THEORY-DRIVEN ENTREPRENEURIAL SEARCH

Ankur Chavda
Joshua S. Gans
Scott Stern

WORKING PAPER 32318

NBER WORKING PAPER SERIES

THEORY-DRIVEN ENTREPRENEURIAL SEARCH

Ankur Chavda
Joshua S. Gans
Scott Stern

Working Paper 32318
http://www.nber.org/papers/w32318

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
April 2024

Joshua Gans has drawn on the findings of his research for both compensated speaking engagements and consulting engagements. He has written the books, Information Wants to be Shared, The Disruption Dilemma, Prediction Machines, Power and Prediction on the economics of AI for which he receives royalties. He is also chief economist of the Creative Destruction Lab, a University of Toronto-based program that helps seed-stage companies, from which he receives compensation. He conducts consulting on anti-trust and intellectual property matters with an association with Charles River Associates and his ownership of Core Economic Research Ltd. He also has equity and advisory relationships with a number of startup firms. The views expressed herein are those of the authors and do not necessarily reflect the views of the National Bureau of Economic Research.

At least one co-author has disclosed additional relationships of potential relevance for this research. Further information is available online at http://www.nber.org/papers/w32318

Theory-Driven Entrepreneurial Search
Ankur Chavda, Joshua S. Gans, and Scott Stern
NBER Working Paper No. 32318
April 2024
JEL No. D81,D83,O32

ABSTRACT

How should theory-based entrepreneurs search for strategies to implement their ideas? The theory-based view of strategy posits that decision-makers hold theories about their environment premised on beliefs that should be actively tested. This causal framework, which underlies the theory-based view, also has implications for entrepreneurial search: the process by which entrepreneurs uncover strategies to implement their ideas. In this paper, we develop a Bayesian model where entrepreneurs update their beliefs as they conduct entrepreneurial search. We find several optimal behaviors for theory-based entrepreneurs such as reverting to a previous strategy after finding a relatively poor strategy and continuing to search after finding a relatively good strategy, which are missing when entrepreneurs lack such a theory-based approach. As these predictions align with examples of successful entrepreneurs, our findings both provide a method to empirically identify skilled entrepreneurs and demonstrate the usefulness of applying the theory-based view to entrepreneurial behavior more generally.

Ankur Chavda
HEC Paris
1 rue de la Libaration
Jouy-en-Josas 78351
France
chavda@hec.fr

Scott Stern
MIT Sloan School of Management
100 Main Street, E62-476
Cambridge, MA 02142
and NBER
sstern@mit.edu

Joshua S. Gans
Rotman School of Management
University of Toronto
105 St. George Street
Toronto ON M5S 3E6
and NBER
joshua.gans@rotman.utoronto.ca

1 Introduction

A main tenant of the theory-based view of strategy is that learning involves not only understanding whether beliefs about the environment are true or false but also how the re-evaluation of those beliefs changes causally linked assumptions and conjectures: learning alters one's theory of how the world works (Felin and Zenger, 2017; Ehrig and Schmidt, 2022). Although the theory-based view applies to a broad range of strategic decision-making, it is especially relevant to entrepreneurs, who must learn about environments where limited prior knowledge is available (Kirzner, 1997), as they often develop new markets, invent new technologies, or apply new organizational forms (Aulet and Murray, 2013; Botelho et al., 2021). Rather than address the uncertainty inherent in entrepreneurship (Bhide, 2003; Shane, 2003) by simply trying a particular action and moving on to the next, the theory-based view would have entrepreneurs interpret their actions as experiments that alter their beliefs and future behavior (Felin and Zenger, 2009). The entrepreneur's goal is not just to discover a good action but also to uncover information that better aligns their beliefs with the true state of their environment. Based on these updated beliefs, the revised theory then improves decision-making about which steps to take next (Felin et al., 2023).

A core contribution of the theory-based view to entrepreneurship is its guidance on how to learn about the entrepreneurial environment through experimentation. Being cognizant of the key premises that underlie a theory allows an entrepreneur to pinpoint the weakest link in the theory's causal logic and improve their overall confidence in the theory by testing that weakest link (Ehrig and Schmidt, 2022). Moreover, fully developing not only a particular theory's causal logic but also the several counter-theories that could be generated if one or more premises proved to be false can enable entrepreneurs to efficiently test multiple theories simultaneously (Camuffo et al., 2023). However, the theory-based view's value is not limited to a phase of experimental theorizing, when an entrepreneur would proactively select which steps to take next to validate their theory (Felin and Zenger, 2009). The theory-based view also has implications for core entrepreneurial behaviors such as the search for strategies because the existence or lack of a falsifiable theory fundamentally affects the decision-making of entrepreneurs.

Consider how the search for strategies differs between an entrepreneur taking a theory-based approach and one that does not. Given an entrepreneurial idea, an entrepreneur lacking theory will seek to uncover the best strategy they can find for their idea, given the constraints they face around the costs of experimenting with strategies, which might include prototyping costs or delayed time to market. Their goal is practical: exploit their idea to the largest extent reasonably possible (Felin et al., 2023). This Practice-based entrepreneur acknowledges there is uncertainty about how well each possible strategy may play out and will experiment to understand the payoff for different strategies. However, their evaluation of the idea, essentially the distribution of potential payoffs achievable by discoverable strategies, remains fixed because they lack a theory underlying that evaluation. There is no causal chain based on their environmental assumptions that can be broken and reevaluated through experimentation which would alter their evaluation of the idea. In effect, while an Practice-based entrepreneur can learn about specific strategies during the search process, they do not learn about their environment and, thus, the distinct value of their idea. Practice-based entrepreneurs experiment on their strategies, not their underlying ideas.

In contrast, a Theory-based entrepreneur holds a more abstract understanding of their environment, making experimentation proactive because it involves learning about the specific strategy tested *and* their environment and, consequently, their idea. This abstraction can involve varying degrees of sophistication, from simple hypothesis testing of different theories (Camuffo et al., 2020) to a causal model that links alternative theories through premises and objections to those premises (Ehrig and Schmidt, 2022), to a theory that is not depended on but instead inclusive of the environment, allowing entrepreneurial behavior to shape that environment (Rindova and Martins, 2023). But regardless of the level of abstraction along this continuum, a Theory-based entrepreneur will not only have a main working theory about their environment in whose premises they believe most strongly but also have counter-theories that stem from the possibility that one or more of those premises are false. Now, when they experiment with a particular strategy, they learn not only about their strategy but also about the likelihood each of their theories is correct. They can compare the experimental payoff of their tested strategy against the range of possibilities generated by their main working theory as well as against those generated by their counter-theories. Given a large discrepancy between the strategy's outcome and what their working theory would predict, they might question the premises behind their theory, increasing their confidence in a counter-theory that would have better predicted the observed outcome. As a result, their working theory might switch to one of their counter-theories. Hence, the learning generated by search can by itself lead to changing beliefs about which theory is most likely true, outside of an explicit experimental theorizing phase of entrepreneurial activity (Felin and Zenger, 2009).

In this paper, we present a model of entrepreneurial search that allows for this change in confidence among theories that is central to the theory-based view of strategy. In doing so, we bridge the gap between prior work on entrepreneurial search theory, which has mostly been concerned with how search should be conducted given an existing theory (e.g., Agrawal et al. (2021)) and the theory-based view of strategies, which is concerned with divining the truth among multiple theories. We model how entrepreneurs search for strategies to implement their ideas. An idea's potential manifests in the distribution of strategies discoverable for an idea: higher-value strategies are more likely to be found for ideas with higher potential. Entrepreneurs hold competing theories about their ideas, with each theory having implications for the idea's potential. Under one theory about how the entrepreneur's environment functions, the idea is good, and they will likely find a high payoff strategy. Under another theory, the idea is bad, and a high payoff strategy would be hard to find. In this setup, the exploration of a specific strategy can update the entrepreneur's fundamental beliefs about their environment. If the entrepreneur starts out with a theory derived from their beliefs that suggests a high-value idea but their first strategy yields disappointing results, the entrepreneur may change their beliefs, lending support to a counter-theory where the idea is of low value. Hence, the search for strategies can affect their beliefs, which alters the theory through which they evaluate their idea.

Our modelling approach uses a Bayesian specification, which allows for simultaneous search and learning (as in Gans et al. (2019)), sometimes referred to as "second-order probabilities" (e.g. Camerer and Weber (1992)). The value of a particular strategy is drawn from a known set of probability distributions, each corresponding to a specific theory, with the entrepreneur having a prior probability that any given theory is the correct one based on the relative strengths of their beliefs and objections to those beliefs. We contrast the search strategies used by a Theory-driven entrepreneur who understands that their underlying theory is based on premises that may be wrong and an Practice-based entrepreneur who does not have that understanding, presuming that there is a static distribution of strategy outcomes. In effect, a Theory-driven entrepreneur has a deeper understanding of the causal structure generating search outcomes than an Practice-based entrepreneur: Theory-driven entrepreneurs are essentially more skilled at entrepreneurial search.

For an example of each type of entrepreneur, consider the contrast between the McDonald brothers, who founded McDonald's and Ray Kroc, who turned McDonald's into a worldwide phenomenon. Both sets of entrepreneurs were aware of the runaway success of the first McDonald's restaurants: the McDonald's restaurants famously needed eight milkshake machines to meet demand, while a typical restaurant of the time needed just one (Kroc and Anderson, 1987). Ray Kroc thought McDonald's early success foreshadowed even greater future growth, while the McDonald brothers thought their restaurant was already a once-in-a-lifetime triumph. Ray effectively viewed the success of McDonald's early strategy as a Theory-driven entrepreneur would. Its wild success suggested the premise that McDonald's was a traditional restaurant to be false, lending support to the theory that McDonald's was a new kind of business, i.e., fast food, with much larger potential than a traditional restaurant.

In contrast, the McDonald brothers saw early success as an excellent draw given their idea of founding a traditional restaurant, behaving as Practice-based entrepreneurs and becoming satisfied with their discovered high payoff strategy.

As in this example, our results show that whether an entrepreneur is Theory-driven or Practice-based has a profound impact on optimal search behaviors. While an entrepreneur drawing from a single distribution, i.e., lacking theory, will search sequentially using a monotonic stopping rule – that is, eventually implementing the first strategy whose outcome is above a threshold – an entrepreneur who understands they are drawing from a set of distributions, i.e., multiple theories and counter-theories, may employ a non-monotonic stopping rule – continuing search after a high draw in situations where a lower draw would have resulted in stopping. This has several major consequences. First, it implies Theory-based entrepreneurs can be more tenacious when searching for strategies. When other Practice-based entrepreneurs might be satisfied with a discovered strategic payoff and commit to that strategy in order to lock in their gains, a Theory-based entrepreneur pursuing the same fundamental idea might continue to mine that idea for an even larger payoff. Second, it suggests that Theory-based entrepreneurs could 'pivot' differently from Practice-based entrepreneurs. Practice-based entrepreneurs will always pursue new strategies when they uncover that an existing strategy fails their expectations. In contrast, a Theory-based entrepreneur would be willing to revert to a prior strategy despite having previously dismissed its payoff as being unsatisfactory. Third, we also show only Theory-driven entrepreneurs value expert opinion for their ability to bootstrap beliefs about the true distribution of strategies. This suggests that heterogeneity in preferences for venture capitalists among entrepreneurs may, in part, stem from the mindset of the entrepreneur. Theory-driven entrepreneurs are more likely to value venture capitalists with prior experience in their idea's context because they can learn from those venture capitalists. Thus, the type of entrepreneur not only affects their search behavior, but also how they match with venture capital firms.

This paper makes several contributions to the literature on entrepreneurship. First and most directly, entrepreneurs in high-growth industries are likely to benefit from mimicking the optimal search behaviors described in our paper for Theory-based entrepreneurs. Such industries are often characterized by a high degree of uncertainty, increasing the likelihood an entrepreneur's premises may be wrong and therefore increasing the value of re-evaluating those premises based on the information gained during search (Gans et al., 2018; Agrawal et al., 2021; Gans, 2023). Secondly, in such uncertain environments, one can expect higher-skilled entrepreneurs to take a Theory-based approach and naturally exhibit the optimal search behaviors predicted in our study. This effectively allows for the identification of highly skilled entrepreneurs, addressing one of the main empirical research challenges in entrepreneurship, as directly observing entrepreneurial skills is difficult. Prior empirical work attempting to measure the effect of entrepreneurial skill on venture success relies on controlling for other known success factors, equating any residual effect to skill (Guzman and Stern, 2020). However, other uncontrolled factors may be driving such residual effects in lieu of skill (Frederiks et al., 2019). Future empirical work could instead use our theoretically predicted entrepreneurial behaviors as an indicator for skill, thereby more directly examining correlations between skill and venture outcomes. Third, we identify another mechanism by which venture capitalists may provide value to entrepreneurs (Kaplan and Strömberg, 2001; Bernstein et al., 2016): aid in resolving theory uncertainty, which only Theory-based entrepreneurs will appreciate.

Our paper also contributes to the broader innovation literature studying Knightian uncertainty, regardless of whether the context is specifically entrepreneurial in nature (Knight, 1921). From a formal theoretical perspective, the exploration of Knightian uncertainty largely follows the framework of (Bewley, 2002), who considers optimal search when search results in possibilities that the decision maker previously did not consider (i.e., placed zero probability weight on occurring). These studies of 'extreme' Knightian uncertainty focus on the possibility of unanticipated information impacts search and decision-making (Chou and Talmain, 1993; Nishimura and Ozaki, 2004; Schröder, 2011; Karni and Vierø, 2017; Cerreia-Vioglio et al., 2020) . Our modelling approach moves away from the idea that learning can uncover information that is literally outside the realm of possibility, as the kind of boundless, unheard-of opportunity attained by "European exploiters among primitive people (sic)" (Knight, 1921, III.IX.35) used to motivate extreme versions of Knightian uncertainty no longer exists. Moreover, individuals are constrained by the envelope of the set of beliefs they hold, even if they can generate new theories by changing those beliefs (Felin et al., 2023). To illustrate, when Eric Schmidt took the helm of Google in its early days, he knew very well what wild success would look like: it meant turning the startup into his arch-nemesis, Microsoft (Gomes, 2000). The range of outcomes was bounded, and the approximate outcome of Eric's theory about Google's future was known. We contend that our less extreme approach to Knightian uncertainty better reflects real-world cases of uncertainty and can yield additional insights on optimum behavior in uncertain environments beyond the ones for entrepreneurial search presented here.

To yield these contributions, our paper brings together three distinct but related approaches to conceptualizing entrepreneurial decision-making. First, a domain of workplaces malleable theory at the centre of entrepreneurial decision-making (Hsieh, 2010; Felin and Zenger, 2009). This work takes the view that the Knightian style uncertainty faced by entrepreneurs (Knight, 1921) makes finding optimal decisions across the landscape of potential options intractable without some theory to guide decision-making. These theories however are not fixed, they are revised over time as entrepreneurs test their theories and learn which theoretical premises hold true and which are falsified. Entrepreneurship is effectively described as a cycle of theory refinement, where the causal structure of a theory helps entrepreneurs identify which theoretical beliefs need further testing and how to proceed given both positive and negative evaluations of those beliefs (Ehrig and Schmidt, 2022; Wuebker et al., 2023). In our paper, we draw on this idea of theory-driven decision-making but relax the literature's assumption that entrepreneurial learning primarily occurs in an 'experimentation' phase where entrepreneurs are focused on testing the validity of their theory (Felin and Zenger, 2009; Wuebker et al., 2023). Instead, we allow for entrepreneurial activity to be less structured, as suggested by recent empirical studies (Arikan et al., 2020; Rapp and Olbrich, 2023) that find learning occurs across a wide variety of entrepreneurial activities, including entrepreneurial search as studied here.

Second, we rely on the formal literature on search, which evaluates optimal search behavior under distributional uncertainty (Chow and Robbins, 1963; McCall, 1965). A key insight of that literature is that, in contrast to a search with a known distribution where the search only stops above a pre-determined 'cut-off', uncertainty over the distribution results in a more complex stopping criterion, for example, leading to the possibility of stopping search after a sufficiently low draw (Rothschild, 1974). Work building on Rothschild (1974) primarily focuses on identifying cases of distributional uncertainty where a cut-off rule could nonetheless be established (Weitzman, 1978; Adam, 2001). This seach theory has been applied to several domains outside of entrepreneurship such as job search (Johnson, 1978; Morgan, 1983; Talmain, 1992; Adam, 2001), pricing (Rosenfield and Shapiro, 1981), and strategic investments by established firms (Tonks, 1983; Bernanke, 1983; Jovanovic and Rob, 1990; Keller et al., 2005). Our model also builds on Rothschild (1974) but instead explicitly focuses on cases when the cut-off rule fails to hold, in contrast to much of the prior work in this area.

Finally, there is a literature that focuses directly on how entrepreneurs discover and exploit new opportunities in an uncertain environment (Kirzner, 1997; Shane and Venkataraman, 2000). This uncertainty distinguishes the entrepreneurial case from the more general search for strategies literature (Levinthal, 1997; Rivkin, 2000), as entrepreneurs must both determine the properties of their environment while simultaneously identifying the most promising opportunities within that environment. Since the pursuit of a given opportunity by construction involves bringing together resources in a novel way to create and capture economic value, understanding how entrepreneurs assess novel opportunities (and different ways of pursuing that opportunity) is a central element of entrepreneurship (Sarasvathy, 2001; Felin and Zenger, 2009; Smith and Cao, 2007). A central insight in the study of opportunity identification and pursuit is that uncertainty impacts nearly every aspect of the discovery and exploitation of opportunities (Shane and Venkataraman, 2000), including the identification of promising domains and specific opportunities (Ward, 2004; Fiet et al., 2005; Baron and Ensley, 2006; Fiet, 2007; Fiet and Patel, 2008; Gruber et al., 2008), the interplay between learning and entrepreneurial action (Gaglio and Katz, 2001; Choi et al., 2008; Rapp and Olbrich, 2023; Agrawal et al., 2021; Gans, 2023), and the valuation of specific opportunities (Norton and Moore, 2002; Chen et al., 2018; Gambardella et al., 2022).

In relation to this last literature, our focus here is on how, given an uncertain idea or opportunity, entrepreneurs discover alternative ways to execute that idea, in other words, search for new strategies. We consider how entrepreneurs uncover such strategies (Alvarez and Barney, 2007) as opposed to how they plan to execute on an opportunity per se (Delmar and Shane, 2003; Shane and Delmar, 2004) or the importance of limiting search to domains of high expertise or potential learning (Fiet et al., 2005; Fiet, 2007). Prior work on the search for strategies effectively focused on cases where the aforementioned cut-off rule results from assumptions about the type of uncertainty faced by an entrepreneur (Mosakowski, 1997; Angus, 2019). We argue that our knowledge of entrepreneurial search can be enhanced by exploring the implications of relaxing this assumption, as it predicts several behaviors observed in entrepreneurs that have thrived under significant environmental uncertainty, including shifting the "anchor" used as a baseline for search (Bhardwaj et al., 2006) and changing search behavior due to learning about the distribution of possible outcomes (Arentz et al., 2013; Goldstein et al., 2020).

The rest of this paper proceeds as follows. Section 2 then uses a specific distributional example to outline our main results. Section 3 provides a general statement of our results by building a formal model. Section 4 presents a further implication about how an idea's value may actually be lower when pursued by a Theory-based entrepreneur. A final section concludes.

2 Basic Model

In Section 3, we will provide formal theorems of our main results that apply to a general set of probability distributions over the potential outcomes resulting from the search for strategies. However, in this section, we provide the intuition behind our results, contrasting the behavior of a Theory-driven and Practice-based entrepreneur under a specific distributional example. In addition, we highlight anecdotes of entrepreneurial behavior which match our theoretical predictions.

We consider an entrepreneur who has an idea for a new venture. Following Gans et al. (2019), a set of possible strategies can be implemented to generate a payoff for the venture. The entrepreneur is resource-constrained and so can only implement one strategy. However, the entrepreneur can 'test' any strategy in sequence, one at a time, before choosing which one to implement.

Specifically, suppose that an entrepreneur will implement a strategy after testing one or more of them if that strategy is predicted to yield a positive (non-zero) return (i.e., the entrepreneur's outside option has an expected payoff of 0, and the entrepreneur is risk-neutral). There is a large set of strategies. By paying a cost, c, a strategy i can be tested, and the entrepreneur receives a signal, x_i, that is a positive real number that gives the expected return the entrepreneur will receive should strategy i be implemented. For this example, we set $c = 0.15$.

There is a set of probability distributions, $\mathcal{D}$, that govern the actual return for strategies. For this illustration, we assume that there are two possible distributions (i.e., $|\mathcal{D}| = 2$). In one distribution, the idea is an 'ordinary' idea and involves a distribution of returns which is normal with mean 1 and variance 1/2; i.e., $N(1, 1/2)$. For the other distribution, the idea is 'exceptional' and involves a distribution with the same variance but a higher mean of 3; i.e., $N(3, 1/2)$.

We compare two entrepreneurial types: Theory-driven and Practice-based. A Theory-driven entrepreneur is aware of several possible theories that may potentially be true, although they know only one of those theories is actually true. Each of those theories implies a specific distribution of payoffs for discoverable strategies. Based on their existing beliefs, they hold priors on the likelihood each of those theories is true. In our example, the entrepreneur has two competing theories through which they evaluate their idea. Under one theory, the idea is exceptional, with strategies distributed $N(3, 1/2)$, and under the other, the theory is ordinary, with strategies distributed $N(1, 1/2)$. Here, we assume that a Theory-driven entrepreneur has a prior probability of 0.75 on the theory, which makes the idea ordinary, and a prior of 0.25 on the theory where the idea is extraordinary.

By contrast, an Practice-based entrepreneur lacks any sense of theory and, therefore, holds a fixed belief about the distribution of returns for their idea. To ensure we are making an apples-to-apples comparison between these two entrepreneurial types, we set the fixed distribution held by the Practice-based entrepreneur equal to the initial mixture distribution held by the Theory-driven entrepreneur. Effectively, both types of entrepreneurs start from the same place, but the Theory-based one learns about the generating distribution while the Practice-based one does not. Given the distributional assumptions used in our example, Figure 1 plots the distribution of strategies for each of these types of ideas (Panel A) as well

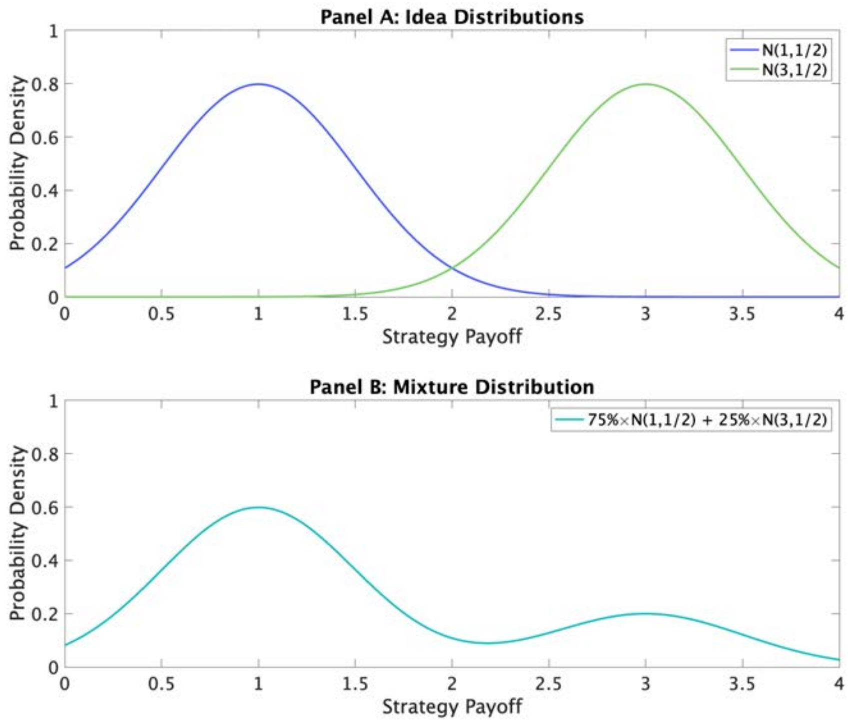

Figure 1: Distributions of Strategy Outcomes

as the mixture that results from their beliefs (Panel B).[1]

2.1 Don't rest on your laurels

Consider first an Practice-based entrepreneur who assumes that draws come from the mixture distribution in Figure 1B. As is well known (e.g., McCall (1965)), the optimal search behaviour will be to set a cut-off value and choose to implement a strategy as soon as one is found that exceeds that value. It can be readily calculated that so long as draws have outcomes less than around 2.4, it is worthwhile to continue searching [2]. Once a strategy with an outcome above that threshold is discovered, search stops. At that point, the probability of finding a better strategy is too low to warrant continued expenditure of c. Figure 2A depicts this stopping rule by showing the region where discovered strategy outcomes lead to

[1]The weight placed on the exceptional distribution is done here to make the example clearer. It is useful to note that the Practice-based entrepreneur's distributional prior can be very close to the Theory-driven's prior for ordinary ideas; e.g., when there is a very low prior probability held by the Theory-driven entrepreneur that the idea is exceptional. Even such a relatively small change can lead to significant changes in potential observed search behaviour.

[2]One must solve for the value of x where $x = \int_x t dF(t) + \int^x x dF(t) - c$ given distribution with cdf $F(t)$

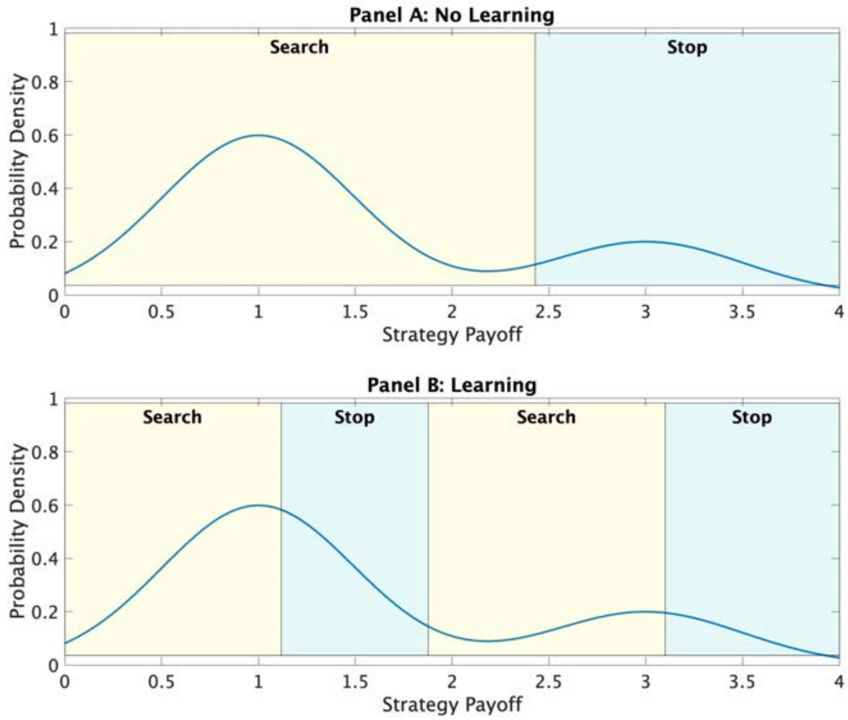

Figure 2: Optimal Stopping Decision

continued search versus stopping for an Practice-based entrepreneur.

Turning to a Theory-driven entrepreneur, a different – non-monotone – search strategy is optimal with the conditions for stopping versus continual searching arising at multiple rather than a single point. This is because discovering a high outcome strategy can, counter-intuitively, lead to *more* search when search additionally produces knowledge about the underlying distribution of potential outcomes. As depicted in Figure 2B, search may continue until outcomes in excess of 3.1 arise [3]. Thus, there exists a range of high-value outcomes where the Theory-driven entrepreneur will continue searching while the Practice-based entrepreneur stops and implements the last strategy tested. Intuitively, the Practice-based entrepreneur does not learn from those draws that it is likely that the true distribution is anything other than the mixture distribution, whereas the Theory-driven entrepreneur's posterior probability that the distribution is exceptional increases. That induces the return to further search as their assessment that higher valued outcomes are likely has increased.

This theorized Theory-driven behavior is perhaps best exemplified by the founding of

[3]Theory-driven stopping rule calculated by finding the fixed point in the value function $v(z;\lambda) = \max\{z, \int v(\max\{z,x\};\pi(\lambda,x))dF(x;\lambda) - c\}$, where λ is the prior and $\pi(\lambda,t)$ updates that prior based on draw t. Contraction mapping code available upon request.

Amazon. Prior to founding Amazon, Jeff Bezos was tasked with exploring the idea of using the internet to reach customers, specifically for stock and bond trading, as he worked at D.E. Shaw, a Wall Street hedge fund (Stone, 2013). These early explorations led to some successes; for example, D.E. Shaw created a stock trading site that was later sold to Merrill Lynch (Buckman and Gasparino, 1999). However, Bezos viewed the potential of such early strategies leveraging the internet as indicative that more was possible. While an Practice-based entrepreneur might have settled for the first good Internet strategy, as it might have seemed attractive compared to non-internet businesses, Bezos continued searching for better strategies. He considered 20 different potential products (ranging from music to electronics to toys) before selecting books as the vertical in which to launch Amazon (Bezos, 1997).

To summarize, while an Practice-based entrepreneur's search rules are monotonic, this may not be the case for a Theory-driven entrepreneur. Note that this also implies that if there are relatively low-valued draws, the Theory-driven entrepreneur chooses to implement those strategies rather than continue the search as they assess that it is likely that the idea is ordinary. By contrast, an Practice-based entrepreneur may continue searching. Thus, a Theory-driven entrepreneur does not rest on their laurels after discovering a good strategy but may relent after discovering a mediocre strategy.

2.2 Test then go back

Because an Practice-based entrepreneur has a monotone search strategy, as soon as they locate a strategy above the optimal stopping threshold, that last strategy tested is the strategy implemented. Figure 3 plots this behaviour where the 25% prior on an exceptional distribution generates the same underlying distribution as Figure 1A. Thus, an Practice-based entrepreneur sets a threshold of just below 2.5. Prior to any draws, the expected payoff from an additional test is positive, specifically just below 2.5. If the first draw yields a value of 2, that expected payoff falls to 0.5 as 2 represents a more favourable alternative return than the outside option of 0. As it is positive, another search is worthwhile. Suppose this yields a draw of 0.5. The returns from additional searches remain unchanged. Thus, the Practice-based entrepreneur keeps searching until a draw of more than around 2.5 is found.

By contrast, suppose that a Theory-driven entrepreneur received a draw of 2 followed by another of 0.5. The first draw lowers the expected return to another test from 1.6 to 0.16. Importantly, 2 could be almost equally generated by either an ordinary or exceptional distribution. Thus, the entrepreneur's prior on the distribution changes very little. The marginal test return is positive, but, in this case, a draw of 0.5 causes the posterior probability that the idea is exceptional to fall well below the prior. 0.5 is unlikely to come from an

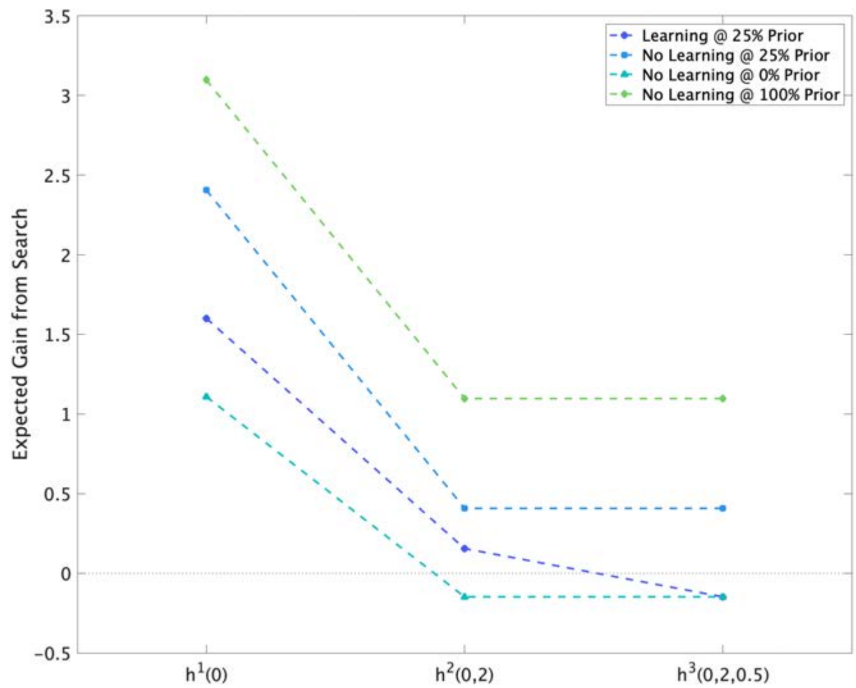

Figure 3: Sequence of Stopping

exceptional distribution, so there is a high probability that the distribution is ordinary. The previously tested strategy of 2 is an excellent outcome for that distribution. Thus, not only does the expected return to future testing fall below 0, making it optimal to stop, the best strategy to implement was the first, not the second one, tested. The Theory-driven entrepreneur is then observed to test and go back to implement a previously tested strategy, something that an Practice-based entrepreneur would not do.

As an example, consider Google's early shifts in strategy to execute its ideas on internet search. Originally a stand-alone search portal, Google's founders in 1999 struggled to find a revenue model and decided it would best to sell Google's underlying technology, Backrub, to an established search portal. However, this strategy failed, as Altavista, Yahoo! and Excite, three of the largest search portals of the period, declined the purchase (Levy, 2011). Upon learning that the market for their technology was limited, they returned to their original strategy of being a standalone portal that used Backrub to serve search results to consumers. This kind of reversion to a previous strategy is predicted for Theory-driven entrepreneurs but not Practice-based ones.

Note that this reversion behavior cannot be replicated for an Practice-based entrepreneur simply by changing the prior belief that the idea is exceptional. Figure 3 also plots the cases when the entrepreneur has a 0% and 100% prior of an exceptional distribution to illustrate this point. Therefore, models of entrepreneurial search that rely solely on heterogeneous priors without incorporating learning about the underlying distribution cannot account for an entrepreneur returning to a previously discovered strategy. However, a Theory-driven entrepreneur who does learn in this way can justify cutting losses and stopping search after realising that the best strategy for an idea has likely already been found.

2.3 Value of True Knowledge

By definition, the Practice-based entrepreneur is certain of their idea's distribution of strategies and, therefore would not value additional knowledge about the distribution of those strategies. In contrast, a Theory-driven entrepreneur, in fact, would be willing to pay to reduce the uncertainty surrounding the distribution of strategies. The value of knowing the true distribution derives from two sources: first, preventing unnecessary searches when the best-discovered strategy is, in fact, a good one conditional on the true distribution and second, continuing search when the best-discovered strategy is poor conditional on the true distribution. This value varies with the entrepreneur's original belief and their current best option, as depicted in Figure 4 based on the above example.

If the best-discovered strategy has a relatively high outcome, above 3, knowledge of the true distribution makes little difference for a Theory-driven entrepreneur. Regardless of whether the idea is exceptional or ordinary, the entrepreneur should stop searching; therefore, learning about the distribution is irrelevant. Similarly, when the discovered outcome is relatively low, below 1, understanding the true distribution makes little difference to the Theory-driven entrepreneur, as they will search regardless and thereby learn about the underlying distribution anyway. Hence, Theory-driven entrepreneurs with intermediate best-discovered strategies will place higher value on knowing the true distribution.

This interplay between an ordinary and an exceptional idea is further elucidated through Figure 4's visualization of how the belief that an idea is exceptional affects the value of the true knowledge about the distribution. This value is low when the entrepreneur is fairly certain the idea is either ordinary or exceptional, the two extremes of belief. In Figure 4, when the entrepreneur's belief is lower than 5% or greater than 95%, they view making a mistake in the search process, either stopping too soon or too late, as unlikely and therefore discount the value of knowing the true distribution of strategies. Overall, this suggests that Theory-driven entrepreneurs, in general, will value the advice of potential experts such as venture capitalists more than Practice-based entrepreneurs and that within Theory-driven

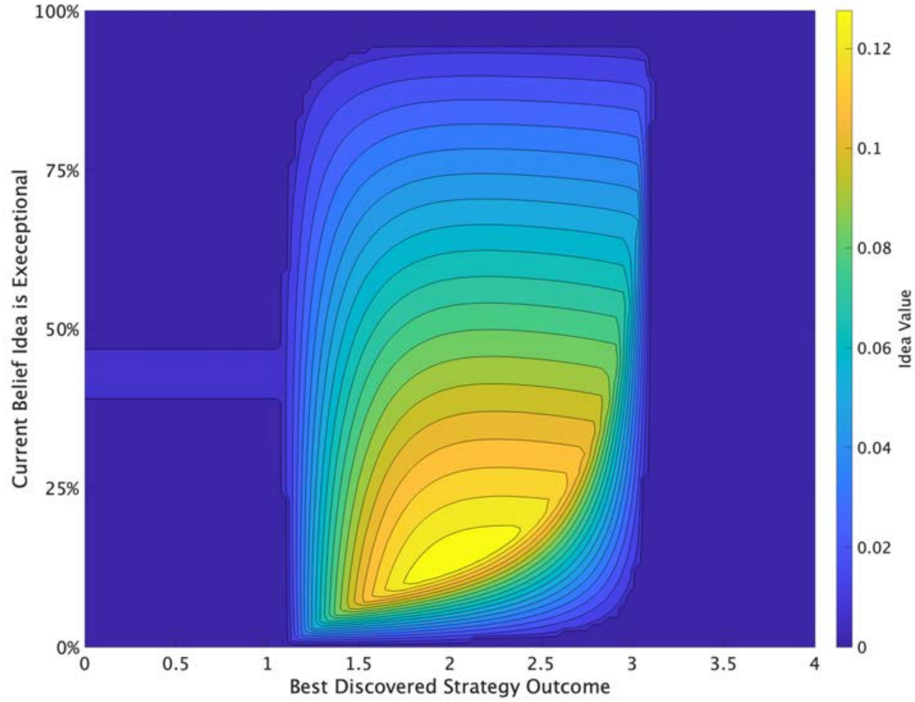

Figure 4: Gain from learning true distribution

,

entrepreneurs, there can exist significant heterogeneity in how much they value that advice.

Such expert advice likely enabled Zappos's success in the online shoe space. Tony Hsieh, the founder of LinkExchange, was a key early investor in Zappos. At the time, Zappos was struggling with its online shoe strategy of drop shipping: getting shoe manufacturers to share their inventory information and ship shoes directly to customers on behalf of Zappos. Hsieh, who at the time ran a venture capital firm invested in several internet startups, saw greater potential in Zappos than the founders did. He encouraged them to take a more aggressive strategy, having their own inventory and delivering phenomenal end-to-end customer service (Hsieh, 2010). Hsieh's expert knowledge of what was possible on the internet pushed Zappos into searching again, leading to its later success and Hsieh joining the company as its CEO.

3 General Results

We now demonstrate that the results for the illustrative example just examined hold in a general setting, in particular, for any set of distributions satisfying some regularity assumptions. Proofs for each proposition here are in the Appendix.

Our model involves an entrepreneur testing strategies $i \in \mathcal{S}$ sequentially at a cost per test of c. After testing each strategy, the entrepreneur decides whether to implement any tested strategy or continue testing by selecting a new strategy from $\mathcal{S}$.

The distribution of strategy outcomes at each stage is not independent and not identical, represented by random variables $X_i \in \mathbb{R}_+$ whose probability distributions form a set $\mathcal{D}$. Having conducted n tests, the entrepreneur can choose to implement a strategy with the highest realisation x_i amongst the n tested strategies. This results in an expected payoff of:

$$y = \max\{x_1, \ldots x_n\} - nc \tag{1}$$

We impose the following regularity assumptions on the set of distributions.

Assumption 1 *There exists a random variables $\overline{X}$ and $\underline{X}$ with distributions such that $\overline{X}$ first order stochastically dominates all marginal distributions $Pr(X_i = x_i \mid X_1 = x_1, \ldots, X_{i-1} = x_{i-1})$ generated from $\mathcal{D}$ and $\underline{X}$ is similarly dominated by all marginal distributions generated from $\mathcal{D}$.*

Assumption 2 *$\overline{X}$ and $\underline{X}$ have finite first moments.*

Assumption 3 *$\overline{X}$ has a finite second moment.*

Assumptions 1 and 2 ensure $y < \infty$ while Assumption 3 additionally establishes that $E[y] < \infty$, both necessary for the value of entrepreneur's idea to be finite. Essentially, this implies an entrepreneur's search process does eventually stop. Note that although we formulate search as a maximisation across all the draws as random variables, we can arbitrarily add an initial outside option set as x_0 without changing any of our results.

The next assumption formally defines a Theory-driven entrepreneur.

Assumption 4 *A Theory-driven entrepreneur's beliefs about the true distribution of strategy payoffs for the idea can be represented by an element of a concave set Λ of probability measures over $\mathcal{D}$.*

Definition 1 *A Theory-driven entrepreneur updates beliefs about the true distribution of strategy payoffs for an idea after each search while an Practice-based entrepreneur believes that the distribution is fixed during search.*

Assumption 4 ensures that a Theory-driven entrepreneur can find a new belief after search that is consistent with their prior beliefs and the result of that search.

With these assumptions, we can state the following result that confirms that learning about distributions is valuable to a Theory-driven entrepreneur. (Note that, by assumption, such learning has no value to an Practice-based entrepreneur.)

Proposition 1 *Under Assumptions 1 to 4, the value of an idea weakly increases when the idea's true distribution is revealed.*

Proposition 1 confirms that a Theory-driven entrepreneur would value an idea more if they knew the true distribution of the idea's value. Such information is valuable in two respects. First, the true distribution may be less favorable than the distribution implied by the entrepreneur's prior. The revelation would allow search to be terminated earlier, lowering the overall cost of search with a proportionality smaller drop in expected payoff, yielding a net benefit. Alternatively, the true distribution may be more favorable than the entrepreneur's prior implied. The revelation would then prevent premature stopping. Hence, knowing the distribution can only improve the value of the idea to the entrepreneur.

Next, we impose a set of parametric restrictions to allow us to characterise entrepreneurial strategy. Let $\mathcal{D}$ be indexed by θ and the distribution of beliefs over those distributions, Λ, in turn by indexed by λ. After each search, λ gets updated conditional on the discovered strategy's value x_i through the function $\pi(\lambda, x)$

Assumption 5 *The family of distributions $\mathcal{D}$ displays first-order stochastic dominance in θ, and the family of distributions Λ displays first-order stochastic dominance in λ.*

Assumption 6 *Belief λ conditional on discovering a strategy of value x updates through $\pi(\lambda, x)$, which is increasing in both its arguments.*

Assumption 7 *Given current belief λ, there exists discoverable strategies with outcomes x', and x'' such that $x' < x''$ and $\lambda' \equiv \pi(\lambda, x') < \lambda'' \equiv \pi(\lambda, x'')$*

Assumption 5 and 6 ensure the expected value of an idea is weakly increasing in outcome z and parameter λ, as well as weakly decreasing in cost c. Assumption 7 states that the necessary variation in belief may result from the next search.

This leads to our first result regarding the monotonicity of search:

Proposition 2 *Under Assumptions 1-7, there exists a discovered strategy outcome z and search cost c such that it is optimal to stop after the discovery of a lower outcome strategy but continue searching after the discovery of a higher outcome strategy.*

For Theory-driven entrepreneurs, search has two effects on the expected gain from additional search: a best outcome effect and a learning effect. The best outcome effect can only make the additional search less valuable; if a newly discovered strategy has a higher outcome than any previously discovered strategy, the entrepreneur is less likely to conduct further search simply because the possibility of finding a better outcome than the new best has been reduced. When the distribution of outcomes is known with certainty, i.e. for Practice-based entrepreneurs, this is the only effect of search, resulting in the simple cutoff-stopping rule. However, for Theory-driven entrepreneurs, the learning effect is also present. Search provides information about the likelihood strategies follow a particular distribution and, therefore, insight into whether a particular theory is true. Proposition 2 states when the outcome of a strategy suggests a poor draw from a good distribution rather than a good draw from a poor distribution, the learning effect can outweigh the best outcome effect, increasing the value of search and therefore resulting in additional search.

In addition, the same set of assumptions allows for the reverse logic: termination when the outcome of a strategy suggests a good draw from a poor distribution rather than a poor draw from a good distribution.

Proposition 3 *Under Assumptions 1 to 7, there exists a best-discovered strategy outcome z, and search cost c such that the best strategy outcome upon the termination of search will not be the last discovered strategy.*

Proposition 3 is the result of the downward revision of an entrepreneur's beliefs about the idea's value. After attaining a surprisingly good draw leading to additional search, the next draw might be relatively poor. This makes the Theory-driven entrepreneur realise the previous draw was the best they are likely to get, causing a reversion to that previously discovered strategy and the termination of search. The entrepreneur effectively learns from failure.

4 Uncertainty and the Value of Ideas

Given an idea, does approaching that idea with a Theory-driven versus Practice-based mindset change the value of that idea? The short answer is yes. However, which mindset values the idea more is not straightforward. While we are not able to generally characterise the difference in the value placed on the idea between a Theory-driven versus an Practice-based entrepreneur, here we continue the example from Section 2 to illustrate the nuances between these two approaches to entrepreneurial search and how an idea is therefore valued.

We already observed in Figure 3 that for a 25% prior probability, the expected return to the Theory-driven entrepreneur was lower than that for an Practice-based entrepreneur. Figure 5 shows the expected return to evaluating an idea on the right Y-axis for different levels of the prior probability that the idea comes from an exceptional distribution. At

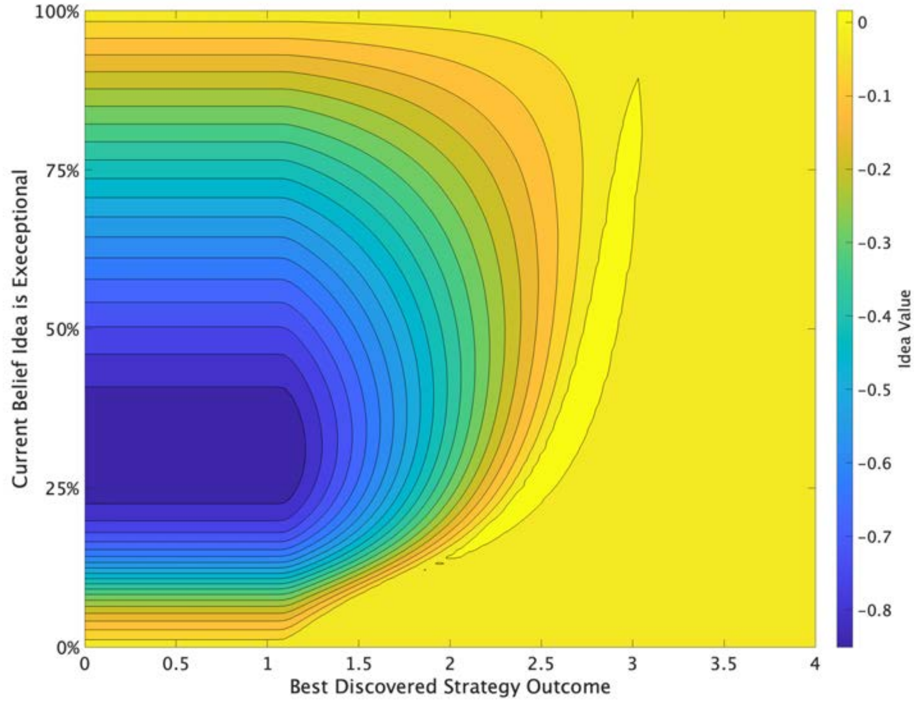

Figure 5: Difference between Theory-driven and Practice-based idea values

lower best-discovered strategies, the Theory-driven entrepreneur values the idea less than the Practice-based one.

Intuitively, this is related to the potential non-monotonicity of the Theory-driven entrepreneur's search strategy. For an Practice-based entrepreneur, additional draws do not alter their beliefs regarding the overall distribution of value. Therefore, if a draw is for a relatively low value, they will keep searching and testing new strategies: the value of the idea remains unchained after such low draws. By contrast, a Theory-driven entrepreneur receiving low-value draws will update their beliefs regarding the distribution – in this case, placing more weight on the idea of being ordinary rather than exceptional. Thus, they are more likely to expect to stop searching after receiving low draws. This both lowers the frequency we would observe the Theory-driven entrepreneur searches relative to the Practice-based entrepreneur and lowers the Theory-driven entrepreneur's value of the idea relative to the Practice-based entrepreneur.

To see why this is the case, let us briefly suppose that the two relevant distributions for X_i are $N(1,0)$ and $N(3,0)$, respectively. In other words, the 'distributions' are really point values of 1 or 3. In this case, a Theory-driven entrepreneur with a single draw will know the distribution. Thus, if their prior that an idea is exceptional is p, their (ex-ante) expected return to exploring ideas is $(1-p)1+p3-c$, which we assume to be positive. By contrast, an Practice-based entrepreneur will assume that the expected value of any given draw is $(1-p)1+p3$ regardless of what draws have already come (so long as $0<p<1$). Thus, the Practice-based entrepreneur's (ex-ante) expected return is $(1-p)((1-p)1+p3-c)+p3-c$. This exceeds the Theory-driven entrepreneur's expected return because the Practice-based entrepreneur expects to keep searching following a draw of 1.

Notably, in the original example, the Theory-driven entrepreneur does not always value the idea less than the Practice-based one. Figure 6 highlights the region where the Theory-driven entrepreneur values the idea more than the Practice-based entrepreneur. This region exists specifically because of the ability of the Theory-driven entrepreneur to revert to a previous strategy: search is done with recall. For the pairs of beliefs and best-discovered strategies in this region, the Practice-based entrepreneur would stop, accepting that best-discovered strategy. But for the Theory-driven entrepreneur, they would take another draw, as the downside risk of being on the ordinary distribution and spending an extra cost of search is low relative to the upside potential of gaining an even better strategy if the distribution is, in fact, exceptional. The fact that the entrepreneur can recall ensures the worst-case outcome is limited for the Theory-driven entrepreneur.

In fact, for this specific example, Figure 6's region of increased value to the Theory-driven entrepreneur does not exist without recall. Without the ability to revert to a previous strategy, in this example, an Practice-based entrepreneur would value all ideas higher, regardless of prior or best-discovered strategy, than a Theory-driven entrepreneur.

These behaviors suggest we should observe Theory-driven entrepreneurs be less exploratory relative to Practice-based ones when prior best-discovered strategies are of relatively low value; Practice-based entrepreneurs in this scenario would appear more exploratory, engaging in more search. However, the reverse holds when prior best-discovered strategies are relatively high, as was the case in our McDonald's example when viewing the McDonald brothers as Practice-based entrepreneurs and Ray Kroc as Theory-driven.

5 Conclusion

Entrepreneurs differ from each other in numerous ways, from their demographic backgrounds to work history to personal networks (Evans and Leighton, 1990; Hurst and Lusardi, 2004; Gompers et al., 2005; Sørensen, 2007; Lerner and Malmendier, 2013). These differences are material in that they have been shown to influence entrepreneurial outcomes, affecting the success or failure of startups. Of particular interest to scholars focused on policy

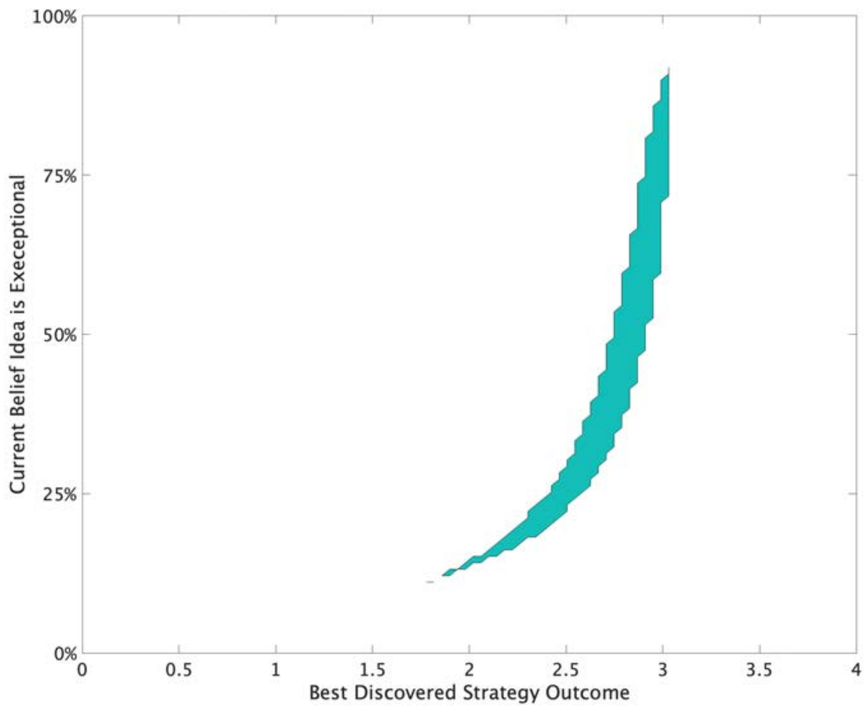

Figure 6: Region where Theory-driven entrepreneur values idea more than Practice-based entrepreneur

or pedagogy related to entrepreneurship are the characteristics that can change, as altering those characteristics, either through government policy or individual instruction, can improve entrepreneurial outcomes and the economy-wide benefits known to be associated with entrepreneurship, such as innovation and job growth (Gans et al., 2002; Decker et al., 2014). This study focuses on one such potentially mutable entrepreneurial characteristic: the mindset towards uncertainty held by the entrepreneur.

We show a Theory-driven approach towards uncertainty differs from an Practice-based approach in several ways. Theory-driven entrepreneurs are more likely to continue search after discovering a strategy with an abnormally large payoff; they interpret their success as indicative of the presence of even better strategies, as opposed to simple luck. They also may revert to a previous strategy if they learn they are wrong about the existence of even better strategies, in contrast to Practice-based entrepreneurs who always stop and execute the last discovered strategy. Theory-driven entrepreneurs have one additional reason to value the advice of outside experts such as venture capitalists: such experts can prove their worth by revealing an idea's true distribution of strategy outcomes. Finally, Theory-driven entrepreneurs, in some conditions, will value ideas less than Practice-based ones do since they are open to the possibility that the idea is actually worse than originally thought. Hence, we add a major consideration for potential entrepreneurs searching for the best strategy: is the industry being entered likely to benefit from theory, in other words, could one's initial beliefs potentially be wrong? If so, holding a theory-based approach is likely to be fruitful, necessitating a different set of behaviors than traditionally practised.

These findings contribute to the growing literature on entrepreneurial strategy: how the optimal decisions and behaviors for entrepreneurs differs from managers in established enterprises due to the nature of entrepreneurship. In our case, the uncertainty surrounding entrepreneurship in new industries is what drives our results and makes them unique to entrepreneurial settings. This builds on previous work studying the choices entrepreneurs need to make when entering markets with entrenched incumbents or when deciding between strategies that cannot be easily ranked (Gans et al., 2018, 2019). It also points to the potential in future work that characterises what is unique about entrepreneurial settings and explores the consequent implications for optimal entrepreneurial behavior.

6 Appendix: Proofs

6.1 Preliminary results

Lemmas 1 and 2 are required for Propositions 1 to 3. Lemma 1 proves search stops at a finite value. Lemma 2 adds that the value is in expectation also finite.

Lemma 1 *Under Assumptions 1 and 2, search stops at a finite value.*

Proof. The probability y of Equation (1) is less than some constant u and is therefore finite can be written as the intersection of all events where the draws of random variables were bounded by u plus the cost of search:

$$\Pr(y \leq u) = Pr(x_1 < u + c, x_2 < u + 2c, \ldots)$$

Using marginal distributions, this becomes:

$$\Pr(y \leq u) = \prod_{n=1}^{\infty} \Pr(x_n \leq u + nc | x_1 \leq u + c, \ldots, x_{n-1} < u + (n-1)c)$$

For each n we know based on Assumption 1 and the definition of first-order stochastic dominance (see for example Shaked and Shanthikumar (2007) (1.A.1)):

$$\Pr(\underline{x} \leq u + nc) \geq Pr(x_n \leq u + nc | x_1 \leq u + c, \ldots) \geq \Pr(\overline{x} \leq u + nc)$$

since at best the draws will be from $\overline{X}$ and at worst from $\underline{X}$. From Chow and Robbins (1963) Lemma 8, we know the stopping problem with independent identical distributions $\underline{X}$ and $\overline{X}$ having finite first moments (Assumption 2) has solutions $\underline{y} < \infty$ and $\overline{y} < \infty$ so can write:

$$\Pr\left(\underline{y} \leq u\right) = \prod_{n=1}^{\infty} \Pr(\underline{x}_n \leq u + nc) \geq Pr(y \leq u) \geq \prod_{n=1}^{\infty} \Pr(\overline{x}_n \leq u + nc) = Pr(\overline{y} \leq u) \tag{A1}$$

By the squeeze theorem this implies $\Pr(y \leq u) = 1$ since both $\Pr\left(\underline{y} \leq u\right) = 1$ and $\Pr(\overline{y} \leq u) = 1$. ∎

Lemma 2 *Under Assumptions 1, 2 and 3, the expected value of search is finite.*

Proof.

We can write:

$$E[y]=\int_0^{\infty}(1-\Pr(y<t))dt-\int_0^{\infty}\Pr(y<t)dt$$

From (A1) and Assumption 1 and the definition of first-order stochastic dominance, we know:

$$\int_0^{\infty}(1-\Pr(y<t))dt-\int_0^{\infty}\Pr(y<t)dt\le\int_0^{\infty}(1-\Pr(\overline{y}<t))dt-\int_0^{\infty}\Pr(\overline{y}<t)dt$$

where $\overline{y}$ is drawn from $\overline{X}$. Therefore, $E[\overline{y}]$ is an upper bound for $E[y]$. By Assumptions 2 and 3, we know $E[\overline{y}]$ has a finite upper bound, see Chow and Robbins (1963) Theorem 3. Similarly, $E[\underline{y}]$ where $\underline{y}$ is drawn from $\underline{X}$ is a lower bound for $E[y]$. So $E[y]$ must be finite. ■

Corollary 1 *Lemmas 1 and 2 hold when an outside option is present.*

Proof. Add a random variable X_0, which always has its outcomes mapped to the constant outside option z. Let stopping happen at stage $n=0$ if the outside option is taken without any draws. Since X_0 is independent of the other draws, we can write

$$\Pr(y\le u)=Pr\left[x_1\le u+c,x_2\le u+2c,\ldots\right]Pr[z\le u]$$

If $z>u$, we can always pick another finite $u'>z$ since z is finite, making $\Pr[z\le u]=1$ and allowing us to follow the rest of Lemma 1 proof that $y<\infty$. The same logic applies to Lemma 2 as well, using the above formulation of $\Pr(y\le u)$. ■

6.2 Proof of Proposition 1

We will show that the proposition follows from the concavity of Bayesian risk. We loosely conform to the notation used in Wijsman (1970) and assume the conditions for Lemma 2 hold so the value functions below exist.

Let each θ in a set $\mathcal{D}$ index the possible distributions an idea's strategy may have with the measure $\theta(dx)$ representing the probability a strategy has value x given the distribution is θ. Let the current best-discovered strategy have value z. Conditional the decision maker knowing a specific θ^* is true, the value of the idea would be:

$$v(z;\theta^*)=\begin{cases}z, & \text{if } z\ge\int v(\max\{z,x\};\theta^*)\theta^*(dx)-c\\ \int v(\max\{z,x\};\theta^*)\theta^*(dx)-c, & \text{otherwise}\end{cases}\tag{A2}$$

In contrast, suppose that while true distribution was θ^*, the decision maker only has beliefs about θ^*, allowed to evolve after each draw. Let Λ be a concave set of probability distributions over $\mathcal{D}$, as per Assumption 4. Each element λ of Λ representing the decision maker's beliefs about the likelihood each element of $\mathcal{D}$ is the correct strategy distribution for an idea, represented by the measure $\lambda(d\theta)$. After each search, the decision maker's beliefs update according to $\pi(\lambda, x)$, where x is the value of the discovered strategy.

In this case, a decision maker having prior λ would gain the following value from the idea:

$$v(z, \lambda; \theta^*) = \begin{cases} z, & \text{if } z \geq \int \int \hat{v}(\max\{z, x\}, \pi(\lambda, x))\theta(dx)\lambda(d\theta) - c \\ \int v(\max\{z, x\}; \theta^*)\theta^*(dx) - c, & \text{otherwise} \end{cases} \tag{A3}$$

Note Equation (A3) differs from Equation (A2) only in the decision rule for which payoff is selected at each stage, not the actual payoff itself. Equation (A3) uses the decision maker's current beliefs about the value of search to determine whether search stops or continues. As Equation (A2) optimizes this decision by selecting the maximum payoff, Equation (A3) must necessarily be weakly less than Equation (A2), as under Equation (A3) the decision maker may stop when continuing search would be optimal or continue search when stopping would be optimal.

We can write the value placed on an idea by the decision maker as the expected value of Equation (A3) given the probability distribution λ over the possible distributions $\mathcal{D}$:

$$v(z, \lambda) = \int v(z, \lambda; \theta)\lambda(d\theta) = E[v(z, \lambda; \theta)] \tag{A4}$$

Contrast this to the value of the idea if that decision-maker with prior belief λ was given knowledge of the true distribution. With the true distribution, the decision maker can follow the optimum decision rule of Equation (A2) rather than (A3):

$$v(z) = \int v(z; \theta)\lambda(d\theta) = E[v(z; \theta)] \tag{A5}$$

As the integrand of Equation (A4) is weakly less than the integrand of Equation (A5) while both Equation (A4) and Equation (A5) share the same measure, it follows that Equation (A4) must be weakly less than Equation (A5). Therefore, a decision maker must be weakly willing to pay for true knowledge of the strategy distribution for an idea, as this weakly increases the value of the idea for their from Equation (A4) to Equation (A5).

6.3 Lemma 3: Comparative statics under Assumptions 1 to 6

We now derive some comparative statics results used in the proofs of Propositions 2 and 3 below. Smith and McCardle (2002) provide in their Proposition 5 a framework to determine the comparative statics for dynamic programming problems such as the value function of Equation (1). Our setup can be described in their framework as:

- Each decision-making stage is index by k with $k = 0$ being the last stage
- Possible actions a_k at each stage k are either *stop* or *continue*
- A state vector x_k containing four elements:
  - The state s_k of the search at stage k which can either be *run* if search is ongoing or *halt* if search has terminated
  - The current best option z_k, initialized at 0
  - The current belief about the distribution λ_k
  - The cost of search c_k, which is constant in our setup
- A reward function $r_k(a_k, x_k)$ as follows

	$s_k = halt$	$s_k = run$
$a_k = stop$	0	z_k
$a_k = continue$	0	$-c_k$

- A transition function for the state vector $\tilde{x}_{k-1}(a_k, x_k)$ that accounts for a random element, in our case the discovered item y_k from search at stage k

x_k	a_k	$\tilde{x}_{k-1}$
$s_k = run$	*stop*	*halt*
	continue	*run*
$s_k = halt$	*stop*	*halt*
	continue	*halt*
z_k	*stop*	$\max\{z_k, y_k\}$
	continue	$\max\{z_k, y_k\}$
λ_k	*stop*	$\pi(\lambda_k, y_k)$
	continue	$\pi(\lambda_k, y_k)$
$-c_k$	*stop*	$-c_k$
	continue	$-c_k$

- *Increasing* as the closed convex cone property under consideration

Smith and McCardle's basic intuition is that each stage of a dynamic program can be broken down into a payoff or reward for that stage plus an expected value of continuation. As the combination of increasing functions is also increasing, the dynamic program will be increasing in its arguments as long as the reward and continuation value are both non-decreasing, conditional on the action taken at each stage. The reward function in our setup is weakly increasing in $-c_k$ and z_k, regardless of the action taken. λ_k has no effect on the reward function, so we can also refer to it as weakly increasing. Therefore, the first part of their requirements are met.

To show the continuation value is increasing, an increase in the parameters x_k must cause a FOSD shift in the distribution of $\tilde{x}_{k-1}$. First consider increasing z_k from z_k' to z_k''. The probability of getting a value of z_{k-1} below z_k' is zero for both cases. However, there is some mass between z_k' and z_k'' for the distribution conditional on z_k' that is missing for the distribution conditional on z_k''. For values of z_{k-1} above z_k'', both have the same cumulative density function. Therefore the difference in mass of z_{k-1} between z_k' and z_k'' implies $\tilde{z}_{k-1}(z_k'') \underset{FOSD}{>} \tilde{z}_{k-1}(z_k')$. As $\tilde{z}_{k-1}$ is the only future state affected by z_k, this is sufficient to show an increase in z_k produces a FOSD shift in $\tilde{x}_{k-1}$.

In addition, the distributions of λ_{k-1} and z_{k-1} must be FOSD in λ_k, since both are a function of λ_k. Increasing λ_k will cause a FOSD shift in y_k due to Assumption 5, following Theorem 1.A.6 in Shaked and Shanthikumar (2007). Since z_{k-1} is an increasing function of y_k, this will cause a corresponding FOSD shift in $\tilde{z}_{k-1}$ due to Theorem 1.A.3 in Shaked and Shanthikumar (2007).

Increasing λ_k has an effect on both the λ_k parameter in $\pi(\lambda_k, y_k)$ and the distribution of y_k, so both aspects need to be accounted for. For any $\lambda_k'' > \lambda_k'$, we know $\pi(\lambda_k'', y_k) > \pi(\lambda_k', y_k)$ for each y_k due to Assumption 6. Theorem 1.A.17 in Shaked and Shanthikumar (2007) gives us therefore that the distribution of $\pi(\lambda_k'', y_k)$ will be greater than $\pi(\lambda_k', y_k)$ conditional on y_k be distributed conditional on λ_k'. Theorem 1.A.6 of Shaked and Shanthikumar (2007) provides that under Assumption 5 the distribution of $\tilde{y}_k$ is FOSD in λ_k. As $\pi(\lambda_k'', y_k)$ is an increasing function of y_k under Assumption 6, we must therefore have the distribution of $\pi(\lambda_k'', y_k)$ conditional on $\tilde{y}_k$ being distributed conditional on λ_k'' FOSD dominating the distribution conditional on λ_k'. Therefore, the distribution of $\pi(\lambda_k, y_k)$ conditional on λ_k'' overall is FOSD over the distribution conditional on λ_k' and more generally the distribution of $\tilde{\lambda}_{k-1}$ is FOSD increasing in λ_k.

Finally, as $\tilde{c}_{k-1}$ is trivially distributed with a single point mass at c_k, increasing $-c_k$ also leads to a FOSD shift in the next period distribution.

As a last condition for v_k to be non-decreasing in z_k, λ_k and $-c_k$, we need v_k to exist as $k \to \infty$. This is given by Lemma 2, as v_∞ is essentially y of Equation (1).

In addition, when recall is not allowed, the transition function for z_k is simply y_k. Therefore, without recall, the same comparative statics as above trivially hold.

Moreover, when learning does not take place, the transition function for λ_k returns simply λ_k, so again, the same comparative statics hold.

6.4 Proof of Proposition 2

Following Assumption 7's terminology, suppose draw x'' is taken. Given the previous best option z and existing beliefs λ, define:

$$c^* \equiv \int v\left(\max\left\{z, x'', x\right\}, \pi\left(\lambda'', x\right)\right) f\left(x; \lambda''\right)\right) dx - \max\{z, x''\}$$

so at a cost, c^*, the decision maker is indifferent between stopping and continuing after drawing x'', making continuing search optimal, fulfilling the second part of the proposition.

Assumption 7 provides there exists a $x' < x''$ such that $\pi(\lambda, x') < \pi(\lambda, x'')$. Above we showed that the value function is decreasing in λ and therefore

$$c^* > \int v\left(\max\left\{z, x', x\right\}, \pi\left(\lambda', x\right)\right) f\left(x; \lambda'\right) dx - \max\{z, x'\}.$$

This implies stopping occurs at x', fulfilling the first part of the proposition.

6.5 Proof of Proposition 3

Let $z = x''$ and similarly to Proposition 2 define the cost as the indifferent point between continuing and stopping search after drawing x'' when z already equals x'':

$$c^* \equiv \int v\left(\max\left\{x'', x\right\}, \pi\left(\lambda'', x\right)\right) f\left(x; \lambda''\right)\right) dx - x''.$$

With this cost of search, if x' is instead drawn when $z = x''$, stopping occurs due to Assumption 7, with the searcher reverting to the strategy having outcome x'' which is less than x'.

$$c^* > \int v\left(\max\left\{x'', x\right\}, \pi\left(\lambda', x\right)\right) f\left(x; \lambda'\right) dx - x''.$$

References

Adam, K. (2001). Learning while searching for the best alternative. *Journal of Economic Theory*, 101(1):252–280.

Agrawal, A., Gans, J. S., and Stern, S. (2021). Enabling entrepreneurial choice. *Management Science*, 67(9):5510–5524.

Alvarez, S. A. and Barney, J. B. (2007). Discovery and creation: Alternative theories of entrepreneurial action. *Strategic Entrepreneurship Journal*, 1(1-2):11–26.

Angus, R. W. (2019). Problemistic search distance and entrepreneurial performance. *Strategic Management Journal*, 40(12):2011–2023.

Arentz, J., Sautet, F., and Storr, V. (2013). Prior-knowledge and opportunity identification. *Small Business Economics*, 41:461–478.

Arikan, A. M., Arikan, I., and Koparan, I. (2020). Creation opportunities: Entrepreneurial curiosity, generative cognition, and knightian uncertainty. *Academy of Management Review*, 45(4):808–824.

Aulet, W. and Murray, F. (2013). A tale of two entrepreneurs: Understanding differences in the types of entrepreneurship in the economy. *Available at SSRN 2259740*.

Baron, R. A. and Ensley, M. D. (2006). Opportunity recognition as the detection of meaningful patterns: Evidence from comparisons of novice and experienced entrepreneurs. *Management Science*, 52(9):1331–1344.

Bernanke, B. S. (1983). Irreversibility, uncertainty, and cyclical investment. *The Quarterly Journal of Economics*, 98(1):85–106.

Bernstein, S., Giroud, X., and Townsend, R. R. (2016). The impact of venture capital monitoring. *The Journal of Finance*, 71(4):1591–1622.

Bewley, T. F. (2002). Knightian decision theory. part i. *Decisions in Economics and Finance*, 25(2):79–110.

Bezos, J. (1997). Jeff bezos special libraries (sla) interview. Technical report, Special Libraries (SLA).

Bhardwaj, G., Camillus, J. C., and Hounshell, D. A. (2006). Continual corporate entrepreneurial search for long-term growth. *Management Science*, 52(2):248–261.

Bhide, A. (2003). *The origin and evolution of new businesses*. Oxford University Press.

Botelho, T. L., Fehder, D., and Hochberg, Y. (2021). Innovation-driven entrepreneurship. Technical report, National Bureau of Economic Research.

Buckman, R. and Gasparino, C. (1999). Merrill lynch agrees to buy d.e. shaw online-trading unit. *Wall Street Journal*. Febuary, 19.

Camerer, C. and Weber, M. (1992). Recent developments in modeling preferences: Uncertainty and ambiguity. *Journal of Risk and Uncertainty*, 5:325–370.

Camuffo, A., Cordova, A., Gambardella, A., and Spina, C. (2020). A scientific approach to entrepreneurial decision making: Evidence from a randomized control trial. *Management Science*, 66(2):564–586.

Camuffo, A., Gambardella, A., and Pignataro, A. (2023). Theory-driven strategic management decisions. *CEPR DP 17664v2*.

Cerreia-Vioglio, S., Hansen, L. P., Maccheroni, F., and Marinacci, M. (2020). Making decisions under model misspecification. *arXiv preprint arXiv:2008.01071*.

Chen, J. S., Croson, D. C., Elfenbein, D. W., and Posen, H. E. (2018). The impact of learning and overconfidence on entrepreneurial entry and exit. *Organization Science*, 29(6):989–1009.

Choi, Y. R., Lévesque, M., and Shepherd, D. A. (2008). When should entrepreneurs expedite or delay opportunity exploitation? *Journal of Business Venturing*, 23(3):333–355.

Chou, C.-f. and Talmain, G. (1993). Nonparametric search. *Journal of Economic Dynamics and Control*, 17(5-6):771–784.

Chow, Y. S. and Robbins, H. (1963). On optimal stopping rules. *Zeitschrift für Wahrscheinlichkeitstheorie und verwandte Gebiete*, 2(1):33–49.

Decker, R., Haltiwanger, J., Jarmin, R., and Miranda, J. (2014). The role of entrepreneurship in us job creation and economic dynamism. *Journal of Economic Perspectives*, 28(3):3–24.

Delmar, F. and Shane, S. (2003). Does business planning facilitate the development of new ventures? *Strategic Management Journal*, 24(12):1165–1185.

Ehrig, T. and Schmidt, J. (2022). Theory-based learning and experimentation: How strategists can systematically generate knowledge at the edge between the known and the unknown. *Strategic Management Journal*, 43(7):1287–1318.

Evans, D. S. and Leighton, L. S. (1990). Some empirical aspects of entrepreneurship. In *The economics of small firms*, pages 79–99. Springer.

Felin, T., Gambardella, A., Novelli, E., and Zenger, T. (2023). A scientific method for startups. *Journal of Management*.

Felin, T. and Zenger, T. R. (2009). Entrepreneurs as theorists: on the origins of collective beliefs and novel strategies. *Strategic Entrepreneurship Journal*, 3(2):127–146.

Felin, T. and Zenger, T. R. (2017). The theory-based view: Economic actors as theorists. *Strategy Science*, 2(4):258–271.

Fiet, J. O. (2007). A prescriptive analysis of search and discovery. *Journal of Management Studies*, 44(4):592–611.

Fiet, J. O. and Patel, P. C. (2008). Entrepreneurial discovery as constrained, sytematic search. *Small Business Economics*, 30:215–229.

Fiet, J. O., Piskounov, A., and Patel, P. C. (2005). Still searching (systematically) 1 for entrepreneurial discoveries. *Small Business Economics*, 25(5):489–504.

Frederiks, A. J., Englis, B. G., Ehrenhard, M. L., and Groen, A. J. (2019). Entrepreneurial cognition and the quality of new venture ideas: An experimental approach to comparing future-oriented cognitive processes. *Journal of Business Venturing*, 34(2):327–347.

Gaglio, C. M. and Katz, J. A. (2001). The psychological basis of opportunity identification: Entrepreneurial alertness. *Small Business Economics*, 16:95–111.

Gambardella, A., Camuffo, A., Maccheroni, F., Marinacci, M., and Pignataro, A. (2022). Microfoundations of low-frequency high-impact decisions. CEPR Discussion Papers 17392, C.E.P.R. Discussion Papers.

Gans, J. S. (2023). Experimental choice and disruptive technologies. *Management Science*.

Gans, J. S., Hsu, D. H., and Stern, S. (2002). When does start-up innovation spur the gale of creative destruction? *RAND Journal of Economics*, 33(4):571–586.

Gans, J. S., Scott, E. L., and Stern, S. (2018). Strategy for start-ups. *Harvard Business Review*, 96(3):44–51.

Gans, J. S., Stern, S., and Wu, J. (2019). Foundations of entrepreneurial strategy. *Strategic Management Journal*, 40(5):736–756.

Goldstein, D. G., McAfee, R. P., Suri, S., and Wright, J. R. (2020). Learning when to stop searching. *Management Science*, 66(3):1375–1394.

Gomes, L. (2000). Novell edges street estimates, but its revenue slides 17%. *Wall Street Journal*. August 17.

Gompers, P., Lerner, J., and Scharfstein, D. (2005). Entrepreneurial spawning: Public corporations and the genesis of new ventures, 1986 to 1999. *The Journal of Finance*, 60(2):577–614.

Gruber, M., MacMillan, I. C., and Thompson, J. D. (2008). Look before you leap: Market opportunity identification in emerging technology firms. *Management Science*, 54(9):1652–1665.

Guzman, J. and Stern, S. (2020). The state of american entrepreneurship: New estimates of the quantity and quality of entrepreneurship for 32 us states, 1988–2014. *American Economic Journal: Economic Policy*, 12(4):212–43.

Hsieh, T. (2010). *Delivering Happiness: A Path to Profits, Passion, and Purpose*. Hachette UK.

Hurst, E. and Lusardi, A. (2004). Liquidity constraints, household wealth, and entrepreneurship. *Journal of Political Economy*, 112(2):319–347.

Johnson, W. R. (1978). A theory of job shopping. *The Quarterly Journal of Economics*, 92(2):261–277.

Jovanovic, B. and Rob, R. (1990). Long waves and short waves: Growth through intensive and extensive search. *Econometrica: Journal of the Econometric Society*, pages 1391–1409.

Kaplan, S. N. and Strömberg, P. (2001). Venture capitalists as principals: Contracting, screening, and monitoring. *American Economic Review*, 91(2):426–430.

Karni, E. and Vierø, M.-L. (2017). Awareness of unawareness: A theory of decision making in the face of ignorance. *Journal of Economic Theory*, 168:301–328.

Keller, G., Rady, S., and Cripps, M. (2005). Strategic experimentation with exponential bandits. *Econometrica*, 73(1):39–68.

Kirzner, I. M. (1997). Entrepreneurial discovery and the competitive market process: An austrian approach. *Journal of Economic Literature*, 35(1):60–85.

Knight, F. H. (1921). *Risk, uncertainty and profit*, volume 31. Houghton Mifflin.

Kroc, R. and Anderson, R. (1987). *Grinding it out: The making of McDonald's*. Macmillan.

Lerner, J. and Malmendier, U. (2013). With a little help from my (random) friends: Success and failure in post-business school entrepreneurship. *The Review of Financial Studies*, 26(10):2411–2452.

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

Levy, S. (2011). *In the plex: How Google thinks, works, and shapes our lives*. Simon & Schuster.

McCall, J. J. (1965). The economics of information and optimal stopping rules. *The Journal of Business*, 38(3):300–317.

Morgan, P. B. (1983). Search and optimal sample sizes. *The Review of Economic Studies*, 50(4):659–675.

Mosakowski, E. (1997). Strategy making under causal ambiguity: Conceptual issues and empirical evidence. *Organization Science*, 8(4):414–442.

Nishimura, K. G. and Ozaki, H. (2004). Search and knightian uncertainty. *Journal of Economic Theory*, 119(2):299–333.

Norton, W. I. and Moore, W. T. (2002). Entrepreneurial risk: have we been asking the wrong question? *Small Business Economics*, 18:281–287.

Rapp, D. J. and Olbrich, M. (2023). From knightian uncertainty to real-structuredness: Further opening the judgment black box. *Strategic Entrepreneurship Journal*, 17(1):186–209.

Rindova, V. and Martins, L. (2023). Ways of knowing: Four types of theories as guides for strategy making. In *Theory-Based View of Strategy*, Milan, Italy. Bocconi University.

Rivkin, J. W. (2000). Imitation of complex strategies. *Management Science*, 46(6):824–844.

Rosenfield, D. B. and Shapiro, R. D. (1981). Optimal adaptive price search. *Journal of Economic Theory*, 25(1):1–20.

Rothschild, M. (1974). A two-armed bandit theory of market pricing. *Journal of Economic Theory*, 9(2):185–202.

Sarasvathy, S. D. (2001). Effectual reasoning in entrepreneurial decision making: existence and bounds. In *Academy of Management Proceedings*, volume 2001, pages D1–D6. Academy of Management Briarcliff Manor, NY 10510. No. 1.

Schröder, D. (2011). Investment under ambiguity with the best and worst in mind. *Mathematics and Financial Economics*, 4:107–133.

Shaked, M. and Shanthikumar, J. G. (2007). *Stochastic orders*. Springer Science & Business Media.

Shane, S. and Delmar, F. (2004). Planning for the market: business planning before marketing and the continuation of organizing efforts. *Journal of Business Venturing*, 19(6):767–785.

Shane, S. and Venkataraman, S. (2000). The promise of entrepreneurship as a field of research. *Academy of Management Review*, 25(1):217–226.

Shane, S. A. (2003). *A general theory of entrepreneurship: The individual-opportunity nexus*. Edward Elgar Publishing.

Smith, J. E. and McCardle, K. F. (2002). Structural properties of stochastic dynamic programs. *Operations Research*, 50(5):796–809.

Smith, K. G. and Cao, Q. (2007). An entrepreneurial perspective on the firm-environment relationship. *Strategic Entrepreneurship Journal*, 1(3-4):329–344.

Sørensen, J. B. (2007). Bureaucracy and entrepreneurship: Workplace effects on entrepreneurial entry. *Administrative Science Quarterly*, 52(3):387–412.

Stone, B. (2013). *The Everything Store: Jeff Bezos and the Age of Amazon*. Random House.

Talmain, G. (1992). Search from an unkown distribution an explicit solution. *Journal of Economic Theory*, 57(1):141–157.

Tonks, I. (1983). Bayesian learning and the optimal investment decision of the firm. *The Economic Journal*, 93(Supplement):87–98.

Ward, T. B. (2004). Cognition, creativity, and entrepreneurship. *Journal of Business Venturing*, 19(2):173–188.

Weitzman, M. (1978). *Optimal search for the best alternative*, volume 78. Department of Energy. No. 8.

Wijsman, R. A. (1970). Continuity of the bayes risk. *The Annals of Mathematical Statistics*, 41(3):1083–1085.

Wuebker, R., Zenger, T., and Felin, T. (2023). The theory-based view: Entrepreneurial microfoundations, resources, and choices. *Strategic Management Journal*, 44(12):2922–2949.