NBER WORKING PAPER SERIES

ENABLING ENTREPRENEURIAL CHOICE

Ajay K. Agrawal
Joshua S. Gans
Scott Stern

Working Paper 27379
http://www.nber.org/papers/w27379

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
June 2020

We thank Shannon Liu and Amir Sariri for helpful comments. 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 a financial relationship of potential relevance for this research. Further information is available online at http://www.nber.org/papers/w27379.ack

Enabling Entrepreneurial Choice
Ajay K. Agrawal, Joshua S. Gans, and Scott Stern
NBER Working Paper No. 27379
June 2020
JEL No. D83,O3

ABSTRACT

Entrepreneurs must choose between alternative strategies for bringing their idea to market. They face uncertainty regarding both the quality of their idea as well as the efficacy of each strategy. While entrepreneurs can reduce this uncertainty by conducting tests, any single test conflates the signal of the efficacy of the particular strategy and the quality of the idea. Resolving this conflation requires exploring multiple strategies. Consequently, entrepreneurial choice is enhanced by finding ways to lower the cost of testing multiple strategies, receiving guidance as to the types of tests likely to reduce signal conflation, and optimally sequencing tests based on prior beliefs. This creates a role for judgment that may be provided by third parties such as mentors and investors. We hypothesize that institutions that lower the cost of transmitting and aggregating judgment spur entrepreneurial success.

Ajay K. Agrawal
Rotman School of Management
University of Toronto
105 St. George Street
Toronto, ON M5S 3E6
CANADA
and NBER
ajay.agrawal@rotman.utoronto.ca

Joshua S. Gans
Rotman School of Management
University of Toronto
105 St. George Street
Toronto ON M5S 3E6
CANADA
and NBER
joshua.gans@gmail.com

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

1 Introduction

There are often multiple potential strategies for bringing a single idea to market. For instance, during the late 1990s, two start-ups, Webvan and Peapod, entered the market for online grocery shopping but did so with very different commercialization strategies. Whereas Peapod pursued the idea of online grocery shopping through cooperative partnerships with existing physical supermarkets, Webvan pursued an integrated solution in which they would compete directly with existing supermarkets. This simple example highlights two related but distinct sources of uncertainty facing entrepreneurs: whether their entrepreneurial idea is of high or low quality (i.e., whether online grocery shopping is a valuable concept) and whether the strategies they are considering are better or worse paths for bringing that idea to the marketplace (i.e., whether this idea is likely to be successful through cooperation, competition, or either approach).

Since the ultimate success of a venture depends both on the quality of the idea as well as the chosen strategy, reducing uncertainty about both of these dimensions is critical. Reducing uncertainty is particularly important for start-up ventures since outright failure is the modal outcome (more than 80% of all new ventures fail within one year of founding (Shane, 2008)). It is, therefore, not surprising that a central question for entrepreneurs is how to design and interpret tests that allow them to assess both the quality of their ideas and the value of particular strategies (Murray and Tripsas, 2004; Ries, 2011; Kerr et al., 2014).

Though conceptually straightforward, the impact of experimentation on entrepreneurial decision-making is subtle. On the one hand, from the perspective of a canonical model of entrepreneurial investment where the only source of uncertainty is whether the venture is pursuing a good 'opportunity,' early experiments that yield positive information induce the entrepreneur to pursue the next stage of development, while negative results prompt the termination of the venture (Nanda and Rhodes-Kropf, 2016a, 2016b). However, if entrepreneurs are simultaneously learning about both the quality of their idea as well as the strength of their strategy, then 'negative' news might result, not in termination, but a 'pivot' away from the current strategy towards an alternative strategic approach (Ries, 2011). This type of adaptive learning, even in the face of negative information, is consistent with the idea that interim periods of failure are common even among those ventures that ultimately succeed (Azoulay and Shane, 2001).

The simultaneous existence of two very different responses to negative information – discontinuation and persistence – highlight the subtle interplay between entrepreneurial experimentation, the nature of learning and entrepreneurial choice. While most existing models of entrepreneurial investment highlight the option value of experiments (where negative information results in discontinuation), the potential for multiple strategies allows for the possibility that even significant periods of negative information need not discourage additional search. In the extreme, if an entrepreneur is completely confident in the underlying value of their idea, then they have a significant incentive to continue to test strategies for that idea until they realize a stage such that the expected return to additional search is less than the costs of additional search. As famously quipped by Thomas Edison, "I have not failed. I've just found 1000 ways that do not work." (Elkorne, 1967)

Of course, it is possible that entrepreneurs could undertake experimentation in such a way as to first resolve uncertainty about the idea and then, conditional on establishing the value of the idea, explore strategies to maximize the return from that idea. In such a case, the core decision from early tests would be whether to abandon the venture or commit to a search for a strategy. If one designs a test that exclusively provides information about the value of a given strategy (conditional on the value of the idea), then negative information simply provides a prompt for further search rather than abandonment of the venture itself.

However, most entrepreneurial experimentation and learning involves a *combined* test of both the idea and strategy. To make this idea concrete, consider the choices facing Walt Disney in 1928 (Gabler, 2006). Disney's core idea was to create animated characters in full-fledged stories that would be appropriate for children (the leading animated characters of the day, such as Felix the Cat or Betty Boop were featured in episodic shorts that arguably were more aimed at adults than children). Disney's chosen strategy for implementing this idea was to build a team at his studio in California and create shorts featuring Oswald the Lucky Rabbit, which were then distributed by a leading industry player, Charles Mintz. In March 1928, Disney met Mintz in New York to renegotiate his contract, at which point Mintz informed him that he had hired away several key animators to found a new studio and that the initial contract Disney had signed had vested Mintz rather than Disney with the copyright and control over the Oswald character. In essence, the entrepreneurial experiment of the "Oswald" filmed resulted in an outcome that placed Disney's fledging venture in danger of failing or perhaps being acquired at a low price by Mintz himself. The choice facing Disney at that moment depended on his assessment as to whether this meaningful failure was the result of a bad idea or its particular implementation. His inference that failure was due to bad strategy prompted Disney to create a new character (Mickey Mouse looks quite like Oswald just with rounder ears) and pioneer sound cartoons to great acclaim in "Steamboat Willie" (over which Disney made sure to retain the copyright). The ultimate success of Disney was grounded in the combination of a good underlying idea and his eventual ability to identify a good strategy for that idea (Disney often noted that "it all started with a mouse.").

The primary goal of this paper is to provide a simple but general model of entrepreneurial experimentation and choice when entrepreneurial experimentation inherently involves learning simultaneously about the underlying idea as well as particular strategies. Specifically, we build a model where an entrepreneur is uncertain about the quality of an idea but can reduce the degree of uncertainty by undertaking entrepreneurial experiments. However, these experiments also involve some degree of implementation that conflates the nature of the learning that is possible. For instance, if the entrepreneur were to launch a new service to a particular customer segment, then if that app receives rave reviews, then (at least for that customer segment), there is unambiguous feedback that both the app idea and the targeted customer segment is good. As well, if the reviews are uniformly negative (i.e., the customers did not like the app), the entrepreneur can efficiently learn that this particular concept is poor. But intermediate feedback likely conflates whether the concept is strong (but the app idea has value) or the customer segment was the only people who find value from the concept.

This simple model yields a number of important insights. First, and perhaps most obviously, if possible, it is optimal to select experiments that reduce uncertainty about the idea prior to undertaking experiments that reduce uncertainty about a particular strategy. Second, even when each individual experiment conflates the signalled value of idea and strategy, multiple sequential tests allow an entrepreneur to gain information about the value of their idea. This is because the values realized across different experiments are correlated insofar as they share a common idea (which contributes to the signal they receive at the end of each experiment). Putting these ideas together, it is possible that an entrepreneur will ultimately implement a test that has a low probability of success in order to provide critical feedback about whether they have a good or bad idea (and so help them choose whether to continue to search or abandon the venture).

The paper brings together two distinct and productive lines of research that have, nonetheless, mostly developed independently of each other. On the one hand, there is a significant literature in entrepreneurial finance that considers how the ability to experiment shapes the staged nature of entrepreneurial financing (see Admati and Pfleiderer, 1994, for the initial treatment). Most notably, building on a line of research in real options, a recent and important line of work of Nanda and Rhodes-Kropf (2016b) illustrates how the ability to experiment allows early-stage projects to proceed (even if their NPV is negative) since negative information received during the experimental phase allows the project to be shut down without further investment (Ewens, Nanda and Rhodes-Kropf (2018) provide support for the predictions of this model by examining the impact of low-cost cloud computing on entrepreneurial financing in software versus other fields).

On the other hand, there is an extensive literature (covering several fields) about how to search and choose among strategies given that one strategy will be implemented. For example, a long line of research considers the nature of optimal stopping rules in search problems (e.g., see the canonical treatment of Weitzman (1979) and the application to entrepreneurship by Bergemann and Hege, 2005). This literature has both been inspired by (Kirzner, 1997) and inspired research in entrepreneurial search (Fiet and Patel, 2008; Gans, Stern and Wu, 2019). Manso (2011) examines incentives to experiment in the search for innovation and how this shapes optimal contracts within firms. However, there has been, perhaps surprisingly, relatively little attention to the nature of entrepreneurial choice when entrepreneurial experimentation involves learning simultaneously about the value of the underlying idea and the value of particular strategies. Creating a tractable model that incorporates both of these sources of uncertainty offers novel insight into both the optimal experimentation for a single entrepreneur as well as the impact of mentors and educational programs meant to spur improved entrepreneurial experimentation and decision-making.

2 Baseline Model

We focus on the choices facing an individual entrepreneur. It is presumed that the entrepreneur has a single 'idea.' An idea is an opportunity to exploit a new technology or combination of technologies in some market. To commercialize an idea (that is, bring it to a market), the entrepreneur needs to formulate, and then undertake, a strategy culminating in a venture's launch. Our primary modelling assumption is that it is easier for the entrepreneur to change their strategy than their idea; something we model in the extreme, presuming the idea cannot be changed at all. In other words, entrepreneurial opportunities are scarce; strategies to commercialize them are not (Erkal and Scotchmer, 2009). While we assume here that the venture has a single idea and many commercialization options, it is useful to note that the main results of the discussion to follow would be unchanged if there were multiple ideas each with multiple commercialization options that offer a correlated signal of the underlying value of particular ideas.

The entrepreneur's main choice is what strategies to undertake to commercialize their idea. In making this choice, an entrepreneur faces two main sources of uncertainty. First, how 'valuable' is the core idea of their venture? Second, how 'effective' (in terms of creating and capturing value) are particular strategies associated with commercializing that idea? The nature of that uncertainty and, in particular, how those two sources interact plays an important role in how entrepreneurs make choices of whether to implement particular strategies, whether to continue with their venture and how to test for outcomes before making key commitments or expending significant resources.

To illustrate this, consider the following environment. Suppose that an idea has two possible values, v and 0 (generated by states, h and l respectively) which can be realized by implementing an effective (g or 'good') strategy. If, instead, an ineffective (b or 'bad') strategy is used, the value realized is 0. The entrepreneur's prior probability that the idea is of high value is p_0. Strategies are drawn from a set, S, of possibilities with element, s_i ($i \in \{1, \dots, n\}$) where $n > 1$. For any given strategy, the prior probability that it is effective is e_0 and effectiveness is independent across strategies and of the value of the idea. To fully launch and implement a particular strategy costs C. Thus, if the entrepreneur were to launch based on a randomly selected strategy, their expected return would be:

$$V_0 = e_0 p_0 v - C$$

To keep the exposition simple, we will focus on the interesting case where $V_0 = 0$; that is, that choosing to randomly pursue a strategy is feasible but not especially profitable.[2] This means that the entrepreneur, if possible, would prefer to have more information before implementing any particular strategy but will pursue the venture if this is not possible.

The entrepreneur can gather information prior to the implementation of a strategy by conducting a 'test' of a given commercialization option.[3] A 'test' generates a signal of the value

---

[2] If $V_0 > 0$, then the entrepreneur will choose to pursue the venture with a randomly selected strategy and no further information while if $V_0 < 0$, the opposite is true. However, as Nanda and Rhodes-Kropf (2016b) show, even with these assumptions, so long as the costs of learning are not too high, the entrepreneur will choose to gather information prior to choosing whether to continue with the venture or not. Thus, an assumption of $V_0 = 0$ does not alter that basic property while simplifying exposition below.

[3] Sometimes this is referred to as an experiment (Nanda and Rhodes-Kropf, 2016b) or as exploration (Manso, 2011). We use the 'test' terminology to distinguish our notion that a commercialization option paired with an idea is being

of an idea and a particular strategy that was used as part of the test. The signal of the idea can be H (for high) or L (for low), and of the strategy it can be G (for good) or B (for bad). The cost of conducting a test of a particular strategy is c ($< C$). Critically, $\Pr[h|H]$, $\Pr[l|L]$, $\Pr[g|G]$ and $\Pr[b|B]$ are all strictly less than 1. The benefit of testing rather than implementing is that the entrepreneur can decide to switch to another strategy or abandon the venture prior to implementation should the test signals deliver 'bad' news. The constraint on testing is that it is difficult to pursue many tests given the resource constraints often faced by entrepreneurial firms. Thus, each test will be followed by a decision point at which the venture might be abandoned, continue to implement the tested strategy, implement an alternative strategy or potentially continue with further tests.

It is useful to note that this environment captures the informal axioms for entrepreneurial strategy of Gans, Stern and Wu (2019). They argue that what makes strategic choices by an entrepreneurial venture distinctive relative to an established firm is that: (1) there is more than one path to create and capture value from an idea; (2) constraints prevent the pursuit of more than one alternative at once; (3) the parameters of the probability distribution governing the value of an idea are not known by the entrepreneur and (4) that commitment-free learning can only generate noisy estimates of the value of an idea and a given strategic alternative, and the relationship between the two. (1) is captured by our assumption that $n > 1$. (2) is captured by assuming that only one strategy can be tested or implemented at a given time. (3) is captured by the uncertainty over the idea's value. These three assumptions generate potential value for testing prior to implementation (Nanda and Rhodes-Kropf, 2016b).

Axiom (4) is not captured by the model thus far as we have not described the signal space following a test. We will do that formally in the next section. However, it is useful to reflect here what would happen if a test could reveal a signal of the value of an idea as distinct from any one strategy. This might occur, for instance, if the entrepreneur had a high degree of confidence in a particular strategy (such that $e_0 = 1$). In this case, the only reason that testing that strategy would fail is because the idea was a poor one. Of course, by definition, if the entrepreneur is confident in the performance of more commercialization options, then entrepreneurial strategy itself – as a process of choosing the best strategy – would have no bite. It is where both $p_0, e_0 \ll 1$ that the

tested as opposed to information gathering that is part of a commercialization option (such as A-B testing or iterating to improve technological performance).

choice of strategy matters but it is precisely in this situation that it is difficult to devise tests that would give rise to signals of an idea's value or a strategy's effectiveness that were not, in some sense, conflated by one another. One reason for this, as emphasized by Gans, Stern and Wu (2019) is that a proper test of a strategy's effectiveness often involves commitments to implement the strategic options – say, to foster network effects in building a platform – and thus, the consequent likelihood that uncertainty over the strategy's effectiveness will be resolved. Without such commitments, the strategy's effectiveness is very uncertain, which means that the possibility of separating out an idea's value from a test based on a strategy is not possible. It is this that makes the choice of strategy – both for testing and subsequent implementation – a challenging one for entrepreneurial ventures who are resource-constrained.

In summary, uncertainty means that the challenge is that it is not possible to explore the value of the idea separate from any given strategy – that is, an idea cannot be tested without testing an associated strategy. The good news is that because there are many strategies associated with a given idea, with enough experimentation, the signals of each can be extracted. The bad news is that it is unlikely that enough experimentation can actually be conducted without carefully selecting what to test. It is this that drives the entrepreneurial choice process, and the nature of assistance entrepreneurs need to improve that process.

3 Conflated Signals of Idea and Strategy

We now turn to see how uncertainty changes the nature of the entrepreneurial choice problem. What makes testing valuable is that it is only possible to implement a single strategy. If it was possible to launch multiple strategies, then those could be regarded as perfectly informative experiments. As they are costlier, there may still be value to cheaper tests, but the trade-off would then be in terms of the costs and benefits of more information. When only one strategy can be fully launched, there is an option value to testing in that it allows a strategy to be discarded and another one undertaken in its place.

3.1 Signal space

Our contention here is that uncertainty associated with most entrepreneurial environments makes it difficult to obtain signals that distinguish between the quality of the idea and effectiveness of the strategy. In particular, while some signals are good news on both dimensions –

unambiguously positive outcomes suggest a valuable idea and effective strategy while unambiguously negative outcomes suggest a poor idea and an ineffective strategy – there are intermediate signals that are conflated. Specifically, under a test, if you pick a strategy, that generates both a signal of effectiveness (*G* or *B*) and a signal of value (*H* or *L*). The signals are correlated. In particular, they exhibit the following characteristics:

		Idea Value: *h*	Idea Value: *l*
Strategy Effectiveness	*g*	(*H*,*G*) wp 1	(*L*,*G*) wp λ_s (*H*,*B*) wp (1- λ_s)
	b	(*H*,*B*) wp λ_v (*L*,*G*) wp (1 - λ_v)	(*L*,*B*) wp 1

For $\lambda_v < 1$, then, if a strategy is bad, this may mask the signal that an idea's value is high, while, for $\lambda_s < 1$, if the value of the idea is low, a strategy may present itself as bad when it is, in fact, good. Thus, as $\lambda_v \to 1$, a high-value idea is more clearly signalled while as $\lambda_s \to 1$, a good strategy is more clearly signalled. In that regard, if experiments can be improved or selected, they can be directed to have a clearer signal of idea value (a larger λ_v) or strategy effectiveness (a larger λ_s).

If a 'test' of a particular strategy, s_i, is conducted then, using Bayes Law, the priors on the idea value and on that particular strategy are updated as follows:

$$p_1(H,B) = \frac{p_0(1-e_0)\lambda_v}{(1-p_0)e_0(1-\lambda_s)+p_0(1-e_0)\lambda_v} \text{ and } p_1(L,G) = \frac{p_0(1-e_0)(1-\lambda_v)}{(1-p_0)e_0\lambda_s+p_0(1-e_0)(1-\lambda_v)}$$

$$e_1(H,B) = \frac{(1-p_0)e_0(1-\lambda_s)}{(1-p_0)e_0(1-\lambda_s)+p_0(1-e_0)\lambda_v} \text{ and } e_1(L,G) = \frac{(1-p_0)e_0\lambda_s}{(1-p_0)e_0\lambda_s+p_0(1-e_0)(1-\lambda_v)}$$

To give meaning to the signals being provided, we want a signal (such as (*H*, *B*)) to cause the posterior probability of idea quality to rise above its prior and that of strategy effectiveness to fall below its prior. In other words, we want the signal to convey some information regarding the underlying state. Thus, we will assume that the 'test' is weakly informative. That is,

(A1) *Tests are (weakly) informative*. $p_1(H,B) \geq p_0$ (or $\frac{\lambda_v}{1-\lambda_s} \geq \frac{e_0}{1-e_0}$), $p_1(L,G) \leq p_0$ (or $\frac{1-\lambda_v}{\lambda_s} < \frac{e_0}{1-e_0}$), $e_1(H,B) \leq e_0$ (or $\frac{1-\lambda_s}{\lambda_v} \leq \frac{p_0}{1-p_0}$) and $e_1(L,G) \geq e_0$ (or $\frac{\lambda_s}{1-\lambda_v} \geq \frac{p_0}{1-p_0}$).

If tests are (weakly) informative, this constrains the degree to which the signal may not represent the underlying state. Specifically, (A1) holding implies that:

$$\frac{\lambda_s}{1-\lambda_v} \geq \frac{p_0}{1-p_0}, \frac{1-e_0}{e_0} \geq \frac{1-\lambda_s}{\lambda_v}$$

This places lower bounds on (λ_v, λ_s); the degrees to which the signal (L, G) might arise even if the idea has value (h), given by λ_v, and to which the signal (H, B) might arise even if the strategy tested is effective (g).

3.2 Timeline

Here we introduce a simplified model to understand the choices involved in testing and implementing strategic options. While conceptually, testing could be a process involving many iterations, here we focus on an environment where there are just two periods; 0 and 1. In the first period, the entrepreneur selects a strategy at random and can either test it or implement it. If it is implemented, no further strategies can be tested or implemented in the second period. If the strategy is tested, then in the second period, it can be implemented, or another strategy can be selected. Thus, this timeline only allows for, at most, a single round of testing.[4] For simplicity, we assume there is no discounting between the two periods.

3.3 Impact of a test

The payoff for testing a strategy in period 0 is:

$$V_0(\lambda_v, \lambda_s) = p_0 e_0 (v - C) + p_0(1 - e_0)\big(\lambda_v V_1(H, B) + (1 - \lambda_v) V_1(L, G)\big) \\ + (1 - p_0) e_0 \big((1 - \lambda_s) V_1(H, B) + \lambda_s V_1(L, G)\big) - c$$

There are four potential outcomes corresponding to each of the signals that might be generated:

1. For a signal of (H, G) which the entrepreneur expects to occur with probability, $p_0 e_0$, the entrepreneur learns that, with probability 1, both the idea is of high value and the strategy is effective. In this case, the entrepreneur optimally chooses to implement that strategy, earning $v - C$.
2. For a signal of (L, B) which the entrepreneur expects to occur with probability $(1 - p_0)(1 - e_0)$, the entrepreneur learns that, with probability 1, both the idea is of low value and the strategy is ineffective. In this case, as the idea is of low value, the entrepreneur optimally abandons the venture and earns a payoff of 0.

[4] Gans, Stern and Wu (2019) in a related environment allow for multiple rounds of testing in a stylized example. We believe that the main insights regarding entrepreneurial choice can be captured with just a single round of testing as we have here.

3. For a signal of (H, B) which the entrepreneur expects to occur with probability $p_0(1-e_0)\lambda_v + (1-p_0)e_0(1-\lambda_s)$, the entrepreneur expects to earn $V_1(H,B)$ in period 1. By (A1), this signal increases the entrepreneur's posterior probability that the idea is of high value but reduces the posterior probability that the chosen strategy is effective. Thus, as $e_1(H,B) < e_0$, they would be better off implementing another random strategy rather than the tested strategy. Thus, $V_1(H,B) = \max\{e_0 p_1(H,B)v - C, 0\}$. However, by our earlier assumption that $V_0 = 0$, $e_0 p_1(H,B)v - C > 0$. This implies that following an (H, B) signal, the entrepreneur chooses to switch strategies and implement an alternative strategy.
4. For a signal of (L, G) which the entrepreneur expects to occur with probability $p_0(1-e_0)(1-\lambda_v) + (1-p_0)e_0\lambda_s$, the entrepreneur expects to earn $V_1(L,G)$ in period 1. By (A1), this signal decreases the entrepreneur's posterior probability that the idea is of high value but increases the posterior probability that the chosen strategy is effective. This suggests that the entrepreneur might be better off implementing the tested strategy than another random strategy rather than the tested strategy as $e_1(L,G) > e_0$. However, this ignores the data generating process for a signal. If it is the case that the tested strategy is really effective, this can only have arisen with an (L, G) signal that also implies that the idea is of low value. Thus, if the tested strategy is implemented, then the expected return is either $0 \times v - C$ or $1 \times 0 - C$. As these are both negative, them the optimal response to (L, G) is to implement another random strategy or abandon the venture so that $V_1(L,G) = Max\{e_0 p_1(L,G)v - C, 0\}$. However, by our earlier assumption that $V_0 = 0$, $e_0 p_1(L,G)v - C < 0$. This implies that following an (L, G) signal, the entrepreneur chooses to abandon the venture.

In summary, this testing process reveals three intuitive responses and one non-obvious or counter-intuitive response in terms of the optimal response to a signal of (L, G). When (L, G) is received, it is possible that the idea is of low value or the strategy is ineffective, but both of these cannot be true. Thus, if the venture were to continue, then it would be preferable to implement an alternative strategy as this is the *only* state in which there would be a possibility of value capture for the venture. Specifically, if (L, G) arises because the value of the idea is low, the venture should be abandoned rather than continuing with the tested strategy. Alternatively, if (L, G) arises when the idea is of value, it is an indicator that the strategy tested was ineffective. Whether (L, G) triggers a strategy switch or venture abandonment then depends upon whether $e_0 p_1(L, G)v \geq C$ or not.

The above results are summarized in the following proposition showing that *only a clear, positive signal* results in a tested strategy being implemented:

Proposition 1. *Under (A1), it is only when the signal is (H, G) is received that a tested strategy is implemented. With signals (H, B) and (L, G), selecting another strategy at random to implement than to implement the tested strategy. With (i) a signal (L, B) or (ii) a signal (L, G) under the assumption that $V_0 = 0$, the venture is abandoned.*

Proposition 1 is a result of the asymmetry between ideas and strategies. Every test involves the same idea as this cannot be changed while it is possible to perform tests with different strategies. As the signal on the two dimensions of uncertainty is conflated, changing strategies is doing the work in creating a clearer picture of the value of the underlying idea. Thus, the challenge for entrepreneurs in strategy selection is part of the process of exploring whether an idea itself has value.

Proposition 1 simplifies the actions that arise following the ambiguous signals of (H, B) and (L, G). Given this, the ex ante expected payoff to the entrepreneur becomes (recalling that $V_1(L, G) = 0$):

$$V_0(\lambda_v, \lambda_s) = p_0 e_0(v - C) + \big(p_0(1 - e_0)\lambda_v + (1 - p_0)e_0(1 - \lambda_s)\big)V_1(H, B) - c$$
$$= p_0 e_0(1 + (1 - e_0)\lambda_v)(v - C) - (1 - p_0)e_0(1 - \lambda_s)C - c$$

With a (weakly) informative signal, the value of testing a strategy is twofold. First, if s_i is an ineffective strategy and the idea is of low value, the test provides a clear signal of that and further testing and/or a launch can be prevented. This provides a direct saving of costs, $(1 - e_0)(1 - p_0)C$. Second, if s_i is signalled to be ineffective, then it can be switched for another strategy that can be implemented.[5] Thus, so long as c is not too high, there will always be value to testing a strategy in period 0.[6]

[5] In general, there will also be a third option (if $V_0 > 0$) where if a strategy is signaled to be effective it can either be implemented (including with the risk that the idea may be of low value) or the project can be abandoned with further savings of $e_0(1 - p_0)C$.

[6] Various factors impact on the cost of experimentation. For instance, the introduction of Amazon Web Services (AWS) in 2006, which allowed users to "rent" space on the cloud on a low-cost fractional basis, enabled start-ups to host online experiments at very low cost and grow with their demand (Ewens, Nanda and Rhodes-Kropf, 2018). This allowed a wide range of new ventures, such as Airbnb, Dropbox, and Uber to undertake small-scale customer experiments at low cost and with rapid customer feedback.

4 Optimal Sequencing of Strategies

Thus far, we have considered the entrepreneurial choice problem when there is no information that would distinguish one strategy from any other ex ante. Each is identical and, moreover, have independent distributions with regard to their effectiveness. While moving away from these assumptions can complicate analyses, here we consider two such deviations – irreversibility and heterogeneous priors – that would generate a motive for identifying strategies in terms of which ones are worthwhile testing before others.

4.1 Irreversibility

Strategies can differ in terms of whether they raise the costs or even render unviable other strategies. The classic example arises from the disclosure problem (Arrow, 1962). In this situation, an entrepreneur discloses the core idea publicly in order to, say, test derived products with a niche market. In the process, because the idea is disclosed, it makes it easier for imitators to commercialize the idea using different strategies (say, alternative niche or mass markets). In other words, the very act of testing a strategy makes being able to create and, notably, capture value from other strategies ineffective (Gans, Stern and Wu, 2019).

There are, of course, ways of mitigating the disclosure problem. For instance, intellectual property protection may foreclose on imitation possibilities (Gans, Hsu and Stern, 2002) and, by implication, may allow an entrepreneur to more freely engage in testing. However, obtaining such protection entails its own costs (Gans, Murray and Stern, 2017). In other situations, different tactics can be used under certain circumstances to mitigate the disclosure problem. For instance, if there are competitive options, then an entrepreneur can use that to ensure that direct expropriation of disclosed ideas (say, during licensing negotiations) do not take place (Anton and Yao, 1994; Arora, 1995 and Gans and Stern, 2000; Gans and Stern, 2003).

As a general matter, suppose that there are two strategies, s_1 and s_2, available to the entrepreneur for testing. Suppose that, for some reason, if s_2 is tested, there is a probability, k, that s_1 can no longer be implemented while the reverse is not true. In this case, it is easy to see that, in order to obtain a better signal about the value of the underlying idea, other things being equal, it is optimal to test s_1 prior to s_2. Having an alternative sequence increases the risk that it will not, in fact, be possible to test or even implement any other strategy. Thus, it is important to consider the full costs of testing (including impact on other strategies) before choosing a strategy to test.

Moreover, as will be discussed further below, expert judgment can assist in identifying the risk, *k*, associated with any given proposed strategy.

4.2 Heterogeneous priors

Another way strategies might differ ex ante is in terms of the entrepreneur's priors that a given strategy will be effective. Suppose that there are two strategies, s_1 has a prior of e_0 and s_2 with a prior $e_0' > e_0$. Which strategy should the entrepreneur test?

If we let $V_1(H,B;s_i)$ and $V_1(L,G;s_i)$ be the expected payoffs following an (*H*, *B*) and (*L*, *G*) signal arising from testing s_i, then we can write the overall payoffs from testing s_1 and s_2, respectively as:

$$\begin{aligned} V_0(\lambda_v,\lambda_s;s_1) &= p_0e_0(v-C) + \big(p_0(1-e_0)\lambda_v + (1-p_0)e_0(1-\lambda_s)\big)V_1(H,B;s_1) \\ &\quad + (p_0(1-e_0)(1-\lambda_v) + (1-p_0)e_0\lambda_s)V_1(L,G;s_1) - c \\ V_0(\lambda_v,\lambda_s;s_2) &= p_0e_0'(v-C) + \big(p_0(1-e_0')\lambda_v + (1-p_0)e_0'(1-\lambda_s)\big)V_1(H,B;s_2) \\ &\quad + (p_0(1-e_0')(1-\lambda_v) + (1-p_0)e_0'\lambda_s)V_1(L,G;s_2) - c \end{aligned}$$

Note that $p_1(H,B;s_2) < p_1(H,B;s_1)$ and $p_1(L,G;s_2) < p_1(L,G;s_1)$; that is, *testing the strategy more likely to be effective, causes you to have a lower posterior probability that the idea is good in the event that a signal is conflated*. This implies that:

$$V_1(H,B;s_1) > V_1(H,B;s_2) \Leftrightarrow e_0'p_1(H,B;s_1)v - C > e_0p_1(H,B;s_2)v - C$$

$$V_1(L,G;s_1) \geq V_1(L,G;s_2) \Leftrightarrow Max\{e_0'p_1(L,G;s_1)v - C, 0\} \geq Max\{e_0p_1(L,G;s_2)v - C, 0\}$$

Without loss in generality, let's assume that $V_1(L,G;s_2) = 0$. Then we can show that $V_0(\lambda_v,\lambda_s;s_1) > V_0(\lambda_v,\lambda_s;s_2)$.[7]

---

[7] $V_0(\lambda_v,\lambda_s;s_1) > V_0(\lambda_v,\lambda_s;s_2)$

$$\begin{aligned} &\Rightarrow p_0e_0(v-C) + \big(p_0(1-e_0)\lambda_v + (1-p_0)e_0(1-\lambda_s)\big)V_1(H,B;s_1) \\ &\quad + (p_0(1-e_0)(1-\lambda_v) + (1-p_0)e_0\lambda_s)V_1(L,G;s_1) \\ &\quad > p_0e_0'(v-C) + \big(p_0(1-e_0')\lambda_v + (1-p_0)e_0'(1-\lambda_s)\big)V_1(H,B;s_2) \\ &\Rightarrow \big(p_0(1-e_0)\lambda_v + (1-p_0)e_0(1-\lambda_s)\big)(e_0'p_1(H,B;s_1)v - C) \\ &\quad - \big(p_0(1-e_0')\lambda_v + (1-p_0)e_0'(1-\lambda_s)\big)(e_0p_1(H,B;s_2)v - C) \\ &\quad + (p_0(1-e_0)(1-\lambda_v) + (1-p_0)e_0\lambda_s)(e_0'p_1(L,G;s_1)v - C) > p_0(e_0'-e_0)(v-C) \\ &\Rightarrow \big((1-p_0)e_0'e_0\lambda_s - p_0(e_0'-e_0)(1-\lambda_v)\big)v > \big(p_0(1-e_0')(1-\lambda_v) + (1-p_0)e_0'\lambda_s - (1-p_0)(e_0'-e_0)\big)C \\ &\Rightarrow \big((1-p_0)e_0'e_0\lambda_s - p_0(e_0'-e_0)(1-\lambda_v)\big)v > e_0p_0(1-e_0')(1-\lambda_v)v - (1-p_0)(e_0'-e_0)C \\ &\Rightarrow e_0'\big((1-p_0)e_0\lambda_s - p_0(1-e_0)(1-\lambda_v)\big)v > -(1-p_0)(e_0'-e_0)C. \end{aligned}$$

Where the second to last substitution comes from the assumption that $V_1(L,G;s_2) = 0$ and the final inequality follows by (A1).

This shows that, in sequencing, it is better to test first the strategy with the *lower* prior that it will be successful than the reverse. This is because that test *provides a clearer signal of the value of the idea itself.* That information has a higher value when matched with a strategy more likely to be successful. In effect, it is better to test a riskier strategy.[8]

5 The Value of Judgment

This paper has highlighted several factors that impact on the ability of entrepreneurs to choose which strategy to implement. First, any given test involves a particular strategy, and the outcome of that test is a conflated signal between the value of the strategy and the idea. This will lead entrepreneurs to test a variety of strategies before being confident enough that the idea is sound to invest in any one of them. Second, faced with a trade-off between lower signal quality and irreversibility, entrepreneurs will tend towards a higher commitment form of experimentation when the conflation between signals of idea versus strategy is otherwise high. Finally, entrepreneurs will favor testing strategies that their prior beliefs suggest are riskier so as to obtain a clearer signal of the value of the idea. Regardless, entrepreneurial choice and the ability to use tests to explore strategies and idea value prior to significant investments is facilitated by being able to conduct tests that reduce conflation. That either means a larger number of experiments/tests of distinct strategies or a clearer expression of the signal for any one test.

One determinant of these factors comes from the entrepreneur and their founding team. Whether it be from experience or relying on some unspecified talent, some entrepreneurs may be endowed with sight into a better set of tests in order to deal with fundamental uncertainty. This is certainly implied by studies such as Shane (2000) where, in 3D printing applications, it is shown that prior experience in a particular knowledge domain played a critical role in the direction of exploration by ventures that arose in this space.

The other determinants come from forces external to the entrepreneur. Perhaps the most widely studied of these is the value of geography and the ecosystems that arise in geographic space. As is well-known, entrepreneurship tends to cluster more than most economic activity (see Guzman and Stern, 2015). While the value of such clustering comes from access to talent (Saxenian, 1996; Florida, 2005; Glaeser, Kerr and Ponzetto, 2010) and capital (Sorenson and

[8] A related result is contained in Che & Mierendorff (2019).

Stuart, 2001; Hsu, 2006), and in some cases, basic research (Agrawal and Goldfarb, 2008; Delgado, Porter and Stern, 2010), there is a strong sense that the value of co-location comes from externalities that are, in Alfred Marshall's phrase, "in the air." In other words, it is difficult with regard to resources to decouple cause from effect, but also there is a sense in which there is a tacit dimension to what arises in the culture, expectations and tolerance for factors such as a failure, that can arise in geographical ecosystems.

Our focus here on the factors that facilitate entrepreneurial choice provides an opportunity to ground the benefits that might otherwise be "in the air." As mentioned above, our perspective here suggests that when there is an improved ability to foster entrepreneurial testing, better choices will arise. Moreover, those tests arise from superiority in the clarity of the signals generated by tests while economizing on the need for commitment – in other words, *judgment* that gives entrepreneurs ways of more effectively testing their various options.

A set of examples of such judgment comes when one considers the design of various programs designed to accelerate startup commercialization. In fact, a third of all startups that raised venture capital in 2015 had been through an accelerator program (Pitchbook, 2016). Furthermore, at least some accelerators enhance startup performance. In three of the four cohorts studied by Hallen et al. (2019), accelerator participants raised 47-171% more funds in the subsequent 2-3 years than the matched, "almost-accepted" applicants. However, accelerators are not all the same. Some differ in terms of the length of the engagement. For example, Howell (2017) reports that entrepreneurs learn from their one-time interactions with experts in the context of feedback during business plan competitions. However, longer programs facilitate a longer duration of engagement between entrepreneurs and agents (mentors). Some programs offer more concentrated consultations. For example, Cohen et al. (2018) examine eight accelerators with 37 startups nested within the programs and discover that four of them spread out their interactions while the other four schedule them in a concentrated block. The total number of mentor meetings varies from fewer than 10 up to 75. Finally, others offered a longer program with various meetings and paced learning as well as a model that pools judgment from a number of sources. Lakhani et al. (2019) describe the potential benefits of this model, exploring the case of the Creative Destruction Lab that originated at the University of Toronto.

In the context of the model presented here, judgment is the ability of an agent to suggest to the entrepreneurs, strategic tests with either a lower c, higher λ_v, λ_s, or reduced irreversibility. The central proposition is that the entrepreneur has knowledge of a certain set of tests, but an agent with judgment can enlarge that set and put forward options that are more efficient in facilitating choice.

However, judgment does not come for free. Those agents may incur costs in considering and generating more efficient tests for the entrepreneur to consider. More critically, those agents may hold different priors (p_0, e_0) to the entrepreneur. In each case, these factors may create a non-alignment of interests between that agent and the entrepreneur. This will complicate how advice is given and received.

5.1 Investment in Precision

As the effects here can be somewhat subtle, we amend our baseline model in a minor way to capture the role of differing priors. We assume that, other than that difference, both the entrepreneur and the agent with judgment (whom we term 'mentor'), have aligned interests. That is, they both place the *same* weight on the success of the venture and the costs involved in performing tests.

To capture this, suppose there are two strategies that can be tested with greater precision: one with $\lambda_s = 1$ and the other with $\lambda_v = 1$. In other words, there is a 'strategy-biased' test (yielding a clearer signal of the effectiveness of the strategy) and a different 'idea-biased' test (yielding a clearer signal of the idea's value). Which test should the entrepreneur perform?

The following proposition answers this question.

Proposition 2. *A test with high λ_s is preferred to a test with high λ_v if $(1 - p_0)e_0(1 - \lambda_s)C > (1 - \lambda_v)p_0(1 - e_0)(e_0 v - C)$.*

The primary benefit from a high λ_s is that it allows the entrepreneur to avoid implementing a strategy when the idea does not have value while the benefit from a high λ_v is that it creates an opportunity for the entrepreneur to implement an alternative strategy with a valuable idea. Thus, the difference between more precision in either signal is that, in one case, it leads to more efficient abandonment of the venture while in the other more efficient continuation of it.

To put this another way, *the test preference is biased by the prior*. If your prior is that effective strategies are relatively easy to find but that the idea may not be valuable, you choose a high λ_s so that if the outcome is not (l, g) you receive a clearer signal. By contrast, if your prior is that the idea is valuable but effective strategies are hard to find, you choose a high λ_v that allows you to receive a clearer signal when the 'surprising state' (*l*, *g*) arises.

5.2 *Impact of Differing Priors*

Proposition 2 shows that a mentor will propose a test that is biased towards their priors. That is because an agent, even if they understand that another agent has a different prior (that is, they agree to disagree), will determine outcomes based on their own priors. In other words, a mentor who has the same interest as the entrepreneur in terms of maximizing the expected value of the venture will propose tests that they believe will lead to optimal outcomes given their own beliefs. An entrepreneur, who may only have available tests with precision dictated by (λ_v, λ_s), will use the suggested test because it has a higher precision on at least one element of interest even if this test differs from the test they would have preferred based on their priors.

The differences between the entrepreneur and mentor in terms of their prior beliefs have more bite when those differences would also change the decisions made.[9] Here, the optimal actions following any signal are determined by other prior assumptions (and Proposition 1). Thus, there is no such conflict. However, to illustrate what might happen if there were such a difference, we adopt a slight amendment to the model. In particular, we relax the assumption that $V_0 = 0 \Leftrightarrow C = p_0 e_0 v$ and instead assume that *C* is a freely moving variable.

To highlight entrepreneur and mentor differences in judgment, let e_E be the prior on a strategy by the entrepreneur and e_M be the prior on a strategy by the mentor. We assume that each agent shares the prior, p_0. But that $e_E > e_M$. Given this, it is possible that the entrepreneur, following a signal (*H*, *B*), would continue the venture with another random strategy while the mentor, following that signal, would abandon the venture. This shows what happens when there are differences in opinions between agents, although it could easily be the case that the interests are reversed. Illustrating one case suffices to demonstrate our main point that differing priors are, potentially, a cost of independent judgment.

We consider the following game:

1. The mentor, with prior e_M over all strategies in *S*, proposes a test with either $\lambda_s = 1$, $\lambda_v = 1$ or (λ_v, λ_s) as in the baseline to the entrepreneur

[9] There is a growing literature on the implications of differing priors to management and economics. See, for instance, van den Steen (2010), Che and Kartik (2009), and Che, Dessein & Kartik (2013).

2. The entrepreneur, with prior e_E over all strategies in *S*, can either undertake the proposed test or the baseline test with (λ_v, λ_s).
3. The entrepreneur takes actions and payoffs are realized.

We already argued that, in stage 3, the entrepreneur would take actions based on their own priors and hence, will choose to continue the venture following a signal of (*H*, *B*). Moreover, for the same reason, the entrepreneur, in stage 2, will choose to use the mentor's proposed test over the baseline test, if it is proposed. Thus, the outcomes in this game are wholly determined by the mentor's decision as to what test, if any, to propose to the entrepreneur.

Differing priors mean that, depending on the test used, there is a potential disagreement as to whether the venture should be abandoned or continue after a signal of (*H*, *B*). For notational purposes let's denote our three tests as (λ_v, λ_s), $(\lambda_v, 1)$ or $(1, \lambda_s)$ where the first test is the baseline, and the others have obvious interpretations. In this case, it is easy to see that:

$$V_1(H, B, (\lambda_v, 1), e_i) > V_1(H, B, (1, \lambda_s), e_i) > V_1(H, B, (\lambda_v, \lambda_s), e_i)$$

for $e_i = e_M$ or e_E. Moreover, for any test, $V_1(H, B, (.,.), e_E) > V_1(H, B, (.,.), e_M)$. Thus, if the mentor prefers to continue, this implies the entrepreneur does too but not vice versa. Note, however, that the different tests have different probabilities of the signal (*H*, *B*) being generated. Specifically,

$$\Pr[H, B, (1, \lambda_s)] > \Pr[H, B, (\lambda_v, \lambda_s)] > \Pr[H, B, (\lambda_v, 1)]$$

So long as $V_1(H, B, (\lambda_v, \lambda_s), e_M) \geq 0$, the mentor's preferences will be akin to those given by Proposition 2 but because of the differences in priors, they may not be aligned with the entrepreneur's preferences. In particular, the mentor may propose a test that would not be the test the entrepreneur would have wanted. Nonetheless, both agents agree that the venture should continue if an (*H*, *B*) signal is received.

Conflicts of interest may arise if the agents disagree as to what should happen in the event of an (*H*, *B*) signal following certain tests. To illustrate, suppose that $V_1(H, B, (1, \lambda_s), e_M) < 0 < V_1(H, B, (1, \lambda_s), e_E)$ while $V_1(H, B, (\lambda_v, 1), e_i) \geq 0$ In this case, the mentor will never propose a test $(1, \lambda_s)$ because this would lead the entrepreneur to continue while the mentor believes that it should be abandoned following an (*H*, *B*) signal with that test. Instead $(\lambda_v, 1)$ will be proposed as this test generates actions where interests are aligned.

In some cases, the mentor may not be able to induce the entrepreneur to take its preferred action and so instead will be motivated to propose a test to minimize their perceived harm from the entrepreneur's behavior. Suppose, for instance, that $V_1(H, B, (\lambda_v, 1), e_M) < 0$ while $V_1(H, B, (\lambda_v, \lambda_s), e_E) > 0$. In this case, the mentor would prefer that the entrepreneur never continued the venture following an (H, B) signal, but the entrepreneur is sufficiently optimistic that they continue regardless of the test used. Here, the mentor would again propose $(\lambda_v, 1)$, not to ensure aligned actions are taken (that is not possible) but instead to minimize the probability of an (H, B) signal arising.

There are several conclusions that can be drawn from this analysis. First, differing priors mean that a mentor will consider how those differences impact on the entrepreneur's actions in proposing tests. For this reason, the mentor may not propose tests that the entrepreneur themselves would have selected. Second, the differences in priors reduce the value to the mentor of a test with $\lambda_v = 1$. Recall that such a test would imply that an (L, G) signal gave clear information on whether the tested strategy was effective or not and consequently, clarity on the value of the idea. However, the (H, B) test would be left with conflated signals. That lack of clarity increases the likelihood of an erroneous choice following an (H, B) signal and the differing priors imply that the trade-offs on the types of errors made differ between mentor and entrepreneur. To avoid that conflict, the mentor avoids that test. Finally, if, as would be reasonable, there were costs associated with more precise tests (Pomatto, Strack and Tamuz, 2020), and if those costs were high, then it is possible that the mentor would choose not to propose a precise test even if the entrepreneur would have preferred one despite those costs.

5.3 Other considerations

The analysis of differing priors demonstrates the challenges of mentorship and communicating judgment in innovative settings. However, there are other considerations that can also pose challenges and it is useful to review them here.

First, even if mentors and entrepreneurs place the same weight on the outcomes for the venture, their alternatives – should the venture not proceed – are likely to differ. For instance, an entrepreneur may have alternative employment meaning that its threshold for abandoning the venture is greater than zero. For a mentor, they may have a number of ventures that they could be directing the effort towards. Thus, for them, if the venture does not proceed, they will likely mentor another venture. It is that alternative venture's expected outcomes that determine the mentor's outside option.

Importantly, these considerations can influence the tests proposed by the mentor. For instance, in choosing between a test that more strongly signalled the value of the idea as opposed to the value of the strategy, a mentor may be more interested in the former signal as it determines with more clarity that the venture they are mentoring is worth 'investing' time in so as to find an effective strategy as opposed to choosing another venture. Similarly, a mentor, who understands that an entrepreneur will only continue to work on the venture if the prospects are sufficiently strong will look to propose tests that are more likely to generate earlier signals that mitigate downside risk rather than ones that provide signals of potentially high returns – although even this is likely to depend on other characteristics of their entrepreneur – including their aversion to risk.

Second, what type of mentor might the entrepreneur select to work with if they understand some of the challenges already mentioned here? In terms of differing priors, does an entrepreneur want a mentor that shares their priors? On the one hand, there will be an alignment of interests in terms of tests proposed if the mentor and entrepreneur are like-minded in terms of priors. On the other hand, if the mentor has to engage in effort to propose tests and convince the entrepreneur, as Che and Nartik (2009) show, a mentor may put in less work to formulate a clearer test when the task of convincing the entrepreneur is 'easy.' In this case, some degree of difference in priors may be optimal for the entrepreneur. This becomes even stronger when the entrepreneur takes advice from multiple sources.

5.4 *Empirical implications*

A central proposition is that the entrepreneur has knowledge of a certain set of tests, but an agent with judgment, a 'mentor,' is able to enlarge the set of possible tests and/or more effectively choose from the existing set and thereby increase λ_v, increase λ_s, lower c, or reduce irreversibility. For example, mentors may be able to improve the entrepreneur's selection of experiments, generating clearer signals of idea value (higher λ_v) or strategy effectiveness (higher λ_s). Of course, entrepreneurs must be receptive to learning from mentors' judgment in order for this to be effective. Howell (2017) provides empirical evidence from business plan competition data that, on average, "nascent entrepreneurs are quite responsive to feedback, and furthermore the ability to learn [from mentors] is an important determinant of success."

What factors influence the ability of mentors to provide judgment for entrepreneurs? A nascent but growing line of inquiry explores this question and provides preliminary descriptive evidence of several margins associated with meaningful variance in the effectiveness of transferring judgment from agents to entrepreneurs. Most of these studies are descriptive; they provide anecdotal evidence or descriptive statistics but are not able to provide evidence of causality. For example, Lakhani et al. (2019) describe the process of mentors providing their judgment to entrepreneurs in the context of objective-setting in the context of a particular program for entrepreneurs, the Creative Destruction Lab. In this case, experienced entrepreneurs and venture investors provide their judgment to entrepreneurs by helping them set three objectives in a series of eight-week cycles.

Examining mentorship judgment in this setting, Sariri (2020) classifies the objectives proposed by entrepreneurs and compares them to objectives proposed by mentors. In many cases, one or more of the objectives involve conducting a test. Not only do the mentors specify the objectives of the test (e.g., how and what to measure), but sometimes also provide guidance on how to conduct the test. Tests can be designed to reduce market demand uncertainty (e.g., interviewing customers, obtaining letters of intent, or selling paid pilots), reduce technological uncertainty (e.g., building a lab-scale proof of concept, conducting an onsite demonstration, or performing a head-to-head comparison against a competing product) or understand novelty (e.g., by obtaining a patent, publishing a peer-reviewed paper, or obtaining a customer commitment in the face of competing products). While this is one of the first empirical studies of mentor judgment as it relates to tests, this research does not distinguish between tests that enhance the signal quality on ideas versus strategies; nor does it estimate the relative costs of different tests or their degree of irreversibility. Overall, research on the classification of mentor judgment in terms of its impact on λ_v, λ_s, c, or irreversibility and measuring its costs and benefits is in its infancy.

6 Future Directions

This paper was written for a conference celebrating 65 years of *Management Science*. In doing this, we wanted to review and explore a topic that we believe will impact the next few decades of research into entrepreneurship and strategy: entrepreneurial choice. The entrepreneur's choice, in the face of uncertainty, on how to bring their idea to market is often irreversible and at the same time a primary determinant of success. To examine the nature of this choice, we drew upon insights from disparate literatures – strategy, finance, economics and innovation as well as the operations management of optimal search and decision analysis – to present a simple framework for evaluating the costs and benefits of generating signals regarding the quality of the idea as distinct from the quality of the strategy for bringing that idea to market. The challenge was the fact that learning about the value of an idea requires testing alternative strategies which mean that conflation of signals is an underlying environmental fact requiring various, distinct process choices. We demonstrated the value of multiple testing, of choosing tests to expose, where possible, the value of the idea, to manage irreversibility and finally, to engage with others (i.e., mentors) in a way that both takes advantage of and makes clear their own biases and priors.

While exploring the implications and efficacy of entrepreneurial search processes requires more theoretical analysis than we could provide here, the work here suggests a key hypothesis regarding entrepreneurial success and the growth of sustainable businesses: institutions and other factors that enable entrepreneurial choice – that is, testing with a superior trade-off between experimental costs and the resulting signal to noise ratio – will be a key factor in explaining entrepreneurial success and its dual outcome of fast and efficient failure. Our belief is that exploration and examination of this broader hypothesis will require the combination of more related managerial and social science areas and represents a path forward for general prescriptive measures to improve entrepreneurial outcomes.

Our framework highlights another area for future research: competition. Mentor judgment is a scarce resource. How does competition in the market for mentor judgment impact outcomes? On the supply side, how do individuals or organizations obtain judgment? Gans (2018) suggests that judgment may result from experience. People learn about the efficacy of tests from observing the quality and costs of signals across multiple applications. In other words, judgment is developed from learning by doing (Arrow, 1971). How do market participants compete in the labor market for developing judgment?

On the demand side, how do startup firms compete in the market for obtaining judgment? For example, startups that receive multiple financing offers in their first round of institutional capital often choose to accept offers made by venture capital firms with high-reputations (presumed judgment) despite the fact the capital comes at a higher cost (Hsu, 2004). Specifically, Hsu reports evidence that startups that receive multiple offers are three times more likely to select offers from high reputation VCs even though their capital costs significantly more: a 10–14% discount. While some of this price premium is likely due to other attributes, such as profile, connections, and access to follow-on financing, a portion of the premium is likely due to the market price for judgment. Future research will examine competition in the market for judgment and measure as well as theorize about variation in the willingness to pay for it.

7 References

Admati, A.R., and Pfleiderer, P. 1994. Robust Financial Contracting and the Role of Venture Capitalists. *The Journal of Finance,* 49(2), pp.371-402.

Agrawal, A. and Goldfarb, A., 2008. Restructuring research: Communication costs and the democratization of university innovation. *American Economic Review*, *98*(4), pp.1578-90.

Anton, J. J., and Yao, D. A. 1994. Expropriation and inventions: Appropriable rents in the absence of property rights. *American Economic Review*, pp.190-209.

Arora, A., 1995. Licensing tacit knowledge: intellectual property rights and the market for know-how. *Economics of innovation and new technology*, *4*(1), pp.41-60.

Arrow, K.J. 1962. Economic Welfare and The Allocation of Resources for Invention. *The Rate and Direction of Inventive Activity, Princeton University Press and NBER*.

Arrow, K.J. 1971. The Economic Implications of Learning by Doing. *Readings in the Theory of Growth*, ed F.H. Hahn, Palgrave Macmillan, pp 131-149.

Astebro, T. and Koehler, D., 2007. Calibration Accuracy of a Judgmental Process that Predicts the Commercial Success of New Product Ideas. *Journal of Behavioural Decision Making*.

Azoulay, P. and Shane, S., 2001. Entrepreneurs, contracts, and the failure of young firms. *Management Science*, 47(3), pp.337-358.

Bergemann, D.,and Hege, U., 2005. The financing of innovation: Learning and stopping. *RAND Journal of Economics*, 719–752.

Chavda, A., Gans, J.S. and Stern, S., 2019. A Model of Entrepreneurial Strategic Experimentation. *Working paper.*

Che, Y.K., Dessein, W., and Kartik, N. (2013). Pandering to persuade. *American Economic Review*, *103*(1), 47-79.

Che, Y.K., and Kartik, N. (2009). Opinions as incentives. *Journal of Political Economy*, *117*(5), 815-860.

Che, Y.K. and Mierendorff, K. (2019), Optimal Dynamic Allocation of Attention. *American Economic Review*, *109*(8), pp.2993-3029.

Cohen, S., D.C. Fehder, Y.V. Hochberg, and F. Murray, 2019. The Design of Startup Accelerators. *Research Policy*, 48, pp. 1781-1797.

Delgado, M., Porter, M.E. and Stern, S., 2010. Clusters and entrepreneurship. *Journal of Economic Geography*, 10(4), pp. 495-518

Elkhorne, J.L. 1967. Edison — The Fabulous Drone, in 73 Vol. XLVI, No. 3 (March).

Erkal, N., and Scotchmer, S. (2009). "Scarcity of Ideas and R&D Options: Use it, Lose it or Bank it," (No. w14940). National Bureau of Economic Research.

Ewens, M., Nanda, R. & Rhodes-Kropf, M., 2018. Cost of experimentation and the evolution of venture capital. *Journal of Financial Economics*, 128(3), pp.422-442.

Fiet, J.O. and Patel, P.C., 2008. Entrepreneurial discovery as constrained, systematic search. *Small Business Economics*, 30(3), pp.215-229.

Florida, R., 2005, *Cities and the Creative Class*, Routledge: New York, NY.

Gabler, N., 2006. *Walt Disney*. Vintage: New York.

Gans, J.S., 2018. A Better Way to Bring Science to Market. *Sloan Management Review*. https://sloanreview.mit.edu/article/a-better-way-to-bring-science-to-market/

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), pp. 571-586.

Gans, J.S., F. Murray, S. Stern, 2017. Contracting over the disclosure of scientific knowledge: Intellectual property and academic publication. *Research Policy*, *46*(4), pp.820-835.

Gans, J.S. and Stern, S., 2000. Incumbency and R&D Incentives: Licensing the Gale of Creative Destruction. *Journal of Economics & Management Strategy*, *9*(4), pp.485-511.

Gans, J.S. and Stern, S., 2003. The product market and the market for 'ideas': commercialization strategies for technology entrepreneurs. *Research policy*, *32*(2), pp.333-350.

Gans, J.S., Stern, S. and Wu, J. 2019. Foundations of Entrepreneurial Strategy. *Strategic Management Journal*, 40(5), pp.736-756

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

Glaeser, E.L., Kerr, W.R. and Ponzetto, G.A., 2010. Clusters of entrepreneurship. *Journal of Urban Economics*, *67*(1), pp.150-168.

Guzman, J., and Stern, S., 2015. Where is Silicon Valley? *Science*, 347(6222): 606–609.

Howell, S.T., 2017. Learning from Feedback: Evidence from New Ventures. Working paper.

Hsu, D.H., 2004. What Do Entrepreneurs Pay for Venture Capital Affiliation? *Journal of Finance*, 59, pp. 1805-1844.

Hsu, D.H., 2006. Venture capitalists and cooperative start-up commercialization strategy. *Management Science*, *52*(2), pp.204-219.

Kerr, W.R., Nanda, R. and Rhodes-Kropf, M., 2014. Entrepreneurship as experimentation. *The Journal of Economic Perspectives*, *28*(3): pp.25-48.

Kirzner, I 1997. Entrepreneurial Discovery and the Competitive Market Process: An Austrian Approach. *Journal of Economic Literature*, 35:1, 60-85.

Lakhani, K, H. Luo, and L. Katsnelson, 2019. Market for Judgement: Creative Destruction Lab, working paper, Harvard Business School Case.

Manso, G., 2011. Motivating innovation. *The Journal of Finance*, 66(5), pp.1823-1860.

Murray, F. and Tripsas, M., 2004. The exploratory processes of entrepreneurial firms: The role of purposeful experimentation. *Advances in strategic management*, *21*, pp.45-76.

Nanda, R. and Rhodes-Kropf, M., 2016a. Financing entrepreneurial experimentation. *Innovation Policy and the Economy*, 16(1), pp.1-23.

Nanda, R. and Rhodes-Kropf, M., 2016b. Financing risk and innovation. *Management Science*, 63(4), pp.901-918.

Pitchbook. 2016. One-third of U.S. startups that raised a Series A in 2015 went through an accelerator.

Pomatto, L., P. Strack and O. Tamuz. 2020. The Cost of Information, mimeo., Caltech.

Ries, E., 2011. *The lean startup: How today's entrepreneurs use continuous innovation to create radically successful businesses*. Currency.

Sariri, A., 2020. Learning versus Doing: The Effect of Business Uncertainty on Entrepreneurial Activities. Working paper, University of Toronto.

Saxenian, A., 1996. *Regional advantage*. Harvard University Press: Boston, MA.

Scott, E.L., P. Shu, and R.M. Lubynsky, 2019. Entrepreneurial Uncertainty and Expert Evaluation: An Empirical Analysis. *Management Science*, Articles in Advance, pp 1-22.

Shane, S.A., 2000. Prior knowledge and the discovery of entrepreneurial opportunities. *Organization Science*, 11(4), pp.448-469.

Shane, S.A., 2008. *The illusions of entrepreneurship: The costly myths that entrepreneurs, investors, and policy makers live by*. Yale University Press.

Sorenson, O. and Stuart, T.E., 2001. Syndication networks and the spatial distribution of venture capital investments. *American journal of sociology*, *106*(6), pp.1546-1588.

Van den Steen, E. (2010). Culture clash: The costs and benefits of homogeneity. *Management Science*, *56*(10), 1718-1738.

Weitzman, M.L., 1979. Optimal search for the best alternative. *Econometrica*, 47(3), pp.641-654.