Is Necessity the Mother of Innovation?" *Organization Science* 4 (1): 57–75.
- Borch, Karl. 1968. "Decision Rules Depending on the Probability of Ruin." *Oxford Economic Papers* 20 (1): 1–10.
- Brighton, Henry, and Gerd Gigerenzer. 2008. "Bayesian Brains and Cognitive Mechanisms: Harmony or Dissonance?" In *The Probabilistic Mind: Prospects for Bayesian Cognitive Science*, edited by Nick Chater and Mike Oaksford, 189–208. Oxford, UK: Oxford University Press.
- Brighton, Henry, and Gerd Gigerenzer. 2012. "Are Rational Actor Models 'Rational' Outside Small Worlds?" In *Evolution and Rationality: Decisions, Co-operation and Strategic Behavior*, edited by Samir Okasha and Ken Binmore, 84–109. Cambridge, UK: Cambridge University Press.
- Brighton, Henry, and Gerd Gigerenzer. 2015. "The Bias Bias." *Journal of Business Research* 68 (8): 1772–84.
- Brown, David B., Enrico De Giorgi, and Melvyn Sim. 2012. "Aspirational Preferences and Their Representation by Risk Measures." *Management Science* 58 (11): 2095–2113.
- Brown, David B., and Melvyn Sim. 2009. "Satisficing Measures for Analysis of Risky Positions." *Management Science* 55 (1): 71–84.
- Camerer, Colin, Linda Babcock, George Loewenstein, and Richard Thaler. 1997. "Labor Supply of New York City Cabdrivers: One Day at a Time." *Quarterly Journal of Economics* 112 (2): 407–41.
- Caplin, Andrew, and Mark Dean. 2011. "Search, Choice, and Revealed Preference." *Theoretical*

- Economics* 6 (1): 19–48.
- Caplin, Andrew, and Mark Dean. 2015. "Revealed Preference, Rational Inattention, and Costly Information Acquisition." *American Economic Review* 105 (7): 2183–2203.
- Caplin, Andrew, Mark Dean, and Daniel Martin. 2011. "Search and Satisficing." *American Economic Review* 101 (7): 2899–2922.
- Charnes, Abraham, and William W. Cooper. 1963. "Deterministic Equivalents for Optimizing and Satisficing under Chance Constraints." *Operations Research* 11 (1): 18–39.
- Chetty, Raj, Adam Looney, and Kory Kroft. 2009. "Salience and Taxation: Theory and Evidence." *American Economic Review* 99 (4): 1145–77.
- Chow, Gregory C. 1989. "Rational versus Adaptive Expectations in Present Value Models." *Review of Economics and Statistics* 71 (3): 376–84.
- Conlisk, John. 1996. "Why Bounded Rationality?" *Journal of Economic Literature* 34 (2): 669–700.
- Crawford, Vincent P. 2013. "Boundedly Rational versus Optimization-Based Models of Strategic Thinking and Learning in Games." *Journal of Economic Literature* 51 (2): 512–27.
- Crawford, Vincent P., and Juanjuan Meng. 2011. "New York City Cab Drivers' Labor Supply Revisited: Reference-Dependent Preferences with Rational-Expectations Targets for Hours and Income." *American Economic Review* 101 (5): 1912–32.
- Cyert, Richard M., and James G. March. 1963. *A Behavioral Theory of the Firm*. Englewood Cliffs, NJ: Prentice-Hall.
- Czerlinski, Jean, Gerd Gigerenzer, and Daniel G. Goldstein. 1999. "How Good Are Simple Heuristics?" In *Simple Heuristics That Make Us Smart*, edited by Gerd Gigerenzer, Peter M. Todd, and ABC Research Group, 97–118. Evolution and Cognition series. New York: Oxford University Press.
- De Boer, Luitzen, Juan Gaytan, and Pilar Arroyo. 2006. "A Satisficing Model of Outsourcing." *Supply Chain Management* 11 (5): 444–55.
- De Los Santos, Babur, Ali Hortaçsu, and Matthijs R. Wildenbeest. 2012. "Testing Models of Consumer Search Using Data on Web Browsing and Purchasing Behavior." *American Economic Review* 102 (6): 2955–80.
- DeGroot, Morris H. 1970. *Optimal Statistical Decisions*. Hoboken, NJ: John Wiley and Sons.
- Dembo, Tamara. 1931. "Der Ärger als dynamisches Problem." *Psychologische Forschung* 15 (1–2): 1–144.
- DeMiguel, Victor, Lorenzo Garlappi, and Raman Uppal. 2009. "Optimal versus Naive Diversification: How Inefficient Is the 1/N Portfolio Strategy?" *Review of Financial Studies* 22 (5): 1915–53.
- Diecidue, Enrico, and Jeroen van de Ven. 2008. "Aspiration Level, Probability of Success and Failure, and Expected Utility." *International Economic Review* 49 (2): 683–700.
- Dosi, Giovanni, Richard R. Nelson, and Sidney G. Winter, eds. 2000. *The Nature and Dynamics of Organizational Capabilities*. Oxford, UK: Oxford University Press.
- Dzyabura, Daria, and John R. Hauser. 2011. "Active Machine Learning for Consideration Heuristics." *Marketing Science* 30 (5): 801–19.
- Ellsberg, Daniel. 1961. "Risk, Ambiguity, and the Savage Axioms." *Quarterly Journal of Economics* 75 (4): 643–69.
- Ether, Johanna, Meglena Jeleva, and Jean-Marc Tallon. 2012. "Decision Theory under Ambiguity." *Journal of Economic Surveys* 26 (2): 234–70.
- Farber, Henry S. 2008. "Reference-Dependent Preferences and Labor Supply: The Case of New York City Taxi Drivers." *American Economic Review* 98 (3): 1069–82.
- Farber, Henry S. 2015. "Why You Can't Find a Taxi in the Rain and Other Labor Supply Lessons from Cab Drivers." *Quarterly Journal of Economics* 130 (4): 1975–2026.
- Fishburn, Peter C. 1977. "Mean-Risk Analysis with Risk Associated with Below-Target Returns." *American Economic Review* 67 (2): 116–26.
- Friedman, Daniel. 1998. "Monty Hall's Three Doors: Construction and Deconstruction of a Choice Anomaly." *American Economic Review* 88 (4): 933–46.
- Friedman, Daniel, R. Mark Isaac, Duncan James, and Shyam Sunder. 2014. *Risky Curves: On the Empirical Failure of Expected Utility*. Oxon, UK: Routledge.
- Friedman, Daniel, Kai Pommerenke, Rajan Lukose, Garrett Milam, and Bernardo A. Huberman. 2007. "Searching for the Sunk Cost Fallacy." *Experimental Economics* 10 (1): 79–104.
- Friedman, Milton. 1953. *Essays in Positive Economics*. Chicago: University of Chicago Press.
- Gabaix, Xavier. 2014. "A Sparsity-Based Model of Bounded Rationality." *Quarterly Journal of Economics* 129 (4): 1661–1710.
- Gabaix, Xavier, David Laibson, Guillermo Moloche, and Stephen Weinberg. 2006. "Costly Information Acquisition: Experimental Analysis of a Boundedly Rational Model." *American Economic Review* 96 (4): 1043–68.
- Gallagher, Justin. 2014. "Learning about an Infrequent Event: Evidence from Flood Insurance Take-Up in the United States." *American Economic Journal: Applied Economics* 6 (3): 206–33.
- Gardner, John W. 1940. "The Use of the Term 'Level of Aspiration.'" *Psychological Review* 47 (1): 59–68.
- Garey, Michael R., and David S. Johnson. 1979.

- Computers and Intractability: A Guide to the Theory of NP-Completeness*. New York: W.H. Freeman Publishers.
- Gavetti, Giovanni, Henrich R. Greve, Daniel A. Levinthal, and William Ocasio. 2012. "The Behavioral Theory of the Firm: Assessment and Prospects." *Academy of Management Annals* 6 (1): 1–40.
- Geman, Stuart, Elie Bienenstock, and René Doursat. 1992. "Neural Networks and the Bias/Variance Dilemma." *Neural Computation* 4 (1): 1–58.
- Gershman, Samuel J., Eric J. Horvitz, and Joshua B. Tenenbaum. 2015. "Computational Rationality: A Converging Paradigm for Intelligence in Brains, Minds, and Machines." *Science* 349 (6245): 273–78.
- Gigerenzer, Gerd. 2004. "Striking a Blow for Sanity in Theories of Rationality." In *Models of a Man: Essays in Memory of Herbert A. Simon*, edited by Mie Augier and James G. March, 389–409. Cambridge, MA: MIT Press.
- Gigerenzer, Gerd. 2018. "The Bias Bias in Behavioral Economics." *Review of Behavioral Economics* 5 (3–4): 303–36.
- Gigerenzer, Gerd. 2021a. "Axiomatic Rationality and Ecological Rationality." *Synthese* 198 (4): 3547–64.
- Gigerenzer, Gerd. 2021b. "What Is Bounded Rationality?" In *Routledge Handbook of Bounded Rationality*, edited by Riccardo Viale, 55–70. Oxon, UK: Routledge.
- Gigerenzer, Gerd, and Henry Brighton. 2009. "Homo Heuristicus: Why Biased Minds Make Better Inferences." *Topics in Cognitive Science* 1 (1): 107–43.
- Gigerenzer, Gerd, and Daniel G. Goldstein. 1996. "Reasoning the Fast and Frugal Way: Models of Bounded Rationality." *Psychological Review* 103 (4): 650–69.
- Gigerenzer, Gerd., and Reinhard Selten, eds. 2001. *Bounded Rationality: The Adaptive Toolbox*. Cambridge, MA: MIT Press.
- Gigerenzer, Gerd, Ralph Hertwig, and Thorsten Pachur, eds. 2011. *Heuristics: The Foundations of Adaptive Behavior*. Oxford, UK: Oxford University Press.
- Gigerenzer, Gerd, Peter M. Todd, and ABC Research Group. 1999. *Simple Heuristics That Make Us Smart*. Evolution and Cognition series. New York: Oxford University Press.
- Gilbert, John P., and Frederick Mosteller. 1966. "Recognizing the Maximum of a Sequence." *Journal of the American Statistical Association* 61 (313): 35–73.
- Gilboa, Itzhak, and David Schmeidler. 1995. "CaseBased Decision Theory." *Quarterly Journal of Economics* 110 (3): 605–39.
- Gilboa, Itzhak, and David Schmeidler. 2001a. *A Theory of Case-Based Decisions*. Cambridge, UK: Cambridge University Press.
- Gilboa, Itzhak, and David Schmeidler. 2001b. "Reaction to Price Changes and Aspiration Level Adjustments." *Review of Economic Design* 6 (2): 215–23.
- Gittins, J.C. 1979. "Bandit Processes and Dynamic Allocation Indices." *Journal of the Royal Statistical Society: Series B (Methodological)* 41 (2): 148–64.
- Gode, Dhananjay K., and Shyam Sunder. 1993. "Allocative Efficiency of Markets with Zero-Intelligence Traders: Market as a Partial Substitute for Individual Rationality." *Journal of Political Economy* 101 (1): 119–37.
- Goldstein, Daniel G., R. Preston McAfee, Siddharth Suri, and James R. Wright. 2020. "Learning When to Stop Searching." *Management Science* 66 (3): 1375–94.
- González-Valdés, Felipe, and Juan de Dios Ortúzar. 2018. "The Stochastic Satisficing Model: A Bounded Rationality Discrete Choice Model." *Journal of Choice Modelling* 27: 74–87.
- Greve, Henrich R. 1998. "Performance, Aspirations, and Risky Organizational Change." *Administrative Science Quarterly* 43 (1): 58–86.
- Greve, Henrich R. 2008. "A Behavioral Theory of Firm Growth: Sequential Attention to Size and Performance Goals." *Academy of Management Journal* 51 (3): 476–94.
- Güth, Werner. 2010. "Satisficing and (Un)bounded Rationality: A Formal Definition and Its Experimental Validity." *Journal of Economic Behavior and Organization* 73 (3): 308–16.
- Güth, Werner, M. Vittoria Levati, and Matteo Ploner. 2010. "Satisficing in Strategic Environments: A Theoretical Approach and Experimental Evidence." *Journal of Socio-Economics* 39 (5): 554–61.
- Hahn, Ulrike, and Paul A. Warren. 2009. "Perceptions of Randomness: Why Three Heads Are Better than Four." *Psychological Review* 116 (2): 454–61.
- Haller, Hans. 1985. "The Principal–Agent Problem with a Satisficing Agent." *Journal of Economic Behavior and Organization* 6 (4): 359–79.
- Handel, Benjamin, and Joshua Schwartzstein. 2018. "Frictions or Mental Gaps: What's behind the Information We (Don't) Use and When Do We Care?" *Journal of Economic Perspectives* 32 (1): 155–78.
- Harstad, Ronald M., and Reinhard Selten. 2013. "Bounded-Rationality Models: Tasks to Become Intellectually Competitive." *Journal of Economic Literature* 51 (2): 496–511.
- Hauser, John R. 2014. "Consideration-Set Heuristics." *Journal of Business Research* 67 (8): 1688–99.
- Hauser, John R., and Birger Wernerfelt. 1990. "An Evaluation Cost Model of Consideration Sets." *Journal of Consumer Research* 16 (4): 393–408.
- Hauser, John R., Olivier Toubia, Theodoros Evgeniou,

- Rene Befurt, and Daria Dzyabura. 2010. "Disjunctions of Conjunctions, Cognitive Simplicity, and Consideration Sets." *Journal of Marketing Research* 47 (3): 485–96.
- Heath, Chip, Richard P. Larrick, and George Wu. 1999. "Goals as Reference Points." *Cognitive Psychology* 38 (1): 79–109.
- Heiner, Ronald A. 1983. "The Origin of Predictable Behavior." *American Economic Review* 73 (4): 560–95.
- Heiner, Ronald A. 1989. "The Origin of Predictable Dynamic Behavior." *Journal of Economic Behavior and Organization* 12 (2): 233–57.
- Hertwig, Ralph, and Christoph Engel. 2016. "Homo Ignorans: Deliberately Choosing Not to Know." *Perspectives on Psychological Science* 11 (3): 359–72.
- Hertwig, Ralph, Thorsten Pachur, and Stephanie Kurzenhäuser. 2005. "Judgments of Risk Frequencies: Tests of Possible Cognitive Mechanisms." *Journal of Experimental Psychology: Learning, Memory, and Cognition* 31 (4): 621–42.
- Hertwig, Ralph, and Timothy J. Pleskac. 2010. "Decisions from Experience: Why Small Samples?" *Cognition* 115 (2): 225–37.
- Hey, John D., Yudistira Permana, and Nuttaporn Rochanahastin. 2017. "When and How to Satisfice: An Experimental Investigation." *Theory and Decision* 83 (3): 337–53.
- Hoerl, Arthur E., and Robert W. Kennard. 1970. "Ridge Regression: Biased Estimation for Nonorthogonal Problems." *Technometrics* 12 (1): 55–67.
- Hogarth, Robin M., and Natalia Karelaia. 2007. "Heuristic and Linear Models of Judgment: Matching Rules and Environments." *Psychological Review* 114 (3): 733–58.
- Hume, David. 1739. *A Treatise of Human Nature*. Oxford Philosophical Texts. Oxford, UK: Oxford University Press, 2000.
- Jacobs, Rodney L., and Robert A. Jones. 1980. "Price Expectations in the United States: 1947–75." *American Economic Review* 70 (3): 269–77.
- Jamal, Karim, and Shyam Sunder. 2001. "Why Do Biased Heuristics Approximate Bayes Rule in Double Auctions?" *Journal of Economic Behavior and Organization* 46 (4): 431–35.
- James, W., and Charles Stein. 1961. "Estimation with Quadratic Loss." In *Proceedings of the Fourth Berkeley Symposium on Mathematical Statistics and Probability, vol. 1, Contributions to the Theory of Statistics*, ed. Jerzy Neyman, 361–79. Oakland, CA: University of California Press.
- Jamison, Dean T., and Lawrence J. Lau. 1973. "Semiordeers and the Theory of Choice." *Econometrica* 41(5): 901–12.
- Jung, Jongbin, Connor Concannon, Ravi Shroff, Sharad Goel, and Daniel G. Goldstein. 2020. "Simple Rules to Guide Expert Classifications." *Journal of the Royal Statistical Society: Series A (Statistics in Society)* 183 (3): 771–800.
- Kahneman, Daniel, and Amos Tversky. 1979. "Prospect Theory: An Analysis of Decision under Risk." *Econometrica* 47 (2): 263–92.
- Katsikopoulos, Konstantinos, Özgür Şimşek, Markus Buckmann, and Gerd Gigerenzer. 2020. *Classification in the Wild: The Science and Art of Transparent Decision Making*. Cambridge, MA: MIT Press.
- Kay, John, and Mervyn King. 2020. *Radical Uncertainty: Decision-Making for an Unknowable Future*. London: Bridge Street Press.
- Keynes, John Maynard. 1921. *A Treatise on Probability*. London: Macmillan and Co.
- Kleinberg, Jon, Jens Ludwig, Sendhil Mullainathan, and Ziad Obermeyer. 2015. "Prediction Policy Problems." *American Economic Review* 105 (5): 491–95.
- Knight, Frank H. 1921. *Risk, Uncertainty, and Profit*. Boston: Houghton Mifflin.
- Kool, Wouter, and Matthew Botvinick. 2014. "A Labor/Leisure Tradeoff in Cognitive Control." *Journal of Experimental Psychology: General* 143 (1): 131–41.
- Kőszegi, Botond, and Matthew Rabin. 2006. "A Model of Reference-Dependent Preferences." *Quarterly Journal of Economics* 121 (4): 1133–65.
- Krishnan, K.S. 1977. "Incorporating Thresholds of Indifference in Probabilistic Choice Models." *Management Science* 23 (11): 1224–33.
- Kunreuther, Howard. 1976. "Limited Knowledge and Insurance Protection." *Public Policy* 24 (2): 227–61.
- Lant, Theresa K. 1992. "Aspiration Level Adaptation: An Empirical Exploration." *Management Science* 38 (5): 623–44.
- Lant, Theresa, and Zur Shapira. 2008. "Managerial Reasoning about Aspirations and Expectations." *Journal of Economic Behavior and Organization* 66 (1): 60–73.
- Levinthal, Daniel, and James G. March. 1981. "A Model of Adaptive Organizational Search." *Journal of Economic Behavior and Organization* 2 (4): 307–33.
- Lewin, Kurt. 1935. *A Dynamic Theory of Personality: Selected Papers*. Redditch, UK: Read Books Ltd, 2013.
- Lichtman, Allan J. 2020. *Predicting the Next President: The Keys to the White House*. Lanham, MD: Rowman and Littlefield.
- Lieder, Falk, Thomas L. Griffiths, and Ming Hsu. 2018. "Overrepresentation of Extreme Events in Decision Making Reflects Rational Use of Cognitive Resources." *Psychological Review* 125 (1): 1–32.
- Lieder, Falk, Dillon Plunkett, Jessica B. Hamrick, Stuart J. Russell, Nicholas Hay, and Tom Griffiths. 2014. "Algorithm Selection by Rational Metareasoning as

- a Model of Human Strategy Selection." *Advances in Neural Information Processing Systems*, Vol. 27, ed. Z. Ghahramani, M. Welling, C. Cortes, N. Lawrence, and K.Q. Weinberger, 2870–78. Red Hook, NY: Curran Associates, Inc.
- Lioukas, S.K. 1984. "Thresholds and Transitivity in Stochastic Consumer Choice: A Multinomial Logit Analysis." *Management Science* 30 (1): 110–22.
- Loomes, Graham, Chris Starmer, and Robert Sugden. 2003. "Do Anomalies Disappear in Repeated Markets?" *Economic Journal* 113 (486): C153–66.
- Luan, Shenghua, Jochen Reb, and Gerd Gigerenzer. 2019. "Ecological Rationality: Fast-and-Frugal Heuristics for Managerial Decision Making under Uncertainty." *Academy of Management Journal* 62 (6): 1735–59.
- Luan, Shenghua, Lael J. Schooler, and Gerd Gigerenzer. 2011. "A Signal–Detection Analysis of Fast-and Frugal Trees." *Psychological Review* 118 (2): 316–38.
- Luan, Shenghua, Lael J. Schooler, and Gerd Gigerenzer. 2014. "From Perception to Preference and on to Inference: An Approach–Avoidance Analysis of Thresholds." *Psychological Review* 121 (3): 501–25.
- Lucci, Stephen, and Danny Kopec. 2016. *Artificial Intelligence in the 21st Century: A Living Introduction*. Dulles, VA: Mercury Learning and Information.
- Luce, Duncan R. 1956. "Semiororders and a Theory of Utility Discrimination." *Econometrica* 24 (2): 178–91.
- Malmendier, Ulrike, and Stefan Nagel. 2011. "Depression Babies: Do Macroeconomic Experiences Affect Risk Taking?" *Quarterly Journal of Economics* 126 (1): 373–416.
- Manski, Charles F. 2017. "Optimize, Satisfice, or Choose without Deliberation? A Simple Minimax-Regret Assessment." *Theory and Decision* 83 (2): 155–73.
- Manzini, Paola, and Marco Mariotti. 2007. "Sequentially Rationalizable Choice." *American Economic Review* 97 (5): 1824–39.
- Manzini, Paola, and Marco Mariotti. 2012. "Choice by Lexicographic Semiororders." *Theoretical Economics* 7 (1): 1–23.
- Manzini, Paola, Marco Mariotti, and Christopher J. Tyson. 2013. "Two-Stage Threshold Representations." *Theoretical Economics* 8 (3): 875–82.
- March, James G., and Herbert A. Simon. 1958. *Organizations*. New York: John Wiley and Sons.
- March, James G. 1981. "Footnotes to Organizational Change." *Administrative Science Quarterly* 26 (4): 563–77.
- March, James G. 1991. "Exploration and Exploitation in Organizational Learning." *Organization Science* 2 (1): 71–87.
- Marewski, Julian N., Wolfgang Gaissmaier, Lael J. Schooler, Daniel G. Goldstein, and Gerd Gigerenzer. 2010. "From Recognition to Decisions: Extending and Testing Recognition-Based Models for Multialternative Inference." *Psychonomic Bulletin and Review* 17 (3): 287–309.
- Markose, Sheri M. 2005. "Computability and Evolutionary Complexity: Markets as Complex Adaptive Systems (CAS)." *Economic Journal* 115 (504): F159–92.
- Markowitz, Harry. 1952. "Portfolio Selection." *Journal of Finance* 7 (1): 77–91.
- Martignon, Laura, and Ulrich Hoffrage. 2002. "Fast, Frugal, and Fit: Simple Heuristics for Paired Comparison." *Theory and Decision* 52 (1): 29–71.
- Martignon, Laura, Konstantinos V. Katsikopoulos, and Jan K. Woike. 2008. "Categorization with Limited Resources: A Family of Simple Heuristics." *Journal of Mathematical Psychology* 52 (6): 352–61.
- Mezias, Stephen J., Ya-Ru Chen, and Patrice R. Murphy. 2002. "Aspiration-Level Adaptation in an American Financial Services Organization: A Field Study." *Management Science* 48 (10): 1285–1300.
- Miller, Danny, and Ming-Jer Chen. 1994. "Sources and Consequences of Competitive Inertia: A Study of the U.S. Airline Industry." *Administrative Science Quarterly* 39 (1): 1–23.
- Miller, Joshua B., and Adam Sanjurjo. 2018. "Surprised by the Hot Hand Fallacy? A Truth in the Law of Small Numbers." *Econometrica* 86 (6): 2019–47.
- Mullainathan, Sendhil, and Jann Spiess. 2017. "Machine Learning: An Applied Econometric Approach." *Journal of Economic Perspectives* 31 (2): 87–106.
- Nash, John F. 1950. "The Bargaining Problem." *Econometrica* 18 (2): 155–62.
- Navarro-Martinez, Daniel, Graham Loomes, Andrea Isoni, David Butler, and Larbi Alaoui. 2018. "Boundedly Rational Expected Utility Theory." *Journal of Risk and Uncertainty* 57 (3): 199–223.
- Nelson, Richard R., and Sidney G. Winter. 1973. "Toward an Evolutionary Theory of Economic Capabilities." *American Economic Review* 63 (2): 440–49.
- Nelson, Richard R., and Sidney G. Winter. 1982. *An Evolutionary Theory of Economic Change*. Cambridge, MA: The Belknap Press.
- Oechssler, Jörg. 2002. "Cooperation as a Result of Learning with Aspiration Levels." *Journal of Economic Behavior and Organization* 49 (3): 405–09.
- Papadimitriou, Christos H., and Kenneth Steiglitz. 1998. *Combinatorial Optimization: Algorithms and Complexity*. Mineola, NY: Dover Publications.
- Papi, Mauro. 2012. "Satisficing Choice Procedures." *Journal of Economic Behavior and Organization* 84 (1): 451–62.
- Papi, Mauro. 2013. "Satisficing and Maximizing Consumers in a Monopolistic Screening Model." *Mathematical Social Sciences* 66 (3): 385–89.

- Papi, Mauro. 2018. "Price Competition with Satisficing Consumers." *International Journal of Industrial Organization* 58: 252–72.
- Payne, John W., James R. Bettman, and Eric J. Johnson. 1988. "Adaptive Strategy Selection in Decision Making." *Journal of Experimental Psychology: Learning, Memory, and Cognition* 14 (3): 534–52.
- Payne, John W., Dan J. Laughhunn, and Roy Crum. 1980. "Translation of Gambles and Aspiration Level Effects in Risky Choice Behavior." *Management Science* 26 (10): 1039–60.
- Payne, John W., Dan J. Laughhunn, and Roy Crum. 1981. "Further Tests of Aspiration Level Effects in Risky Choice Behavior." *Management Science* 27 (8): 953–58.
- Pazgal, Amit. 1997. "Satisficing Leads to Cooperation in Mutual Interests." *International Journal of Game Theory* 26 (4): 439–53.
- Pearl, Judea. 1984. *Heuristics: Intelligent Search Strategies for Computer Problem Solving. The Addison-Wesley Series in Artificial Intelligence*. Reading, MA: Addison-Wesley Publishing.
- Persson, Andreas, and Lynette Ryals. 2014. "Making Customer Relationship Decisions: Analytics v Rules of Thumb." *Journal of Business Research* 67 (8): 1725–32.
- Pfeifer, Phillip E. 1994. "Are We Overconfident in the Belief That Probability Forecasters Are Overconfident?" *Organizational Behavior and Human Decision Processes* 58 (2): 203–13.
- Rabin, Matthew. 2013. "Incorporating Limited Rationality into Economics." *Journal of Economic Literature* 51 (2): 528–43.
- Reis, Ricardo. 2006. "Inattentive Consumers." *Journal of Monetary Economics* 53 (8): 1761–1800.
- Reutskaja, Elena, Rosemarie Nagel, Colin F. Camerer, and Antonio Rangel. 2011. "Search Dynamics in Consumer Choice under Time Pressure: An Eye-Tracking Study." *American Economic Review* 101 (2): 900–926.
- Rothschild, Michael. 1974. "A Two-Armed Bandit Theory of Market Pricing." *Journal of Economic Theory* 9 (2): 185–202.
- Rubinstein, Ariel, and Yuval Salant. 2006. "A Model of Choice from Lists." *Theoretical Economics* 1 (1): 3–17.
- Salant, Yuval. 2011. "Procedural Analysis of Choice Rules with Applications to Bounded Rationality." *American Economic Review* 101 (2): 724–48.
- Sauerermann, Heinz, and Reinhard Selten. 1962. "Anspruchsanpassungstheorie der Unternehmung." *Zeitschrift Für Die Gesamte Staatswissenschaft Zeit für die gesamte Staatswissenschaft* 118 (4): 577–97.
- Savage, Leonard J. 1954. *The Foundations of Statistics*. Hoboken, NJ: John Wiley and Sons.
- Schmittlein, David C., Donald G. Morrison, and Richard Colombo. 1987. "Counting Your Customers: Who-Are They and What Will They Do Next?" *Management Science* 33 (1): 1–24.
- Selten, Reinhard. 1998. "Aspiration Adaptation Theory." *Journal of Mathematical Psychology* 42 (2–3): 191–214.
- Selten, Reinhard. 2001. "What Is Bounded Rationality?" In *Bounded Rationality: The Adaptive Toolbox*, edited by Gerd Gigerenzer and Reinhard Selten, 13–36. Cambridge, MA: MIT Press.
- Shannon, Claude E. 1950. "A Chess-Playing Machine." *Scientific American* 182 (2): 48–51.
- Shubik, Martin. 1958. "H.A. Simon, Models of Man: Social and Rational." *Journal of Political Economy* 66 (3): 273–74.
- Simon, Herbert A. 1947. *Administrative Behavior. A Study of Decision-Making Processes in Administrative Organization*. 1st ed. New York: Macmillan.
- Simon, Herbert A. 1955. "A Behavioral Model of Rational Choice." *Quarterly Journal of Economics* 69 (1): 99–118.
- Simon, Herbert A. 1956. "Rational Choice and the Structure of the Environment." *Psychological Review* 63 (2): 129–38.
- Simon, Herbert A. 1957. *Models of Man: Social and Rational*. Hoboken, NJ: John Wiley and Sons.
- Simon, Herbert A. 1959. "Theories of Decision Making in Economics and Behavioral Science." *American Economic Review* 49 (3): 253–83.
- Simon, Herbert A. 1979. "Rational Decision Making in Business Organizations." *American Economic Review* 69 (4): 493–513.
- Simon, Herbert A. 1981. *The Sciences of the Artificial*. 2nd ed. Cambridge, MA: MIT Press.
- Simon, Herbert A. 1985. "Human Nature in Politics: The Dialogue of Psychology with Political Science." *American Political Science Review* 79 (2): 293–304.
- Simon, Herbert A. 1989. "The Scientist as a Problem Solver." In *Complex Information Processing: The Impact of Herbert A. Simon*, edited by David Klahr and Kenneth Kotovsky, 375–98. Hillsdale, NJ: Lawrence Erlbaum Associates Publishers.
- Sims, Christopher A. 2003. "Implications of Rational Inattention." *Journal of Monetary Economics* 50 (3): 665–90.
- Şimşek, Özgür. 2013. "Linear Decision Rule as Aspiration for Simple Decision Heuristics." In *Advances in Neural Information Processing Systems*, Vol. 26, ed. C.J.C. Burges, L. Bottou, M. Welling, Z. Ghahramani, and K.Q. Weinberger, 2904–12. Red Hook, NY: Curran Associates, Inc.
- Şimşek, Özgür, and Marcus Buckmann. 2015. "Learning from Small Samples: An Analysis of Simple Decision Heuristics." *Advances in Neural Information*

- Processing Systems*, Vol. 28, ed. C. Cortes, N. Lawrence, D. Lee, M. Sugiyama, and R. Garnett, 3159–67. Red Hook, NY: Curran Associates, Inc.
- Smith, Adam. 1776. *An Inquiry into the Nature and Causes of the Wealth of Nations*. Oxford: Clarendon Press [1976].
- Smith, Adam. 1759. *The Theory of Moral Sentiments*, edited by D. D. Raphael and A. L. Macfie. Indianapolis, IN: Liberty Fund [1981].
- Smith, Vernon L. 1962. "An Experimental Study of Competitive Market Behavior." *Journal of Political Economy* 70 (2): 111–37.
- Smith, Vernon L. 2008. *Rationality in Economics: Constructivist and Ecological Forms*. Cambridge, MA: Cambridge University Press.
- Solow, Robert M. 1958. "Models of Man—Social and Rational by Herbert A. Simon." *Review of Economics and Statistics* 40 (1): 81–84.
- Starmer, Chris. 2005. "Normative Notions in Descriptive Dialogues." *Journal of Economic Methodology* 12 (2): 277–89.
- Sterman, John D. 1987. "Systems Simulation. Expectation Formation in Behavioral Simulation Models." *Behavioral Science* 32 (3): 190–211.
- Stigler, George J. 1961. "The Economics of Information." *Journal of Political Economy* 69 (3): 213–25.
- Stüttgen, Peter, Peter Boatwright, and Robert T. Monroe. 2012. "A Satisficing Choice Model." *Marketing Science* 31 (6): 878–99.
- Telser, L.G. 1973. "Searching for the Lowest Price." *American Economic Review* 63 (2): 40–49.
- Tibshirani, Robert. 1996. "Regression Shrinkage and Selection via the Lasso." *Journal of the Royal Statistical Society: Series B (Methodological)* 58 (1): 267–88.
- Todd, Peter M., and Geoffrey F. Miller. 1999. "From Pride and Prejudice to Persuasion Satisficing in Mate Search." In *Simple Heuristics That Make Us Smart*, edited by Gerd Gigerenzer, Peter M. Todd, and ABC Research Group, 287–308. Evolution and Cognition series. New York: Oxford University Press.
- Tversky, Amos. 1969. "Intransitivity of Preferences." *Psychological Review* 76 (1): 31–48.
- Tversky, Amos. 1972. "Elimination by Aspects: A Theory of Choice." *Psychological Review* 79 (4): 281–99.
- Tversky, Amos, and Daniel Kahneman. 1974. "Judgment under Uncertainty: Heuristics and Biases." *Science* 185 (4157): 1124–31.
- Tversky, Amos, and Daniel Kahneman. 1992. "Advances in Prospect Theory: Cumulative Representation of Uncertainty." *Journal of Risk and Uncertainty* 5 (4): 297–323.
- von Hayek, F.A. 1937. "Economics and Knowledge." *Economica* 4 (13): 33–54.
- von Hayek, F.A. 1945. "The Use of Knowledge in Society." *American Economic Review* 35 (4): 519–30.
- van Rooij, Robert. 2011. "Revealed Preference and Satisficing Behavior." *Synthese* 179 (S1): 1–12.
- Varian, Hal R. 1980. "A Model of Sales." *American Economic Review* 70 (4): 651–59.
- Varian, Hal R. 2014. "Big Data: New Tricks for Econometrics." *Journal of Economic Perspectives* 28 (2): 3–28.
- von Neumann, John, and Oskar Morgenstern. 1944. *The Theory of Games and Economic Behavior*. Princeton, NJ: Princeton University Press.
- Verrecchia, Robert E. 1982. "Information Acquisition in a Noisy Rational Expectations Economy." *Econometrica* 50 (6): 1415–30.
- Wald, Abraham. 1950. *Statistical Decision Functions*. New York: Wiley.
- Wall, Kent D. 1993. "A Model of Decision Making under Bounded Rationality." *Journal of Economic Behavior and Organization* 20 (3): 331–52.
- Wang, Yudong, Chongfeng Wu, and Li Yang. 2015. "Hedging with Futures: Does Anything Beat the Naïve Hedging Strategy?" *Management Science* 61 (12): 2870–89.
- Weitzman, Martin L. 1979. "Optimal Search for the Best Alternative." *Econometrica* 47 (3): 641–54.
- Wierzbicki, Andrzej P. 1982. "A Mathematical Basis for Satisficing Decision Making." *Mathematical Modelling* 3 (5): 391–405.
- Winter, Sidney G. 1971. "Satisficing, Selection, and the Innovating Remnant." *Quarterly Journal of Economics* 85 (2): 237–61.
- Winter, Sidney G. 2000. "The Satisficing Principle in Capability Learning." *Strategic Management Journal* 21 (10–11): 981–96.
- Witt, Ulrich. 1986. "Firms' Market Behavior under Imperfect Information and Economic Natural Selection." *Journal of Economic Behavior and Organization* 7 (3): 265–90.
- Wübben, Markus, and Florian von Wangenheim. 2008. "Instant Customer Base Analysis: Managerial Heuristics Often 'Get It Right.'" *Journal of Marketing* 72 (3): 82–93.
- Yatchew, Adonis. 1998. "Nonparametric Regression Techniques in Economics." *Journal of Economic Literature* 36 (2): 669–721.
- Yee, Michael, Ely Dahan, John R. Hauser, and James Orlin. 2007. "Greedoid-Based Noncompensatory Inference." *Marketing Science* 26 (4): 532–49.