This is a portion of the eBook at doi:10.7551/mitpress/15918.001.0001

1

The Foundations of Bayesian Entrepreneurship

Ajay Agrawal, Arnaldo Camuffo, Alfonso Gambardella, Joshua Gans, Erin L. Scott, and Scott Stern

1.1 Introduction

The study of entrepreneurship is inherently linked to the study of decision-making. As a phenomenon, potential entrepreneurs choose to pursue some sort of opportunity (or not), undertake activities that allow them to learn about that opportunity (or not), and make choices that allow them to realize the potential of that opportunity (or not). At each stage of this dynamic process, founders face a set of decisions that inevitably shape the future path by which they pursue a particular opportunity and the organization they build in order to take advantage of that opportunity. There is, of course, no requirement that entrepreneurs (or others) make these decisions optimally, but at the same time, there is no requirement that entrepreneurial decision-making is inherently flawed or biased. There is, of course, no requirement that all entrepreneurs make their decisions in the same way, and certain approaches to entrepreneurial decision-making may be far more effective than others. Centering the role of decision-making in entrepreneurship is simply highlighting that entrepreneurs do indeed engage in decision-making (i.e., they are *not* simply following directions) and that their decision-making is consequential for performance. The decisions they make, and possibly how they are made, affect the financial and nonfinancial benefits they accrue.

A central challenge, then, in the field of entrepreneurship is the development of a theoretical framework that is simultaneously broad enough to allow for a wide range of entrepreneurial decision-making approaches yet also specific enough that researchers are able to ask and answer questions of fundamental theoretical and empirical interest and potentially provide guidance for practitioners or policymakers. Striking this balance has proved elusive. Consider the contrast between two broad approaches that characterize a large fraction of existing studies of entrepreneurial decision-making: environmental and behavioral. On the one hand, the environmental approach, grounded in economics

Downloaded from direct.mit.edu/books/oa-edited-volume/chapter-pdf/2595206/c000600_9780262052160.pdf by Università Bocconi user on 04 November 2026

This is a portion of the eBook at doi:10.7551/mitpress/15918.001.0001

and finance, largely abstracts away from how entrepreneurs identify opportunities or how they make decisions and focuses instead on the impact of the microeconomic, strategic, and institutional environment on the choice to enter entrepreneurship or the impact of entrepreneurship on outcomes such as earnings or wealth (Evans and Jovanovic, 1989; Hamilton, 2000; Hurst and Lusardi, 2004; Lazear, 2004; Åstebro et al., 2011). On the other hand, more behavioral approaches have tended to abstract away from the potential for entrepreneurs to optimize and focus more on the role of the background and experiences of entrepreneurs (Shane, 2003), the propensity for entrepreneurs to learn through action and local contingency (Bhide, 1994; Sarasvathy, 2001), and the potential role of biases and heuristics to distort entrepreneurial choice. Both of these broad approaches and their many offshoots have proven valuable in shaping our understanding of the role of the external environment in shaping entrepreneurial choice and the role of entrepreneurial motivation and behavior in enacting those choices.

Over the past decade, a third approach has emerged that focuses squarely on the interplay between entrepreneurial decision-making and the process of entrepreneurial learning (Rosenberg, 1992; Murray and Tripsas, 2004; Sull, 2004). Three related (but distinct) areas are highlighted in this rapidly emerging body of research. First, entrepreneurs are decision-makers making active choices under uncertainty (i.e., entrepreneurship is itself an experiment), including critically the choice to pursue a particular opportunity (Nanda and Rhodes-Kropf, 2016; Manso, 2016). In other words, although entrepreneurs may be subject to biases or constrained by their local environment, the choice to pursue a particular opportunity (and not others) is clearly a choice made by the entrepreneurs themselves, and the construction of the opportunity is itself uncertain (or else there would be no opportunity to pursue). Second, whether and how to pursue an opportunity is shaped by purposeful learning and experimentation as a tool for entrepreneurial decision-making (among many others, Gans et al., 2019, Camuffo et al., 2020, and, importantly, Ries, 2011, in the domain of entrepreneurial practice). Entrepreneurial experimentation allows entrepreneurs to (somehow) reduce the uncertainty associated with the pursuit of an opportunity, and this process of uncertainty reduction shapes the choices that entrepreneurs then make on an ongoing basis. However, as emphasized by Camuffo et al. (2020), we are at an early stage of understanding how entrepreneurs learn from experiments and how different types of experimentation (e.g., more iterative vs. scientific approaches) impact entrepreneurial strategy and performance. Finally, to ultimately realize the potential of the opportunity they are pursuing, entrepreneurs must separately *persuade* others of the value associated with that opportunity. Success in

Downloaded from direct.mit.edu/books/oa-edited-volume/chapter-pdf/2595206/c000600_9780262052160.pdf by Università Bocconi user on 04 November 2026

This is a portion of the eBook at doi:10.7551/mitpress/15918.001.0001

pursuing a given opportunity depends on the ability of the entrepreneur to attract resources and capabilities from others (Stevenson, 1983), and the choices that entrepreneurs make are not in isolation. Instead, the successful realization of an entrepreneurial opportunity depends on choices by other decision-makers—potential investors, employees, and even customers—who will also be making decisions about whether to invest time, effort, or resources to engage with a nascent venture. For each of these domains—how entrepreneurs form their initial beliefs about opportunities, how they learn from proactive experimentation, and how they ultimately persuade others (or not)—this work confronts directly the fact that entrepreneurial decision-making takes place under significant uncertainty and that a primary means by which entrepreneurs realize the opportunity is through the resolution of that uncertainty (in ultimately a public way).

The purpose of this chapter is to synthesize the different streams of work in this third approach under an umbrella that we will refer to as "Bayesian entrepreneurship." This synthesis builds on two distinct but related insights. First, different (and often disparate) studies of the process of entrepreneurial decision-making—the formation of beliefs, entrepreneurial learning and experimentation, and entrepreneurial communication—are each linked to one another. At a broad level, what you believe influences what types of experiments and learning you might undertake, and each of these will then influence what other actors you are able to persuade in order to attract resources and capabilities. Second, a necessary condition on the nature of the pursuit of entrepreneurial opportunity allows us to place additional structure on this linkage: Entrepreneurs *choose* the opportunities they pursue. Although this may seem obvious, the fact that entrepreneurs choose to pursue a *particular* opportunity has far-reaching implications. Simply put, entrepreneurs (almost by construction) hold "contrarian" beliefs about the opportunities they choose to pursue (Masters and Thiel, 2014). Almost by construction, an entrepreneur choosing a particular opportunity must hold a *relatively more positive* belief about the potential for that opportunity. Conversely, if entrepreneurs believed that the value of the opportunity they were pursuing was equally valued by others (who perceived the same chance of success, risk, etc.), then the rationale for pursuing that opportunity would disappear. Bayesian entrepreneurship brings these two insights together: How does the fact that entrepreneurs choose the opportunity they pursue impact entrepreneurial experimentation, strategy, and performance?

Consider the case of Zappos (as memorably recounted in Hsieh, 2010), founded by Nick Swinmurn based on his belief in the potential for online shoe retail. Although there were, of course, many other online retail businesses

Downloaded from direct.mit.edu/books/oa-edited-volume/chapter-pdf/2595206/c000600_9780262052160.pdf by Università Bocconi user on 04 November 2026

This is a portion of the eBook at doi:10.7551/mitpress/15918.001.0001

started during the dot-com "boom," very few other growth-oriented businesses entered the shoe category, as footwear traditionally required customized fittings, bulky shipping, and a high level of returns and exchanges. Despite these observable challenges (which presumably dissuaded others) and his lack of background in the shoe industry (or even retail!), Swinmurn formed a specific and clear hypothesis about why there was such a large opportunity in the footwear domain: Although footwear was a large and growing category ($40 billion), retail was highly fragmented, and so there was an opportunity to create a one-stop online offering that was "a network among all the separate shoe stores." Armed with his contrarian hypothesis, he sought to attract resources (in the form of venture capital) and capabilities (in the form of retail experience) that would help him realize the value of this opportunity. Notably, although he was rejected or dismissed by many, those who joined at an early stage—investors Alfred Lin and Tony Hsieh and former Nordstrom's executive Fred Mossler—were explicitly attracted to join the venture precisely because they shared the same contrarian theory held by Swinmurn (with the understanding that others who did not fund or join disagreed with them). However, this shared belief by the founding team then had consequences for the venture going forward. Specifically, although the founders interpreted their early growth numbers in a positive light (believing it gave them a reason to continue to move forward), larger venture capitalists (most notably Sequoia) looked at the same data and came to the conclusion that the venture was unlikely to succeed (and declined to invest). Finally, despite this feedback from Sequoia, and with cash running perilously low, the team (contrary to prevailing wisdom to reduce expenses to extend the runway) opted to allocate a significant portion of its dwindling resources toward purchasing inventory. This decision was underpinned by the hypothesis that a substantial in-stock inventory, coupled with an exceptional customer service experience, would catalyze e-commerce traction. More precisely, Hsieh characterizes the experiment (which involved investing the remainder of his own personal financial resources) as a test to "either save Zappos or ensure our speedy demise." The experiment itself was a success, and after a record run of observable sales growth and bottom-line performance, more traditional venture capitalists (including Sequoia!) invested in more mature stages of the company prior to its acquisition by Amazon in 2009.

Although most start-ups do not experience the success of Zappos, its history nonetheless captures a number of recurrent themes that are consistent with a Bayesian entrepreneurship approach. First, although Nick Swinmurn may be more optimistic or not than the general population in terms of background effect, there is no doubt that Swinmurn was specifically optimistic about the opportunity associated with online retail. Importantly, he did not have an extensive

Downloaded from direct.mit.edu/books/oa-edited-volume/chapter-pdf/2595206/c000600_9780262052160.pdf by Università Bocconi user on 04 November 2026

This is a portion of the eBook at doi:10.7551/mitpress/15918.001.0001

background or prior career in this domain but instead observed the opportunity as a particularly fruitful (albeit particularly challenging) path to pursue in the broader domain of online commerce. This characteristic feature brings us to our first broad insight into Bayesian entrepreneurship: Since an entrepreneur *chooses* the specific opportunity they are pursuing, a given entrepreneur is *relatively optimistic* about that opportunity relative to other decision-makers. In other words, a Bayesian approach captures the idea that an entrepreneur has a systematically different theory of why an opportunity is valuable relative to others who have chosen *not* to pursue that opportunity. More generally, Bayesian entrepreneurs are endowed with relatively favorable priors about the opportunities they choose to pursue, a state that is also associated with a contrarian hypothesis about how they might create and capture value from that opportunity in a way that others do not foresee.

Second, the Bayesian approach offers distinct and novel insight into the nature of purposeful learning and experimentation on the part of founders. First, at a broad level, the fact that founders have relatively more positive priors does not imply that they are *certain*; outside of degenerate cases, priors are noisy, and the "theory" underlying a given prior may be more or less well developed. As such, the existence of calculable uncertainty about an opportunity leads founders to have a *demand for experimentation*. In particular, the demand for experiments (which are costly in terms of time, effort, and, potentially, opportunity cost) will be shaped by the priors held by the entrepreneur. When founders are more uncertain about their theory of value creation and capture or hold more noisy priors, the demand for experimentation will increase. Just because the founding team at Zappos had more positive beliefs did not mean they were immune to feedback or learning through experimentation; instead, they tailored experiments that allowed them to update their priors in order to both assess whether the venture was likely to be viable and make decisions about their overall entrepreneurial strategy. Moreover, this demand for experiments (i.e., opportunities to learn about the opportunity in order to make better decisions) is fundamentally influenced by beliefs about the opportunity. For example, relative to a disinterested agent, a Bayesian entrepreneur will conduct experiments that are more likely to allow for false positives (a signal to continue despite the opportunity not being valuable) than false negatives (a signal to disband even though the opportunity is worthwhile). This sort of biased experiment (in which entrepreneurs choose a "best foot forward" approach that allows them to only receive negative feedback in the case when the opportunity is, in fact, not worthwhile) can be seen directly in the experiment chosen in Zappos where they undertook a costly experiment that was "optimized to succeed" in order to give the team the confidence that they should continue to pursue the opportunity.

Downloaded from direct.mit.edu/books/oa-edited-volume/chapter-pdf/2595206/c000600_9780262052160.pdf by Università Bocconi user on 04 November 2026

This is a portion of the eBook at doi:10.7551/mitpress/15918.001.0001

Finally, this approach yields novel insight into who (and how) entrepreneurs are able to attract resources and capabilities by providing information to other agents. Notably, entrepreneurs are more likely to persuade those who share their idiosyncratically more positive beliefs about an opportunity. For example, in the Zappos case, the investors and industry experts who were attracted to the venture were attracted to it precisely because they held the same optimistic beliefs as the founder of the original idea. This Bayesian dynamic suggests that the emergence of homophily within start-up organizations might reflect not simply a preference for association among similar types but attraction to an opportunity for those who share similar beliefs (Van den Steen, 2005, 2011). More subtly, the existence of heterogeneous priors by different agents suggests that, when needing to attract resources and capabilities of those whose priors differ from that of the entrepreneur, the experiments that are conducted (and the way that evidence is interpreted) will cater to those with different (more negative) beliefs than they themselves hold. In other words, the natural distortions that arise in attempts at Bayesian persuasion (Kamenica and Gentzkow, 2011) are a natural consequence of a Bayesian entrepreneurship approach.

The remainder of the chapter proceeds as follows. First, we undertake a bit of a deck-clearing exercise by delineating what Bayesian entrepreneurship is *not*, followed by the development of a formal model of a Bayesian entrepreneur. After illustrating some broad properties of the model, we dig into the role of priors, the nature of learning and experimentation, and the role of informativeness. We then briefly highlight a connection to two prior bodies of work (by subsets of the authors) in line with this approach, including a Bayesian approach to entrepreneurial strategy and the scientific method to entrepreneurial decision-making.

1.2 What Bayesian Entrepreneurship Is Not

In order to make useful progress for how a Bayesian entrepreneurship approach might usefully inform theory, empirics, and practice in entrepreneurship, it is useful to start with more definitional questions about what exactly Bayesian entrepreneurship means. At one level, the idea that Bayesian logic is related to entrepreneurial decision-making is intuitive and appealing. Certainly, entrepreneurs pursuing opportunity in the face of uncertainty must have some beliefs (perhaps even something we might call "priors") about the probability of success. A central dogma of entrepreneurship is the informational value of learning and experimentation at the earliest stages of new venture formation. To the extent that Bayesian reasoning involves how priors are updated

Downloaded from direct.mit.edu/books/oa-edited-volume/chapter-pdf/2595206/c000600_9780262052160.pdf by Università Bocconi user on 04 November 2026

This is a portion of the eBook at doi:10.7551/mitpress/15918.001.0001

through learning (and how that updating informs practical decision-making), the utilization of a Bayesian approach toward the study of entrepreneurial decision-making seems like a promising approach.

Despite this appeal, it is useful to note that some of the most central approaches and paradigms for studying entrepreneurship either abstract away from the central logical consequences of a Bayesian approach or (often implicitly) impose assumptions that make a Bayesian approach infeasible (or perhaps meaningless). And many approaches that seem, at first blush, to embrace a Bayesian logic (e.g., the lean start-up movement and its focus on learning) nonetheless do not draw out the consequences of an implicitly Bayesian approach in terms of their practical implementation. As such, before turning to a more precise definition of what we mean by Bayesian entrepreneurship, a useful deck-clearing exercise involves first grappling with what Bayesian entrepreneurship is *not*.

1.2.1 Not Knightian Uncertainty

A long tradition in entrepreneurship is premised on the idea that entrepreneurs face a form of fundamental uncertainty—the most extreme type of uncertainty described by Knight (1921)—in which agents are unable to form any meaningful probability distribution over the space of possible outcomes. The essential argument in favor of Knightian uncertainty is that it is impossible for an entrepreneur to forecast all the potential contingencies and circumstances that might arise as they pursue a particular opportunity. If one cannot foresee the many potential contingencies that might arise, the notion of probabilistic reasoning simply does not apply.

In contrast, a Bayesian approach to entrepreneurship operates on a fundamentally different premise. It acknowledges that while absolute certainty is unattainable and the future inherently uncertain, entrepreneurs can still form rational probability estimates based on available information. Bayesian reasoning accepts the premise of uncertainty but posits that uncertainty can be quantified and updated through experience and new information. This perspective would argue against the Knightian view by suggesting that even in the absence of complete information, the perception of entrepreneurial opportunity (which might reflect prior experiences, idiosyncratic market data, or trends) allows entrepreneurs to nonetheless create a subjective probability distribution over the outcomes of interest. These probability distributions are then updated as new evidence becomes available, reflecting a dynamic learning process. Put another way, the Bayesian approach does not deny the existence of uncertainty but rather provides a rigorous framework for managing it. Although Knightian uncertainty underscores the limits of predictability in entrepreneurial ventures,

Downloaded from direct.mit.edu/books/oa-edited-volume/chapter-pdf/2595206/c000600_9780262052160.pdf by Università Bocconi user on 04 November 2026

This is a portion of the eBook at doi:10.7551/mitpress/15918.001.0001

the Bayesian approach embraces these limits as a starting point for systematic improvement in decision-making under uncertainty.

1.2.2 Not Shared Priors

Traditional equilibrium models in economics and finance often operate under the assumption that all agents, including entrepreneurs and investors such as venture capitalists, share a common perspective on the underlying distribution of possible outcomes. This assumption facilitates the analysis of financial contracting and the impact of external variables (e.g., wealth levels) on decision-making processes. Within this framework, there's an implicit consensus on the risks and rewards associated with entrepreneurial ventures, leaving little room for divergent beliefs or the notion that parties might agree to disagree.

The Bayesian approach, by contrast, fundamentally diverges from this assumption by recognizing and accommodating the heterogeneity of priors among different agents. It acknowledges that each decision-maker may hold an idiosyncratic set of priors shaping their individual probability distributions regarding the outcomes of entrepreneurial endeavors. This diversity of priors is seen not as a hurdle to be overcome but rather as a realistic and valuable aspect of the entrepreneurial process.

Allowing for the possibility of agreeing to disagree, especially in the early stages of venture formation and investment, the Bayesian approach offers a more nuanced understanding of how decisions are made under uncertainty. It opens the door to exploring how these varied priors influence negotiation, risk assessment, and the eventual structuring of financial contracts between entrepreneurs and investors. Far from requiring consensus, the Bayesian framework permits a dynamic interaction between differing views, recognizing that such differences are inherent to the process of innovation and value creation.

1.2.3 Not Effectuation

Entrepreneurial effectuation posits a process where opportunities are not discovered but rather created through iterative actions within the entrepreneur's immediate environment (Bhide, 1994; Sarasvathy, 2001). This contrasts sharply with the Bayesian approach, which emphasizes the identification and evaluation of opportunities based on preexisting theories of value creation and capture. In effectuation, the entrepreneur progresses by leveraging available resources and relationships, adapting to feedback, and iterating on their actions, effectively making the opportunity as they go. This method is more about shaping the future through a series of contingent choices rather than predicting it.

Conversely, the Bayesian perspective is grounded in the notion that entrepreneurs start with a specific hypothesis about an opportunity, informed

Downloaded from direct.mit.edu/books/oa-edited-volume/chapter-pdf/2595206/c000600_9780262052160.pdf by Università Bocconi user on 04 November 2026

This is a portion of the eBook at doi:10.7551/mitpress/15918.001.0001

by their prior experiences and beliefs. They then purposefully engage in experimentation to test the validity of their assumptions, adjusting their strategy based on new evidence. This process is analytical and systematic, relying on probabilistic reasoning to refine the entrepreneur's understanding and approach to the opportunity before them. The key distinction here is not just in the method but in the underlying philosophy of opportunity itself. Whereas effectuation sees opportunity as emergent and cocreated with stakeholders, the Bayesian approach views opportunity as something that can be perceived and assessed through rigorous analysis and experimentation.

It is worthwhile noting, however, that although this chapter draws a clear demarcation between these approaches, the formal modeling methodologies outlined here might well form the basis for a future constructive reconciliation of what is, at this stage, different viewpoints in approaching entrepreneurial studies. That is, our expectation is that a richer set of articulations of what is and is not exogenous to entrepreneurial choice may lay bare the truth in both approaches.

1.2.4 Not Simply Biases and Heuristics

A considerable portion of the literature on entrepreneurship has been dedicated to the exploration of biases such as overoptimism and overconfidence, alongside heuristics such as pattern matching, which are thought to significantly influence entrepreneurial decision-making (Åstebro et al., 2014). These biases and heuristics are often portrayed as systematic deviations from rationality that can lead entrepreneurs to make suboptimal decisions. However, the Bayesian approach to entrepreneurship provides a nuanced perspective on these phenomena. It suggests that what is traditionally labeled as biases might, in fact, be the result of varying priors among decision-makers.

Under a Bayesian lens, the apparent overoptimism of entrepreneurs, for instance, may not universally reflect a bias toward unwarranted positivity. Instead, it can be seen as a natural outcome of the unique information sets, or priors, that these individuals possess about their ventures. Entrepreneurs, almost by definition, believe in the potential of their projects more than the average person does. This belief is not merely a bias but a reflection of their unique perspective and information, which informs their probability distributions differently from those of external observers.

Similarly, the use of heuristics, such as pattern matching, often critiqued for its simplicity and potential for error, can be reinterpreted within the Bayesian framework.[1] These heuristics may represent efficient, albeit imperfect, strategies for updating beliefs in the face of new evidence, especially under constraints of time and information.

Downloaded from direct.mit.edu/books/oa-edited-volume/chapter-pdf/2595206/c000600_9780262052160.pdf by Università Bocconi user on 04 November 2026

This is a portion of the eBook at doi:10.7551/mitpress/15918.001.0001

It is important to clarify that the Bayesian approach does not outright refute the existence or impact of biases and heuristics in entrepreneurship. Rather, it offers an alternative interpretation that these phenomena may sometimes reflect rational updating of beliefs based on individual priors and information. This perspective encourages a deeper investigation into the nature of so-called biases and heuristics, suggesting that they may not always signify errors in judgment but could instead be adaptive responses to the unique challenges faced by entrepreneurs.

By reframing the discussion around biases and heuristics within the context of Bayesian reasoning, we open the door to a more sophisticated understanding of entrepreneurial decision-making. This approach not only acknowledges the complexity of the entrepreneurial environment but also emphasizes the role of subjective experiences and information in shaping entrepreneurs' decisions. Through this lens, we can begin to see these decision-makers not as inherently flawed in their reasoning but rather as individuals navigating uncertainty with the tools and information at their disposal.

1.3 What Bayesian Entrepreneurship Is

Of course, simply saying what Bayesian entrepreneurship is *not* does not actually clarify what Bayesian entrepreneurship is. At its core, Bayesian entrepreneurship is a framework that puts the spotlight on the critical interaction between an entrepreneur's prior beliefs and the capacity for learning and experimentation to update those beliefs. This dynamic process allows entrepreneurs to fundamentally formulate and then test theories about value creation and capture and, in so doing, attract resources and capabilities to a venture and allow that venture to choose and implement an overall entrepreneurial strategy.

The essence of Bayesian entrepreneurship lies in its iterative process of belief adjustment. Entrepreneurs set out hypotheses about their business models, market opportunities, and customer preferences and then engage in experimentation to validate or refute these hypotheses. Each piece of evidence collected through these experiments serves to update the entrepreneur's priors, refining their understanding of the venture's potential success and guiding their subsequent decisions.

This dynamic process of updating beliefs in light of new evidence differentiates the Bayesian approach from more static models of decision-making. It recognizes the inherent uncertainty of entrepreneurial ventures and provides a structured method for navigating this uncertainty. By valuing both the initial set of beliefs and the continual learning process, Bayesian entrepreneurship offers

Downloaded from direct.mit.edu/books/oa-edited-volume/chapter-pdf/2595206/c000600_9780262052160.pdf by Università Bocconi user on 04 November 2026

This is a portion of the eBook at doi:10.7551/mitpress/15918.001.0001

a robust framework for understanding how entrepreneurs can more effectively assess and react to the evolving landscape of opportunities and challenges they face.

From a scholarly standpoint, the Bayesian approach blends the sizable streams of research on entrepreneurial strategy as choice (Gans et al., 2019; Agrawal et al., 2021), entrepreneurial theories (Ehrig and Schmidt, 2022; Felin and Zenger, 2009, 2017), entrepreneurship as experimentation (Kerr et al., 2014; Koning et al., 2022; Lindholm-Dahlstrand et al., 2019), and the scientific approach to entrepreneurial decision-making (Camuffo et al., 2020; Camuffo, Gambardella, Messinese et al., 2024; Zellweger and Zenger, 2023).

From a practitioners' standpoint, the Bayesian approach provides an overarching rationale to understand the widespread adoption, often in combination, of methods such as design thinking (Liedtka, 2018), the business model canvas (Osterwalder and Pigneur, 2010), minimum viable products, and A/B testing (the lean start-up method; Eisenmann et al., 2012; Maurya, 2022; Ries, 2011) that explicitly or implicitly advocate for structured processes of formation, testing, and updating of beliefs in developing new ventures.

To clear the field from conflicting interpretations, we clarify that the Bayesian approach does not intend to be a rigid model of entrepreneurial learning. Instead, we see Bayesian entrepreneurship as providing an actionable and possibly teachable "algorithm" (Minniti and Bygrave, 2001) that entrepreneurs can flexibly use to navigate uncertainty; frame, formulate, and solve iterated choice problems; learn by incorporating new information in their beliefs; and ultimately improve decisions concerning their entrepreneurial projects.

Given this end, there are three key tenets, or core claims, that characterize a Bayesian entrepreneurship approach:

1. **Heterogeneous prior beliefs**: Agents do not share a common prior belief about important elements of the state space—that is, they effectively agree to disagree.
2. **Entrepreneurs have stronger prior beliefs than other agents**, which defines the nature of any disagreement with other agents.
3. **Agents update beliefs using Bayes' rule**—that is, as a general statement, uncertainty is not so pervasive that it is impossible to update their beliefs and learn prior to making certain choices. But specifically, updating uses Bayes' rule where possible.

Next, we elaborate on each of these and show that they direct entrepreneurial research toward a set of research questions that are important to both positive and normative theories of entrepreneurship but that, in the absence of an explicit

Downloaded from direct.mit.edu/books/oa-edited-volume/chapter-pdf/2595206/c000600_9780262052160.pdf by Università Bocconi user on 04 November 2026

This is a portion of the eBook at doi:10.7551/mitpress/15918.001.0001

Bayesian entrepreneurship approach as defined by the three tenets, have been understudied in the discipline.

1.3.1 Prior Beliefs Are Heterogeneous

Much of game theory rests on the assumption that it is not possible for rational agents possessing both common knowledge and common priors to disagree. The common prior assumption is that, given the same information, agents would have identical beliefs about the probabilities of states arising. Aumann (1976) proved a famous result that rational agents would not agree to disagree in the sense that they would understand and have common knowledge (or in Geanakoplos and Polemarchakis, 1982, the weaker notion of mutual knowledge) that they had differing priors. If those priors were the result of different information, then rational agents, understanding that the other agent may know something different, will infer each other's information, leading to agreement.

The conclusion of this set of assumptions is not typically borne out in practical reality where persistent disagreement is observed. This includes disagreement in very high-stakes situations, such as those that surround innovation and entrepreneurship. This situation has not been lost on many economic theorists (notably Yildiz, 2000, and Van den Steen, 2001), who have developed ways to move on from the common priors assumption. Indeed, Harsanyi (1967) noted that rational agents with the same information may well assign different subjective probabilities to states. As Van den Steen (2001) notes, what is required is a foundation for the notion that agents might have differing priors while acknowledging and having common knowledge of the fact that they have different priors.

Van den Steen (2001) outlines potential paths forward that we develop below using the notation in table 1.1. His starting point is a game with agents, $i \in \{1, \ldots, N\}$ and a finite state space. For our purposes here, it is useful to consider

Table 1.1
Key notation definitions

Symbol	Definition
i	Agent index, $i \in \{1, \ldots, N\}$
V	Value of a venture conditional on success
C	Cost of exploiting the entrepreneurial opportunity
x	Random variable where $x \in \{0, V\}$
$\tilde{\mu}_i$	Posterior probability of i
μ_i	Prior probability that i holds that value will be V

Downloaded from direct.mit.edu/books/oa-edited-volume/chapter-pdf/2595206/c000600_9780262052160.pdf by Università Bocconi user on 04 November 2026

This is a portion of the eBook at doi:10.7551/mitpress/15918.001.0001

a state space where the value created by an entrepreneurial idea is either 0 (a failure) or $V(>0)$ (a success). A simple approach is to assume that each agent has a prior, $\mu_i \equiv \Pr_i[V]$. A more elaborate setup involves a common prior μ, but agents observe information that is objectively correct with some probability but over which agents have subjective beliefs regarding that information being correct. The difference in priors comes from differences in the informativeness of the signal. These beliefs, $\tilde{\mu}_i$, can then be used as the foundational, but subjective, different priors in a resulting model.[2]

The second approach is cumbersome relative to the first, but Van den Steen (2001, p. 16) notes that it makes it "easier to judge the 'reasonableness' of priors in the analysis." This allows the researcher (i) to link results to psychological biases and behavioral considerations such as overconfidence, (ii) to create a link between beliefs and some fundamental reality that can assist in evaluating efficiency, and (iii) to tackle the issue of differing information versus differing beliefs and, as a result, could allow interpretations that differences are the result of distinct theories about the world rather than the facts per se. This also allows for a role of experimental communication whereby some agent could produce facts that could convince someone else of their position.

This second approach connects with what we describe below as a theory-driven process. That is, although a Bayesian entrepreneurship approach can take heterogeneous priors as a starting point, it also makes explicit an analysis of how priors are formed in the first place. The theory-driven approach specifies a specific class of processes that elaborates on how entrepreneurs build the future state spaces (and prior beliefs) underlying their ideas. Under that approach, prior heterogeneity stems from both differences in the envisioned state space and how entrepreneurs elicit the prior probability distribution.

As an illustration of how this second approach can be formalized, assume two agents, i and j, have the same information. Other things equal, they should form the same prior, μ. However, based on their idiosyncratic background, knowledge, and experience, they build future state spaces and elicit prior probability distribution differently (e.g., maximum entropy vs. indifference rule vs. other methods; Jaynes, 1968). Their prior belief distributions are $\mu_i = \mu + \phi_i$ and $\mu_i = \mu + \phi_j$, with ϕ_i and ϕ_j subjective parameters corresponding to different subjective future state spaces and idiosyncratic ways in which priors are formed. Under this condition, ex ante $\phi_i \neq \phi_j$, and hence priors can be different ("agree to disagree" condition). However, ex post (over time) priors tend to converge through social processes (e.g., communication, persuasion, as described below) so that $\phi_i = \phi_j$ ("common knowledge" condition; Samuelson, 2004).

Downloaded from direct.mit.edu/books/oa-edited-volume/chapter-pdf/2595206/c000600_9780262052160.pdf by Università Bocconi user on 04 November 2026

This is a portion of the eBook at doi:10.7551/mitpress/15918.001.0001

1.3.2 Entrepreneurs Have Relatively Strong Prior Beliefs

The second tenet of Bayesian entrepreneurship concerns the nature of disagreements between agents. Informally, one might claim that if there are two broad states regarding an entrepreneurial idea, success and failure, then the entrepreneur has stronger prior beliefs about the success of the idea than other agents do.

As a term, however, "optimism" has both positive and normative connotations. The normative connotation is that the entrepreneur's beliefs are mistaken relative to a ground truth probability of success that is lower than the entrepreneur's subjective probability. As that ground truth is neither knowable nor observable to the entrepreneur, even ex post, that normative connotation is not a reasonable claim on which to base analysis.

Instead, this tenet of Bayesian entrepreneurship focuses on the positive component of entrepreneurs'optimism—namely, that entrepreneurs have stronger prior beliefs than other agents, which involves a higher likelihood that their project would be successful by whatever criteria one may consider. This is the approach taken by Van den Steen (2005). Thus, we will use the following definition to rank beliefs:

Definition 1.1. Stronger prior. An individual, i, has a stronger prior belief about an idea than individual j, if $\mu_i > \mu_j$.

This definition does not define belief strength per se but rather *relative strength* without reference to the consequences of those beliefs. In applications, an analyst might make additional assumptions or focus on cases where belief differences are consequential. For instance, if by investing C in exploitation an entrepreneur resolves uncertainty about the venture and generates payoffs of V or 0 as the case may be, a case of interest is where $\mu_i > \mu^*$, where $\mu^* \equiv C/V$. In that case, without any additional information, i would be *predisposed* to choose to exploit an opportunity. Belief strength could also be evaluated relative to available experiments—that is, at a level at which even a negative signal resulting in $\tilde{\mu}_i(0)$ would still involve $\tilde{\mu}_i(0)V \geq C$, implying that i has such strong beliefs that no information would dissuade them from pursuing an entrepreneurial opportunity.

The relative definition of belief strength, however, does not require any such thresholds to be met. It only requires that (a) individuals hold heterogeneous priors and (b) they can be ranked using those priors. Thus, it could also be the case that $\mu_i V < C$ but that μ_i is high enough that it is worth gathering information (i.e., engaging in exploration). Interestingly, a single experiment is insufficient to persuade another to pursue the venture, so another will not explore.

Downloaded from direct.mit.edu/books/oa-edited-volume/chapter-pdf/2595206/c000600_9780262052160.pdf by Università Bocconi user on 04 November 2026

This is a portion of the eBook at doi:10.7551/mitpress/15918.001.0001

The second tenet of Bayesian entrepreneurship takes this definition of the relative strength of prior beliefs and claims, according to this ordering, that an *entrepreneur has stronger prior beliefs about their ideas than all (or most) others*. The claim is a focused one. Bayesian entrepreneurship simply requires that the entrepreneur has stronger beliefs about the idea they are pursuing, even if they might be a pessimist about realizations of other uncertain variables.

The claim that entrepreneurs have stronger beliefs is made without reference as to why this is so. One possibility is that this is a trait of entrepreneurs (see, e.g., Dushnitsky, 2010; Fraser and Greene, 2006; Hmieleski and Baron, 2009; Landier and Thesmar, 2008) and can be simply taken as a starting point. Alternatively, it is possible to imagine that this claim is justified by a selection mechanism. For instance, one could ground opportunity recognition and initial action in the context of evolutionary game theory.[3] Justifying this claim more carefully, both theoretically and empirically, is, of course, a research direction that works within Bayesian entrepreneurship, which makes it more salient and compelling.

This claim defines a central challenge for the entrepreneur. To the extent that they do not control key resources necessary to exploit an opportunity, an entrepreneur will have to convince and reduce the level of disagreement with agents who hold weaker prior beliefs. Interestingly, reducing agreement is desirable, but only for instrumental reasons. In a competitive context, entrepreneurs may benefit if others do not share their priors. Weaker beliefs on the part of others imply that they are less likely to pursue an opportunity in competition with the entrepreneur, and entrepreneurs may be concerned about signals that positively update those beliefs of others (Gans, 2023). This would change their incentives regarding the disclosure of information. Indeed, it may well be that this could be the basis of another micro-foundation for the optimistic priors of entrepreneurs, as being optimistic in a relative sense is a precondition for the entrepreneur to find it worthwhile to pursue an opportunity (Robson and Samuelson, 2011).

A key implication of this challenge is that the probability of a venture's success will not just be driven by differences in the preferences or incentives of entrepreneurs. It will also depend on the skills of the entrepreneur (and indeed, as we will see, others) in being able to interpret data, signals, and experiments in a Bayesian manner and choose those experiments understanding data-generating processes and biases. That is, research should focus on improvements in the cognitive process that will allow entrepreneurs to interpret information by taking into account their priors and preferences.

Downloaded from direct.mit.edu/books/oa-edited-volume/chapter-pdf/2595206/c000600_9780262052160.pdf by Università Bocconi user on 04 November 2026

This is a portion of the eBook at doi:10.7551/mitpress/15918.001.0001

1.3.3 Agents Update Beliefs Using Bayes' Rule

The final tenet of Bayesian entrepreneurship concerns how prior beliefs are updated—that is, how agents learn from new information. Consistent with the name for this approach, those beliefs are updated using Bayes' rule to generate a posterior belief from a prior belief. This is a process we call "Bayesian learning." As we will emphasize below, Bayesian learning can be simple or it can be deep. Simple Bayesian learning presumes that posterior beliefs depend only on the prior and not on how the prior is formed. Deep Bayesian learning takes into account how priors are formed and how this impacts what is actually learned. In what follows, we will initially assume Bayesian learning is simple, leaving the consideration of the technically more challenging exposition of deep Bayesian learning until section 1.5.

Given its prominence for Bayesian entrepreneurship, it is useful here to recall how (simple) Bayesian learning actually works. Consider an entrepreneur who has subjective beliefs that their business idea will be successful. The entrepreneur can simply decide to pursue their idea based on what they know or consider collecting data to learn more about it—that is, a signal or experiment.

For Bayesians, a prior probability, simply referred to as a prior, is the probability distribution of a random variable *prior* to the consideration of new evidence or information. It is a subjective probability encapsulating the beliefs an individual might have about a random variable before gathering any data. It is this prior that is updated using new information and, for the purposes of this research program, that updating procedure is given by Bayes' rule. Thus, if μ is a prior probability attached to an event or state, $\tilde{\mu}_i$ is the updated probability that takes into account new information using Bayes' rule. Specifically,

$$\tilde{\mu}_i = \frac{\Pr[\text{Signal}|\mu_i]\mu_i}{\Pr[\text{Signal}]}.$$

The output of this updating procedure, $\tilde{\mu}_i$, is called i's posterior beliefs. The posterior belief is what a decision-maker will use after collecting information to make a decision. The prior beliefs will guide whether it is worth collecting information or, instead, making a decision without such information. For many applications, a full version of Bayesian updating is not necessarily required, and it is robust to some relaxation; see the discussion in Jakobsen (2023).

Bayes' rule summarizes the three main steps of Bayesian learning, which is applicable to entrepreneurs (Zellweger and Zenger, 2023):

Downloaded from direct.mit.edu/books/oa-edited-volume/chapter-pdf/2595206/c000600_9780262052160.pdf by Università Bocconi user on 04 November 2026

This is a portion of the eBook at doi:10.7551/mitpress/15918.001.0001

1. Forming a prior distribution about an entrepreneurial idea
2. Collecting data (e.g., through an experiment) and determining the likelihood function using the information about the parameters available in the data
3. Combining the prior distribution and the likelihood function—using Bayes' theorem—to form the posterior distribution

The posterior distribution reflects the entrepreneur's updated prior, which balances prior knowledge with observed data and is used to conduct inferences about the idea. Bayesian entrepreneurship can be seen as a way to systematize a series of advances that recently occurred in entrepreneurship theory and practice.

As already noted, entrepreneurs and others understand that entrepreneurs have strong priors, so they agree to disagree along those lines. However, this also molds their expectations regarding the informativeness of experiments and other information received. Importantly, all agents expect that new information will reduce disagreement and hence posteriors will converge with new information (Kartik et al., 2021). However, Bayesians believe that their priors are correct, and given a set of possible signals, their own expected posterior equals their own prior.

1.3.4 How Are Priors Formed?

The three tenets of Bayesian entrepreneurship put priors at the centre of analysis. With this focus the question of where those priors come from is raised. This, itself, becomes an important research question, and later in the paper, we will show that different models of prior formation can have different implications for how entrepreneurs gather information and make choices. Given this, here, we note some of the main approaches to the formation of priors that have been undertaken in the literature to date. This is not an exhaustive list but an indicator of where development has taken place and how these give rise to an important research agenda in Bayesian entrepreneurship.

Entrepreneurs' priors form idiosyncratically and are heterogeneous, so that individuals might have different priors about the same idea or opportunity, even starting from the same available information. The idiosyncrasy and heterogeneity of entrepreneurs' prior beliefs stem from a variety of factors.

First, they derive from individual differences rooted in physical and psychological traits, personality, or other individual factors (Baum et al., 2014). These generate differences in cognitive processes (e.g., memory, thought, perception), emotions, and so forth, which result in heterogeneous levels of self-efficacy and achievement drive (Frese and Gielnik, 2014).

Downloaded from direct.mit.edu/books/oa-edited-volume/chapter-pdf/2595206/c000600_9780262052160.pdf by Università Bocconi user on 04 November 2026

This is a portion of the eBook at doi:10.7551/mitpress/15918.001.0001

Second, they derive from differences in entrepreneurs' knowledge, skills, and abilities. Human capital is unevenly distributed across entrepreneurs in type and quantity (Lazear, 2004; Unger et al., 2011). It stems from unique knowledge of a specific matter (Shane, 2000), education (Martin et al., 2013), experience in a particular knowledge domain, or talent (Eesley and Roberts, 2012), and such differences in endowment shape what entrepreneurs explore (exploration space), why and how they do it, and for how long (Gimeno et al., 1997).

Third, they derive from the context in which they operate, as context also contributes to the definition of the individual–opportunity nexus (Shane, 2003). For example, geographic location (Delgado et al., 2010), the ecosystem in which entrepreneurs operate (Acs et al., 2017), their relevant institutions such as universities (Tartari and Stern, 2021), accelerators (Cohen et al., 2019), as well as the access to investors, shape entrepreneurs' prior beliefs. A similar role is played by the social networks entrepreneurs build and are embedded in (Stuart and Sorenson, 2005).

Fourth, they derive from entrepreneurs' preferences (Roach and Sauermann, 2015), interests (Van de Ven et al., 2007), motivation (Baum and Locke, 2004; Guzman et al., 2020), incentives, and goals.

Finally, entrepreneurs form prior beliefs based on their cognitive processes (Baron, 1998), which include intuition (Blume and Covin, 2011), creativity (Amabile, 1997), imagination (Patvardhan and Ramachandran, 2020), as well as logical reasoning and causal thinking (Camuffo, Gambardella and Pignataro, 2024; Felin and Zenger, 2009, 2017). Logical reasoning and causal thinking (Pearl (2000)) are especially important because, through them, entrepreneurs can intentionally and deliberately form priors through theories (Felin and Zenger, 2017; Ehrig and Schmidt, 2022; Agarwal et al., 2025) and hence *choose* (Agrawal et al., 2021; Gans et al., 2019, 2021) which idea is worth pursuing.

1.4 Experiments Under Simple Bayesian Learning

Bayesian entrepreneurship places entrepreneurial processes for decision-making at the center of entrepreneurial research. That focus gives priority to the types of information that are available to entrepreneurs and the purposes to which that information is put. Here, we will follow a convention in information economics that uses the term "experiment" or "signal" as a formal construct for any information that entrepreneurs might gather that itself requires a choice to generate and then utilize. More broadly, information may include information that is passively acquired by entrepreneurs that itself can be used in Bayesian learning. In this section and the next, we focus on the entrepreneur's

Downloaded from direct.mit.edu/books/oa-edited-volume/chapter-pdf/2595206/c000600_9780262052160.pdf by Università Bocconi user on 04 November 2026

This is a portion of the eBook at doi:10.7551/mitpress/15918.001.0001

choices in incurring the costs of generating experiments to understand the underlying value of their idea, as well as strategies that may be employed in its commercialization. This section will assume that priors are given and entrepreneurs use simple Bayesian learning, whereas the next section will consider deep Bayesian learning, which requires the causal structure associated with any experiments or signals to be appropriately understood by the entrepreneur.

1.4.1 How Priors Impact Experimental Interpretation

Priors are the root primitive for considering a Bayesian approach to entrepreneurial decision-making. They matter precisely because entrepreneurs will engage in various activities to de-risk and refine their commercialization plans, which will be a key driver of their potential success. Not only that, these activities will be designed to influence or persuade others changing the set of resources that may be brought under an entrepreneur's control. To that end, here, we consider how posterior beliefs are formed through experimentation.

We derive a fundamental and, we believe, underappreciated result that entrepreneurs' demand for and the form of experimentation they choose will depend critically on their priors. In particular, stronger priors will generate a demand for experiments designed to yield clear signals that the venture should not be pursued while at the same time providing less clear signals that the venture should proceed.

The key starting point for this is to explicitly consider the role of experiments in entrepreneurial decision-making and how Bayesian learning can identify their demand characteristics. Specifically, we want to demonstrate how priors—both of entrepreneurs and others they interact with—drive the value of experiments. We will consider other issues such as the type of experiments chosen below.

Let's return to our ongoing simple example. Recall that an entrepreneur (E) has a prior, μ_E, that a venture idea will be successful, earning them V, versus not successful, earning them 0. That is, the true underlying state space is $x \in \{0, V\}$. If the cost of launching that venture were C, then the expected return from that venture given the prior of μ_i is:

$$R(\mu_i) = \mu_i V - C.$$

If $R(\mu_i)$ exceeds the return of the entrepreneur's alternative opportunity (something we typically normalize to 0 in theoretical models), then, without more information, an entrepreneur might choose to launch the venture. Using our earlier notation, this implies that the entrepreneur will exploit the opportunity, in the absence of further information, if and only if $\mu_i \geq \mu^* \equiv C/V$.

Downloaded from direct.mit.edu/books/oa-edited-volume/chapter-pdf/2595206/c000600_9780262052160.pdf by Università Bocconi user on 04 November 2026

This is a portion of the eBook at doi:10.7551/mitpress/15918.001.0001

Table 1.2
Experiment maps from signal to state, $\Pr[s|\theta]$

	(Good news) $s=1$	(Bad news) $s=0$
(Success) V	λ_1	$1-\lambda_1$
(Failure) 0	$1-\lambda_0$	λ_0

What if there is a signal, s, that takes on a value of 1 with probability λ_1 if the venture will be successful and a value of 0 with probability λ_0 if the venture is unlikely to be successful—that is, $\Pr[s=1|x=V]=\lambda_1$ and $\Pr[s=0|x=0]=\lambda_0$ (see table 1.2). Note here that $\Pr[s=1|\mu_E]=\lambda_1\mu_E+(1-\lambda_0)(1-\mu_E)$ and $\Pr[s=0|\mu_E]=(1-\lambda_1)\mu_E+\lambda_0(1-\mu_E)$. Note also that if the entrepreneur follows the signal in deciding whether to invest, the rate of entrepreneurial implementation (i.e., starts) will be $\Pr[s=1|\mu_E]$, and the observed rate of successful ventures will be $\lambda_1\mu_E$. Thus, observed success will increase with μ_E if the entrepreneur is in control of whether the venture is launched following an experiment.

It is useful to orient ourselves by translating the information in table 1.2 to some familiar terms.

- Sensitivity or true-positive rate: $\lambda_1\mu_i$
- Specificity or true-negative rate: $\lambda_0(1-\mu_i)$
- False-positive probability: $(1-\lambda_1)(1-\mu_i)$
- False-negative probability: $(1-\lambda_0)\mu_i$

Notice, therefore, that an experiment that involves both a higher λ_1 and a higher λ_0 increases both the sensitivity and specificity and reduces the probabilities of false positives and false negatives. Below, we will consider situations where the entrepreneur must choose experiments that trade off λ_1 and λ_0 and so face the option of reducing either false positives or false negatives but only at the expense of each other. It can be seen here that the prior, μ_i, will likely factor into such trade-offs.

1.4.2 How Priors Impact Experimental Value

The above shows how priors impact the information contained in an experiment. Not surprisingly, for this reason, those priors also impact the experiment's value to a decision-maker. We explore that here.

Suppose that it costs $c << C$ to obtain that signal. To calculate whether it is worthwhile to obtain the signal prior to making a decision to launch the venture, we need to calculate the posterior beliefs contingent on "good news" (i.e., $s=1$)

Downloaded from direct.mit.edu/books/oa-edited-volume/chapter-pdf/2595206/c000600_9780262052160.pdf by Università Bocconi user on 04 November 2026

This is a portion of the eBook at doi:10.7551/mitpress/15918.001.0001

or "bad news" (i.e., $s = 0$). Those are:

$$\tilde{\mu}_i(1) = \frac{\lambda_1 \mu_i}{\lambda_1 \mu_i + (1-\lambda_0)(1-\mu_i)}$$

$$\tilde{\mu}_i(0) = \frac{(1-\lambda_1)\mu_i}{(1-\lambda_1)\mu_i + \lambda_0(1-\mu_i)}.$$

Given this, the expected returns, contingent on the signal realizations, are $R(\tilde{\mu}_i(1)) = \tilde{\mu}_i(1)V - C$ and $R(\tilde{\mu}_i(0)) = \tilde{\mu}_i(0)V - C$.

What level of $\{\lambda_1, \lambda_0\}$ implies that a signal would be useful (even if c is arbitrarily low)? A natural interpretation of usefulness would be that good news should lead to a decision to launch the venture—that is, $R(\tilde{\mu}_i(1)) > 0$—and bad news should lead to a decision to abandon the venture—$R(\tilde{\mu}_i(0)) < 0$. Combining these two conditions implies that:

$$\frac{\lambda_1 \mu_i}{\lambda_1 \mu_i + (1-\lambda_0)(1-\mu_i)} V > C > \frac{(1-\lambda_1)\mu_i}{(1-\lambda_1)\mu_i + \lambda_0(1-\mu_i)} V.$$

A necessary condition for this inequality to hold is:

$$\lambda_1((1-\lambda_1)\mu_i + \lambda_0(1-\mu_i)) > (1-\lambda_1)(\lambda_1\mu_i + (1-\lambda_0)(1-\mu_i))$$

$$\Leftrightarrow \lambda_1\lambda_0(1-\mu_i) > (1-\lambda_1)(1-\lambda_0)(1-\mu_i) \Leftrightarrow \lambda_1 + \lambda_0 > 1.$$

Intuitively, the precision of the signal (i.e., the probability the signal is a correct indicator of the underlying state) needs to be sufficiently high. If these conditions hold, then the signal will determine the decision taken.

The important feature to note regarding the necessary conditions for the informativeness of a signal in terms of the decision to be taken is that it only depends on $\lambda_1 + \lambda_0$ and not on prior beliefs, μ_i, or parameters of the decision context. This means that any agent, regardless of beliefs or preferences, will agree that, following good news, $\tilde{\mu}_i(1) \geq \mu_i$ and, following bad news, $\tilde{\mu}_i(0) \leq \mu_i$. This is a standard finding in information economics that the Blackwell ordering of "informativeness" is independent both of prior beliefs and of preferences (Blackwell, 1953).

Prior beliefs come into play when the entrepreneur considers gathering information and whether the information is sufficient to revise their decision.[4] There are two cases of interest where an entrepreneur is "predisposed to exploit" and "not predisposed to exploit." In the predisposed case, $R(\mu_E) > 0$, implying that, without information, the entrepreneur will launch the venture. In this case, the entrepreneur only finds it worthwhile to acquire information if:

$$(\lambda_1\mu_E + (1-\lambda_0)(1-\mu_E))\,(R(\tilde{\mu}_E(1)) - C) - c \geq R(\mu_E)$$

$$\Leftrightarrow \lambda_1\mu_E V - (\lambda_1\mu_E + (1-\lambda_0)(1-\mu_E))C - c \geq \mu_E V - C$$

Downloaded from direct.mit.edu/books/oa-edited-volume/chapter-pdf/2595206/c000600_9780262052160.pdf by Università Bocconi user on 04 November 2026

This is a portion of the eBook at doi:10.7551/mitpress/15918.001.0001

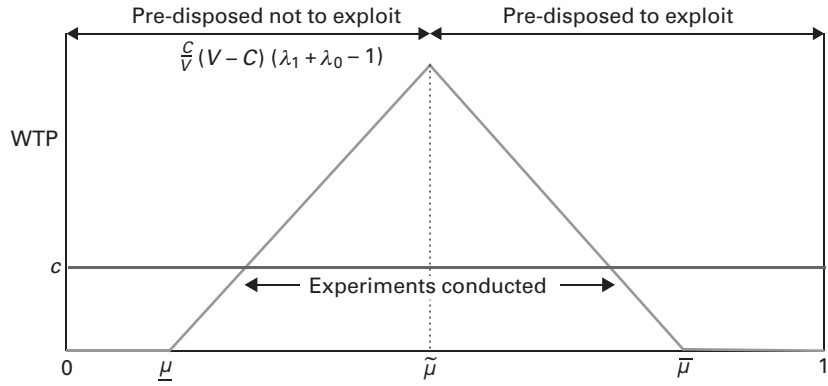


Figure 1.1
Willingness to pay (WTP) for experiments.

$$\Leftrightarrow \bar{\mu}_E(c) \equiv \mu_E \leq \frac{C\lambda_0 - c}{C\lambda_0 + (1-\lambda_1)(V-C)}.$$

Intuitively, if an entrepreneur will launch even if they receive bad news, it is not worth spending to obtain a signal, as it will not be decisive. For the not predisposed case, we have:

$$(\lambda_1\mu_E + (1-\lambda_0)(1-\mu_E))\,(R(\tilde{\mu}_E(1)) - C) - c \geq 0$$

$$\Leftrightarrow \lambda_1\mu_E V - (\lambda_1\mu_E + (1-\lambda_0)(1-\mu_E))C - c \geq 0$$

$$\Leftrightarrow \underline{\mu}_E(c) \equiv \mu_E \geq \frac{c + (1-\lambda_0)C}{(1-\lambda_0)C + \lambda_1(V-C)}.$$

Intuitively, an entrepreneur who is predisposed not to launch can only be persuaded by the signal to launch if their prior is sufficiently strong.

Given this, figure 1.1 shows the willingness to pay for an experiment with parameters $\{\lambda_1, \lambda_0\}$ as a function of the prior, μ.[5] Notice that the willingness to pay peaks at $\mu_E = \tilde{\mu}_E$. At this point, $\tilde{\mu}_E V = C$, and so an uninformed entrepreneur is indifferent between pursuing the venture and not. At this point, the entrepreneur is willing to pay $C/V(V-C)(\lambda_1+\lambda_0-1)$, which is the most for information that will determine whether the venture should go ahead and break the indifference. The experiment cost, c, determines the range of entrepreneurs in terms of their priors who will actually conduct an experiment, $\{\lambda_1, \lambda_0\}$. Note that as the informativeness of the experiment, $\lambda_1 + \lambda_0$, increases, the blue lines shift upwards, and an experiment is more likely to be conducted.

Thus, we see here that although priors do not play a role in determining what action to take in the face of signals that are informative (i.e., where $\lambda_1 + \lambda_0 > 1$),

Downloaded from direct.mit.edu/books/oa-edited-volume/chapter-pdf/2595206/c000600_9780262052160.pdf by Università Bocconi user on 04 November 2026

This is a portion of the eBook at doi:10.7551/mitpress/15918.001.0001

they play a role in determining whether entrepreneurs will gather information before launching a venture. In particular, priors will determine entrepreneurial demand for experiments.

1.4.3 Experimental Choice for Self-Informativeness

One of the tenets of Bayesian entrepreneurship is that entrepreneurs hold strong prior beliefs. However, it is also generally the case that entrepreneurs are resource constrained. That means that they need to secure resources from other people in order to launch and sustain their venture. By definition, those other people do not necessarily share the entrepreneur's beliefs. Thus, in addition to designing experiments to provide information regarding whether further investment in a venture is worth undertaking, the entrepreneur must consider whether experiments are designed in such a way that would persuade others who hold more pessimistic priors to devote resources to the venture.

Before considering the challenge associated with persuading others, it is useful to consider the entrepreneur's choice of experiment when they are seeking to persuade themselves to pursue the venture or not. To see this, we continue the model from section 1.3.1 by endogenizing experimental choice. The experimental choice set is $E \equiv \{\lambda_1, \lambda_0 | \lambda_1 + \lambda_0 \leq \Lambda\}$. Here, Λ parameterizes the scope or size of the experimental choice set. It is assumed that $\Lambda > 1$ so all experiments are informative and $\Lambda < 2$ so that the experimental set does not include the perfect experiment. The experimental choice set is depicted in figure 1.2 and is bounded by the red line. For simplicity, it is assumed here that $c = 0$ and that the entrepreneur is constrained to select one experiment only from the set.

Note that allowing the entrepreneur to choose experiments from this set is a more general model than is commonly assumed in other Bayesian contexts. For instance, Kamenica and Gentzkow (2011) and Bergemann and Morris (2019) consider the case of a sender that communicates information (i.e., the results of an experiment) to a receiver and assume that the set of experiments a receiver may choose from is Bayes plausible. Under this restriction, the expected value of the posteriors equals the prior.[6] In our context, note that the expected value of the posteriors is:

$$\begin{aligned}\mathbb{E}[\tilde{\mu}_i] &= \mu_i \Pr[s = 1|V] + (1 - \mu_i) \Pr[s = 1|0] \\ &= \mu_i \lambda_1 + (1 - \mu_i)(1 - \lambda_0).\end{aligned}$$

This only equals the prior, μ_i, if

$$\frac{\mu_i}{1 - \mu_i} = \frac{1 - \lambda_0}{1 - \lambda_1}$$

Downloaded from direct.mit.edu/books/oa-edited-volume/chapter-pdf/2595206/c000600_9780262052160.pdf by Università Bocconi user on 04 November 2026

This is a portion of the eBook at doi:10.7551/mitpress/15918.001.0001

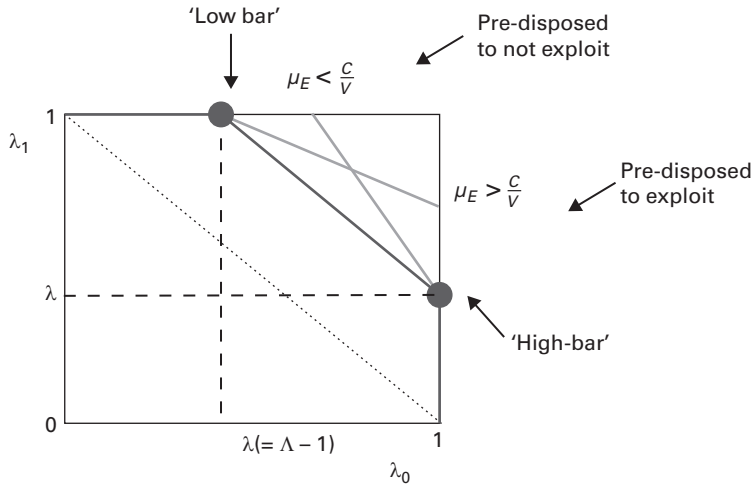


Figure 1.2
Choice of experiment.

that is, the ratio of false positives to false negatives equals the prior odds. Imposing this means that the set of experiments is constrained by the prior itself. Hence, it is not possible to consider how priors change observed experimental choices because that is part of the experimental choice set constraint. However, if this condition does not hold, the signal is biased. As we will now demonstrate, it is possible that entrepreneurial decision-makers might prefer signals that are biased, depending on the level of their priors. Thus, priors will determine the type of experiments employed.

To see this, consider figure 1.2. The blue lines represent two indifference lines for the entrepreneur's expected payoff. That expected payoff is:

$$\Pi \equiv \lambda_1 \mu_E V - (\lambda_1 \mu_E + (1 - \lambda_0)(1 - \mu_E))\, C.$$

Rearranging for experimental space, we have:

$$\lambda_1 = (1 - \lambda_0) \frac{(1 - \mu_E)C}{\mu_E (V - C)}.$$

Note that the marginal rate of substitution between λ_1 and λ_0 is:

$$MRS_{\lambda_1, \lambda_0} = -\frac{(1 - \mu_E)C}{\mu_E (V - C)}.$$

Thus, if $\mu_E V > (<) C$, $MRS_{\lambda_1, \lambda_0} > (<) - 1$.

Downloaded from direct.mit.edu/books/oa-edited-volume/chapter-pdf/2595206/c000600_9780262052160.pdf by Università Bocconi user on 04 November 2026

This is a portion of the eBook at doi:10.7551/mitpress/15918.001.0001

Table 1.3
Extremal experiments

$\{\lambda_1, \lambda_0\}$	**Low bar** $\{1, \Lambda - 1\}$	**High bar** $\{\Lambda - 1, 1\}$
Name	"Best Foot Forward"	"Raising the Bar"
Description	Easy to pass	Hard to pass
Bias	No false negatives	No false positives
Start-up formation rate	$(\Lambda - 1)\mu$	$\mu + (2 - \Lambda)(1 - \mu)$
Start-up failure rate	$(2 - \Lambda)(1 - \mu)$	0
Example	Target lead users (Moore, 2014)	Minimum viable product (Ries, 2011)

Note that, depending upon whether the entrepreneur is predisposed to exploit or not in the absence of an experiment, the highest indifference curve will be tangent with the experimental set at opposite and extreme points. When the entrepreneur is predisposed to exploit, they choose $\{\lambda_1, \lambda_0\} = \{1, \lambda\}$ (where $\lambda = \Lambda - 1$). Gans (2023) terms this a "low-bar" experiment. This is because, for this experiment, it is very unlikely to generate bad news, but if it does, that news sends a clear (indeed perfect) signal that it is not worthwhile to exploit the opportunity. Building a high-quality product to establish a beachhead is a strategy emphasized by Moore (2014). The idea is to target customer segments where you believe you can gain traction among lead users. This is a low-bar experiment because if you build a sufficiently strong product, if that product does not succeed in that segment, this is a strong signal that it will not succeed more broadly (see table 1.3 for a characterization).

By contrast, when the entrepreneur is predisposed not to exploit, they choose $\{\lambda_1, \lambda_0\} = \{\lambda, 1\}$. This represents a high bar because, for this experiment, it is very unlikely to generate good news, but if it does, that news sends a clear (indeed perfect) signal that it is worthwhile to exploit the opportunity. An example of a high-bar experiment is a minimum viable product because it is launched in a way that is stacked against success (Ries, 2011).

As noted already, entrepreneurs are presumed to have a higher prior, μ_E. Thus, we have demonstrated that:

Proposition 1.1. An entrepreneur for whom $\mu_E > C/V$ has an optimal experiment of $\{\lambda_1, \lambda_0\} = \{1, \Lambda - 1\}$.

The proof is in Gans (2023). The result that decision-makers should "bias" their information signals toward their priors was most clearly demonstrated in Che and Mierendorff (2019). Their application was a dynamic process; however, it was considered biased in how media outlets presented information to

Downloaded from direct.mit.edu/books/oa-edited-volume/chapter-pdf/2595206/c000600_9780262052160.pdf by Università Bocconi user on 04 November 2026

This is a portion of the eBook at doi:10.7551/mitpress/15918.001.0001

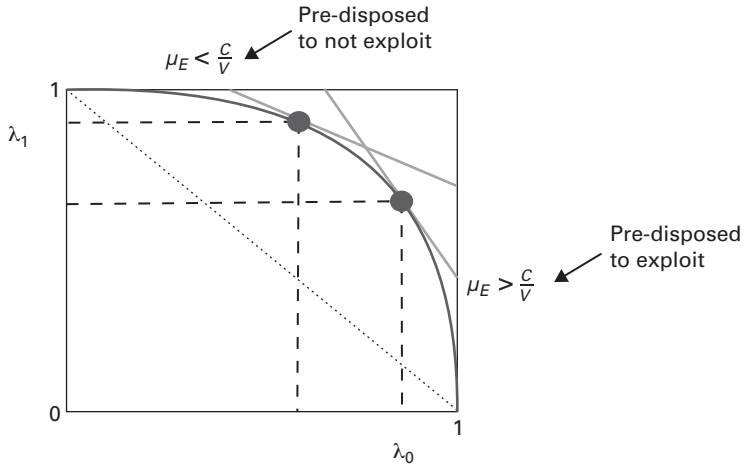


Figure 1.3
Strictly concave experimental frontier.

users. Antecedents also focusing on media and politics while not deriving clear own-bias results include Calvert (1985), Suen (2004), Burke (2008), and Damiano et al. (2020). In all of those cases, however, the experimental choice set was binary and confined to extremal choices. Here, it is demonstrated that these are the optimal choices across all of the suitably constructed experimental choices.

Finally, the construction of the experimental choice set in the example analyzed here is somewhat arbitrary. In particular, $\Lambda(\lambda_1, \lambda_0)$ could be a concave boundary, as depicted in figure 1.3. Notice that the optimal experiment is selected where the slope of the indifference curve, which depends on μ_E, is tangent to the $MRS_{\lambda_1,\lambda_0}$. In this situation, extremal experiments are not chosen, but it can be seen that the same trade-off in terms of the bias of the optimal experiment and its relationship to μ_E remains.

1.4.4 Experimental Choice to Inform Others

We now turn to consider the experimental choice when an entrepreneur has to convince another agent to support the project. The typical situation would involve a venture capitalist supplying C to the venture.

The issue of contrasting self-informative versus other-informative experimental choice is explored by Gans (2022). To consider this, let μ_O be the prior of others who are resource holders or, as we will refer to them, investors. As per our founding assumption, it is assumed that $\mu_E \geq \mu_O$. Given this, we can now define two experimental benchmarks:

Downloaded from direct.mit.edu/books/oa-edited-volume/chapter-pdf/2595206/c000600_9780262052160.pdf by Università Bocconi user on 04 November 2026

This is a portion of the eBook at doi:10.7551/mitpress/15918.001.0001

- Entrepreneurial-optimal experiment: $\arg\max_{\{\lambda_1,\lambda_0\}\in E} \lambda_1\mu_E V - (\lambda_1\mu_E + (1-\lambda_0)(1-\mu_E))C$
- Investor-optimal experiment: $\arg\max_{\{\lambda_1,\lambda_0\}\in E} \lambda_1\mu_O V - (\lambda_1\mu_O + (1-\lambda_0)(1-\mu_O))C$

These optimization problems are linear in λ_1 after substituting in the constraint $\lambda_0 = \Lambda - \lambda_1$. Therefore, the entrepreneurial-optimal experiment involves $\lambda_1 = 1$ if $\mu_E \geq C/V$ whereas the investor-optimal experiment involves $\lambda_1 = 1$ if $\mu_O \geq C/V$. Thus, these choices will coincide if, absent information, both the entrepreneur and others will choose to invest or not choose to invest. By contrast, if $\mu_E > C/V > \mu_O$, the entrepreneur will favor an experiment with $\lambda_1 = 1$, whereas others will favor an experiment with $\lambda_0 = 1$.[7]

This conflict can be resolved in several ways. One way would be to run more than one experiment. This, however, would involve additional experimental costs, c, that may be prohibitive, given the marginal value of a second experiment is less than the value of the first one alone. The other way this is solved is by considering the equity position of investors in the venture. Suppose that the entrepreneur gives the others an equity share of any resulting value of $1-\alpha$ (where $\alpha \in [0,1]$) in return for their contribution of C. For simplicity, it will be assumed that the entrepreneur can finance the experiment cost, c, out of pocket and will determine equity levels after the experiment.[8] Finally, suppose that there is a competitive investment market, which implies that as long as investors expect to earn at least 0, they will accept the deal.

The following proposition demonstrates that the entrepreneur is better off offering the investor-optimal experiment.

Proposition 1.2. Suppose that $\mu_E > C/V > \mu_O$. The entrepreneur chooses an experiment $\{\lambda_1, \lambda_0\} = \{\Lambda - 1, 1\}$.

The proof involves a simple comparison of conditions.[9] Intuitively, although the entrepreneur prefers the entrepreneur-optimal experiment, the amount by which they can reduce the equity ceded to the others is greater if they adopt the investor-optimal experiment. The second factor outweighs the first.

It is instructive to illustrate how this impacts the rate of exploitation as well as the rate of observed entrepreneurial success. To examine this, we need to make an assumption regarding what the true prior, $\hat{\mu}$, of success is. There are two interesting cases: (i) $\hat{\mu} = \mu_E$ or (ii) $\hat{\mu} = \mu_O$.

1. If the entrepreneur is correct ($\hat{\mu} = \mu_E$), then following Proposition 1.2, the venture is launched with probability $\mu_E + (2-\Lambda)(1-\mu_E)$, and it never fails.

Downloaded from direct.mit.edu/books/oa-edited-volume/chapter-pdf/2595206/c000600_9780262052160.pdf by Università Bocconi user on 04 November 2026

This is a portion of the eBook at doi:10.7551/mitpress/15918.001.0001

2. If the venture capitalist is correct ($\hat{\mu} = \mu_O$), then following Proposition 1.2, the venture is launched with probability $\mu_O + (2 - \Lambda)(1 - \mu_O)$, and it never fails.

By contrast, if the entrepreneur's optimal experiment is selected, then:

1. If the entrepreneur is correct ($\hat{\mu} = \mu_E$), then following Proposition 1.1, the venture is launched with probability $(\Lambda - 1)\mu_E$, and it fails with probability $(2 - \Lambda)(1 - \mu_E)$.
2. If the venture capitalist is correct ($\hat{\mu} = \mu_O$), then following Proposition 1.1, the venture is launched with probability $(\Lambda - 1)\mu_O$, and it fails with probability $(2 - \Lambda)(1 - \mu_O)$.

This shows that when the entrepreneur is free to launch the ventures without having to convince others, ventures are more likely to be launched and more likely to fail compared with the case where other-informativeness is required. This is a testable implication of this framework.

This illustrates just one way in which experiment design might be used in a communication context. Gans et al. (2019) and Agrawal et al. (2021) explore others in which experiments might be altered by changing tested strategies alongside a fixed core idea when both are subject to uncertainty. In this case, there is a rationale for continued experimentation following some signal realizations. Specifically, Agrawal et al. demonstrate that an experiment that tests strategies with a lower expected prior first in a sequence can provide a better signal of the underlying idea as opposed to choosing the strategy with the highest prior.

1.4.5 A Comment on Bayesian Persuasion

How do you get someone to do something that is in your interests but not necessarily in the immediate interest of that person? One way is to provide incentives. Another is to make investments that reduce the cost that a person faces in doing what you want them to do. A final way is to persuade them. Persuasion involves providing information that causes people to change their decisions. Such persuasion is Bayesian when the process by which those people use that information is to update their priors according to Bayes' rule.

Entrepreneurs face challenges in providing incentives and changing costs. That is why persuasion can potentially be a means of getting others to take action to improve the venture's prospects. Above, we demonstrate how a clearly defined experiment can be designed to be persuasive. In economics, a literature nominally termed "Bayesian persuasion" has arisen that deals with the use of

Downloaded from direct.mit.edu/books/oa-edited-volume/chapter-pdf/2595206/c000600_9780262052160.pdf by Università Bocconi user on 04 November 2026

This is a portion of the eBook at doi:10.7551/mitpress/15918.001.0001

information to change decisions when the provider of that information may have incentives to selectively choose how information is generated and transmitted. In other words, this literature deals with situations when the receiver of such information does not know the underlying data-generating process of the signals they are receiving. Unlike older literatures that deal with cheap talk, verification opportunities, or plain-old signaling, this newer literature allows the sender far more commitment power.

The new literature was founded by Kamenica and Gentzkow (2011); (see Kamenica, 2019, for a review).[10] The basic model assumes that there is a receiver (e.g., a venture capital investor) who has a payoff function $u(a, x)$, which is a function of the action they take, a (say, invest or not invest), and a state of the world, x (say, the success of the venture). The sender (or entrepreneur) has their own payoff, $\pi(a, x)$. The sender chooses an experiment from a set, E. The receiver knows the structure of the experiment. Then, the experiment is run, generating signals that, in turn, influence the action the receiver takes. It is by knowing the structure of the experiment that the receiver is able to use Bayes' rule to update their prior.

The focus then is on what type of experiment the sender chooses.[11] In the entire Bayesian persuasion literature, it is always assumed that the experimental set is Bayes plausible, meaning, where

$$\frac{\mu}{1-\mu} = \frac{1-\lambda_0}{1-\lambda_1}.$$

As already noted, this does not leave much room for a consideration of priors. Expanding this to the unconstrained set, E, we already saw a version of that above when the entrepreneur chose the investor-optimal experiment in order to influence their own payoff, including the amount of equity the entrepreneur had to cede. However, the Bayesian persuasion literature is about more complicated experiments that, for instance, are designed to gather some types of evidence and not others. But a key feature is that it is the receiver's posterior beliefs (which, as we know, depend on the receiver's prior) that are the focus of the sender's attention. The contribution of the literature is to provide ways of computing the optimal persuasive experiment in various circumstances.

Another research line that follows from this is: What happens when some agents do not know the precise data-generating structure of an experiment?[12] How does this impact other-informativeness and also on incentives to distort persuasive signals? The work of Andrews and Shapiro (2021) makes a start on this interesting question.[13]

Downloaded from direct.mit.edu/books/oa-edited-volume/chapter-pdf/2595206/c000600_9780262052160.pdf by Università Bocconi user on 04 November 2026

This is a portion of the eBook at doi:10.7551/mitpress/15918.001.0001

1.5 Priors and Experiments Under Deep Bayesian Learning

Under simple Bayesian learning, priors are given without any consideration of how they are formed. Moreover, the choice of experiments focuses solely on the stochastic properties of experiments rather than on their precise interpretation. This represents a discord between formal theory and the underlying descriptive structure of that theory. That is, a particular assumption on a data-generating process might motivate the levels of λ_1 and λ_0, but that assumption is embedded solely in those parameters. Moreover, the experiments were aimed at resolving some uncertainty about the likelihood of particular states using a given distribution where choices between experiments might alter the precision of signals in certain directions. In reality, experiments might inform about particular processes and also about the data-generating process—that is, entrepreneurial experiments may cause entrepreneurs to revisit and alter their theories about the world. This section examines what happens when we model entrepreneur prior formation in a deep rather than simple fashion.

1.5.1 Theory-Based Prior Formation

To set the stage, consider an entrepreneurial venture where V will be generated in profit if two events occur based on two attributes that are modeled as the realization of two random variables, $x_t \in \{S, W\}$ and $x_m \in \{Y, N\}$. Specifically, profit is V only if $x_m = Y$. x_t is a random variable that captures the strength of a technological opportunity as strong (S) or weak (W), whereas x_m is a random variable that captures the existence of a particular new market opportunity as existing (Y) or not (N). The total state space is defined by the Cartesian product of these individual attributes. Here, $X = x_t \times x_m$ or:

$$X = \{(S, Y), (S, N), (W, Y), (W, N)\}.$$

x_t represents the logical antecedent that drives the prior of interest. It can be a predictor or truly a "cause," i.e., a hypothetical intervention of the entrepreneur (Pearl and Mackenzie, 2018). for which the entrepreneur sets the parameter. Differently from the other attributes, in this case conditional probabilities unambiguously reflect causality and not correlations. Given this, the probability that the venture yields a profit is:

$$\Pr[Y] = \Pr[S, Y] + \Pr[W, Y]. \tag{1.1}$$

Because the state space is *subjective*, the probabilities of the four states are also subjective and, potentially, rely on sparse data. In the appendix, we operationalize our approach using Dirichlet probability distributions following Camuffo,

Downloaded from direct.mit.edu/books/oa-edited-volume/chapter-pdf/2595206/c000600_9780262052160.pdf by Università Bocconi user on 04 November 2026

This is a portion of the eBook at doi:10.7551/mitpress/15918.001.0001

Gambardella and Pignataro (2024). Entrepreneurs complement their sparse information about the state space with "pseudo-observations," which they build in their head "bootstrapping" form abductive reasoning or using analogies. Hence, they fit preliminary probability distributions on X stemming from synthetic empirical frequencies that occur in worlds that decision makers envision in their mind.

A simple approach to Bayesian learning would focus on the prior of interest—namely, $\mu = \Pr[Y]$. However, here we have posited another attribute, x_t, that potentially drives the level of that prior being its logical antecedent. A deep approach to Bayesian learning would take that into account and ensure that any updating that occurs as the result of experiments incorporates the structure of the relationship between x_m and x_t.

It is at this point that there is a role for theorizing on the part of the entrepreneur. This is a process that respects the potentially sparse nature of information to assist in informing a prior as well as the role that theories and the confidence with which they are held play a role in the strength of those priors. When entrepreneurs build priors using causal reasoning and thinking, they strengthen their beliefs using logical arguments (Carroll and Sorensen, 2021; Ehrig and Schmidt, 2022; Felin and Zenger, 2017). They hold strong beliefs about the value of their venture in the sense of being "logically persuaded" that their understanding of how value will be created is plausible.

The challenge facing entrepreneurs is not simply to learn about not simply probabilities but also states. In entrepreneurial environments, both are potentially unknown. For example, an entrepreneur might think that the development of a *new* technology will generate a *new* market. In order to believe in this idea, they need to build the future state space and form a prior belief (subjective probability distribution) about it. We follow Camuffo, Gambardella and Pignataro (2024) to illustrate how prior beliefs can be formed as theories of value and then tested and updated through experiments.

To define the future state space, entrepreneurs start by identifying *attributes* of the problem, which are elements of the future state space with uncertain realizations. Among the many attributes that they can choose, entrepreneurs focus on the attributes that they believe affect the likelihood of occurrence of a state of interest, such as the existence of a new market. The theory-based approach requires that entrepreneurs explicitly consider the identification of attributes and causal relationships among them as a forward-looking process. Entrepreneurs envision future states and attach subjective probabilities (beliefs) to them. In our example, the entrepreneur wonders whether a new market will exist and thinks that if they can develop an appropriate technology, the rise of this market will be more likely.

Downloaded from direct.mit.edu/books/oa-edited-volume/chapter-pdf/2595206/c000600_9780262052160.pdf by Università Bocconi user on 04 November 2026

This is a portion of the eBook at doi:10.7551/mitpress/15918.001.0001

Entrepreneurs' uncertainty derives not only from the randomness of their subjective probabilities but also from the uncertainty about the parameters of their probability distributions. A specific set of parameters, θ, identifies one distribution and therefore, in our example, one expected probability of new market rise, $\nu(\theta)$. But entrepreneurs do not know the probability distribution of the parameters either. Thus, they also form, cognitively, a subjective probability distribution, $\mu(\theta)$, of the parameters. This defines a family of probability distributions (see appendix A for formal derivation) identified by the different θ in the space of all the possible realizations of the parameters in the example.

Causal reasoning or theories (Karni, 2022) allow us to select a set of probability distributions on which entrepreneurs can concentrate their beliefs. In our example, the entrepreneur believes that the availability of the new technology increases the probability of the emergence of a new market. This logical causal chain can be represented as a Bayesian network (Pearl, 2000), which, in our simple example, is:

$$X_t \longrightarrow X_m$$

and $\Pr[Y|S] > \Pr[Y|W]$. This is the entrepreneur's *theory*, which we define as a restriction of the parameters of the probability distribution or, equivalently, as a restriction of the set of probability distributions that the entrepreneur considers for their future state of interest, $\Pr[Y]$.[14]

Priors and expected values of theories Using Bayes' rule, we can decompose equation 1.1 to find the expected prior of interest, $\mathbb{E}\Pr[Y]$ as:

$$\mathbb{E}[\Pr[Y]] = \mathbb{E}[\Pr[Y|S]\Pr[S]] + \mathbb{E}[\Pr[Y|W](1-\Pr[S])] \tag{1.2}$$

$$= \mathbb{E}\Pr[Y|W] + \mathbb{E}[\Pr[S](\Pr[Y|S] - \Pr[Y|W])]. \tag{1.3}$$

We can use this to posit various theories that capture the underlying parameters of the distribution. For instance, consistent with causal reasoning expressed above, let Θ denote the theory that $\Pr[Y|S] > \Pr[Y|W]$—that is, the market opportunity is more likely to exist if the technological opportunity is strong. A natural alternative theory or null hypothesis is that there is no relationship between the technological opportunity and the existence of the market. Let Θ' denote that alternative, whereby $\Pr[Y|S] = \Pr[Y|W]$. Each theory has an expected value, which is the expected value of the future states of interest under the theory:[15]

$$V_\Theta \equiv \left(\mathbb{E}_\Theta[\Pr[Y|W]] + \mathbb{E}_\Theta[\Pr[S](\Pr[Y|S] - \Pr[Y|W])]\right)V > V_{\Theta'}$$

$$\equiv \mathbb{E}_{\Theta'}[\Pr[Y|W]]V.$$

Downloaded from direct.mit.edu/books/oa-edited-volume/chapter-pdf/2595206/c000600_9780262052160.pdf by Università Bocconi user on 04 November 2026

This is a portion of the eBook at doi:10.7551/mitpress/15918.001.0001

Under the theory, the parameters outside Θ have probability 0. Noteworthy, the more concentrated the parameter set on parameters that imply high expected probability, the higher the subjective probability of occurrence of the future states of interest (in our example, the rise of a new market). This, given $\mathbb{E}_{\Theta}[\Pr[Y]]$, makes the entrepreneur's prior stronger because, thanks to their theory, they are persuaded that the states of interest are more likely to occur. This depends on their reasoning—that is, on a choice of more powerful attributes or causal links (a "better Bayesian network") that increases, to a greater extent, the likelihood of occurrence of the states of interest.[16] In this way, entrepreneurs build, through logical thinking and causal reasoning, prior beliefs about an idea or opportunity.

In our example, the potential development of the new technology raises the probability of the emergence of a new market in the entrepreneur's mind. This theory is subjective and idiosyncratic. It determines the entrepreneur's confidence or optimism, leading them to believe what others might not. This process contributes to explaining why, given individual characteristics and contextual factors, priors across entrepreneurs might be heterogeneous, even about the same opportunity and with the same information available.

Noteworthy, the presence of a null hypothesis is important, as it allows the causal nature of theories to be fully captured.

In our example, special cases of the theory $\Pr[Y|S] > \Pr[Y|W]$ are: (a) if $x_t = S$ happens, then $x_m = Y$ happens—that is, $x_t = S$ is a *sufficient condition* for $x_m = Y$; and (b) if $x_t = S$ does *not* happen, then $x_m = Y$ does *not* happen—that is, $x_t = S$ is a *necessary condition* for $x_m = Y$. In the former case, entrepreneurs theorize that a strong technology implies the emergence of the market—that is, $\Pr[Y|S] = 1$—and that other attributes may also generate high demand—that is, $\Pr[Y|W] > 0$. In the latter case, they recognize that $x_m = Y$ cannot occur without $x_t = S$—that is, $\Pr[Y|W] = 0$—but may not occur, even if $x_t = S$—that is, $\Pr[Y|S] < 1$—which means that other attributes can increase the probability of $x_m = Y$. An even more special case is the one in which $x_t = S$ is a necessary and sufficient condition for $x_m = Y$—that is, $x_m = Y$ happens if and only if $x_t = S$ happens. In this case, entrepreneurs theorize that a strong technology implies the emergence of the market and that the new market will emerge only in presence of a strong technology—that is, $\Pr[Y|S] = 1$—and that no other attributes may also make the new market emerge—that is, $\Pr[Y|W] = 0$.

Null hypotheses and alternative theories represent the way in which entrepreneurs incorporate counterfactual thinking (Pearl and Mackenzie, 2018) in the prior formation process.

Downloaded from direct.mit.edu/books/oa-edited-volume/chapter-pdf/2595206/c000600_9780262052160.pdf by Università Bocconi user on 04 November 2026

This is a portion of the eBook at doi:10.7551/mitpress/15918.001.0001

Priors on priors (or theories) Entrepreneurs might be aware that the future state space they build their prior beliefs on might not be true. This equates to acknowledging that there are low-probability states they do not consider or unforeseen contingencies their prior beliefs do not account for.

In this case, entrepreneurs are aware of being unaware (Karni and Vierø, 2017) and might have "priors on priors," which are priors on whether their theories are true. We represent this prior on prior as a probability, $\omega \in (0, 1)$. Priors on priors imply that entrepreneurs define a null hypothesis, as we have done with Θ' above. This is a different configuration of the parameters against which entrepreneurs compare their theories. In our example, this corresponds to the probability that the new market will rise absent the entrepreneur's attribute and causal link or, equivalently, to the probability the new market will rise under all other imaginable alternative theories.

The expected value of the theory is then:

$$V = \omega V_{\Theta} + (1 - \omega) V_{\tilde{\Theta}}, \tag{1.4}$$

which takes into account that the entrepreneur is not sure whether their theory—that is, the Bayesian network of their theory—is the right one (i.e., the entrepreneur has methodic doubt). This is the prior belief formed by an entrepreneur who is also aware that their theory might be false.[17]

1.5.2 Experimental Choice and Design

Under simple Bayesian learning, experiments are easy to conceive: they are signals with a known structure that inform the entrepreneur—and others—about the true state. Deep Bayesian learning is different. It does more than update prior distributions; it can prompt a reformulation or even a rejection of the very theories that produced those priors. Crucially, experiments vary in their capacity to do this. As we will show, experiments can be designed to distinguish among alternative parameterizations of a given distribution, but the testing of theories themselves typically occurs only within a theory—usually with the goal of refuting that theory—rather than by enabling clean comparisons across competing theories. Nevertheless, we argue that experimentation conducted within a theory can still help entrepreneurs choose which theory to rely on when making key decisions.

Testing within a theory As already noted, a theory involves not only distributional assumptions on a state space but also assumptions regarding the states under consideration and the particular causal structure between states. In addition, an entrepreneur does not simply hold a theory but also has methodic doubt and so will consider there is some probability, $1 - \omega$, that the particular set of

Downloaded from direct.mit.edu/books/oa-edited-volume/chapter-pdf/2595206/c000600_9780262052160.pdf by Università Bocconi user on 04 November 2026

This is a portion of the eBook at doi:10.7551/mitpress/15918.001.0001

parametric restrictions do not hold and an alternative hypothesis (i.e. other sets of attributes and causal links) is true. Thus, when the entrepreneur conducts experiments within the theory, it can inform on the distribution, state space, and also parametric restrictions.

According to Ortoleva (2012), when decision-makers face *drastically* contradicting evidence that disrupts dynamic coherence, which is a set of coherent beliefs that they hold, they will update their prior on prior—that is, the prior on the future state space or decision problem they are working on. These coherent beliefs are analogous to the above-defined entrepreneurial theories in that Ortoleva has in mind logical links contradicted (or possibly supported) by evidence. In our framework, this implies that when the distance between the two distributions is higher than some threshold, entrepreneurs also change their prior on prior, ω, to, say, ω'. Intuitively, this says that if an experiment does not change expectations about parameters radically, entrepreneurs only change their distributions of the parameters. However, beyond the threshold, not only do V_{Θ}, $V_{\Theta'}$, and V change but also entrepreneurs update ω. An entrepreneur who gains extremely negative evidence (beyond a threshold they have in mind) will not only adjust their prior but also question the state space at hand. They will update negatively their prior on prior (i.e., prior ω on Θ).

When this occurs, entrepreneurs do not simply update, based on experimental evidence, their beliefs on the idea but also update their beliefs on the underlying conceptual structure (i.e. their theory. This means that the experiment has either raised the probability of states that were previously neglected, leading the entrepreneur to consider them in their theory or provided evidence about the existence of new attributes, i.e. of states originally unforeseen by the entrepreneur.

Such "upper level" updating can occur in two ways. The first, defined as "reverse Bayesianism" (Karni and Vierø, 2013, 2017), posits that as entrepreneurs obtain awareness that unforeseen contingencies are possible, they will shift probabilities of foreseen contingencies proportionally to them. Specifically, entrepreneurs change their prior beliefs over (i) new states that may emerge when they obtain awareness about previously unforeseen actions or consequences and (ii) previously null states that had been previously considered as not possible. Reverse Bayesianism enables entrepreneurs to adjust the state space they envision (their theory/opportunity) for unknown events. This process is consistent with the methodic doubt of scientists, who acknowledge there might be other explanations/theories to what they study, which they do not know (Camuffo et al., 2023). We can think of reverse Bayesianism as forward-looking Bayesianism.

Downloaded from direct.mit.edu/books/oa-edited-volume/chapter-pdf/2595206/c000600_9780262052160.pdf by Università Bocconi user on 04 November 2026

This is a portion of the eBook at doi:10.7551/mitpress/15918.001.0001

The second, defined as the "hypothesis-testing model" (Ortoleva, 2012), posits that entrepreneurs use standard Bayesian updating but might have non-Bayesian reactions to low-probability events or unforeseen contingencies. Under the hypothesis-testing model, entrepreneurs follow Bayes' rule if they receive signals to which they assign a probability above a threshold. Otherwise, they look at a prior over priors and update these second-level priors using Bayes' rule. Such updating of second-level priors corresponds to forming, testing, and updating beliefs over theories (Karni, 2022) or future state spaces and corresponds to the process underlying the changes in an entrepreneur's confidence that their theory/future state space is true. As a result, Bayesian entrepreneurs form, test, and update their beliefs about their ideas and the generating mechanisms underlying such ideas and their choices.

After the experiment, the expected value under the theory (the prior belief on the opportunity or idea) changes to values V'_{Θ} and V'_{Θ} based on updated distributional assumptions. The conditional expectation V'_{Θ} can be larger or smaller than V_{Θ}, depending on whether the information from the experiment upholds or contradicts prior beliefs. Of course, there will be also an updated expected value under the null hypothesis, $V'_{\Theta'}$, developed analogously to $V_{\Theta'}$. Following the experiment, the expected value of the theory becomes:

$$V' = \omega' V'_{\Theta} + (1 - \omega') V'_{\Theta'}, \tag{1.5}$$

where the update ω' occurs only if the evidence is sufficiently contradictory to question the current future state space.

Testing when there are competing theories There is a subtlety that arises when testing in the context of holding a theory—that is, when entrepreneurs hold $\omega = 1$. If entrepreneurs use all the information they have to form their prior beliefs, then before running an experiment on any subset of parameters, θ, of their theory, their expected posterior (updated probability distribution) on a theory is equal to the current prior belief—specifically, $\mathbb{E}[\omega'] = \omega$. Of course, after the experiment, when they observe its outcome, they make a positive or negative update. But before the experiment, they weigh positive or negative updates in such a way that the expected updated probability distribution and the expected updated prior on prior remain the same as their current beliefs.

If this is not the case, entrepreneurs must have, before the experiment, information that puts greater weight on favorable or unfavorable information about the theory. However, in this case, their current expected probability distribution and prior must be different from their current beliefs. In other words, any information available to the entrepreneur before the experiment will be incorporated in the current probability distributions of the theory and in the prior belief on

Downloaded from direct.mit.edu/books/oa-edited-volume/chapter-pdf/2595206/c000600_9780262052160.pdf by Università Bocconi user on 04 November 2026

This is a portion of the eBook at doi:10.7551/mitpress/15918.001.0001

it. In turn, this implies that, using equation 1.4, the expected update of V is V. Therefore, if the experiment is even minimally costly, it is not worth running it because it does not provide, in expectation, additional information.[18]

This suggests that testing within a theory is, at best, useful for understanding the limitations of that theory itself. It does not create information that would necessarily allow for updating of competing or distinct theories. In part, this is a result of the way we have set up the problem; for instance, two theories may be based on common assumptions about the value of a technological opportunity, and so information collected regarding one may inform others (see the discussion in Gans et al., 2019). But it is also a conclusion that changes when the entrepreneur has engaged in the cognitive processes to construct more than one competing theory. As Gans et al. (2024) argue (i.e., "test two, choose one"), there is value to entrepreneurs considering two distinct choices or plans—here, two alternative theories—precisely so that they compare them. The second theory, in a sense, comes with an alternative set of actions that can be taken, the expected value of which is an outside option when evaluating the initial theory.

Thus, experiments, even if informing only of the value of a specific theory, can be valuable when alternative theories have been articulated.[19] An alternative theory comprises a different state space characterized by different attributes or causal links. Using the relevant parameters and probabilities, the entrepreneur calculates V_2, the unconditional expected value of the alternative theory, in the same way as V, which, for clarity, we call V_1. If entrepreneurs have an alternative theory with expected value V_2, then important changes in V_1—such as those induced by updating priors or priors on priors—are more likely to imply $V_1' < V_2$ and thus a switch to the alternative theory V_2. Experiments that generate unexpected (surprising) outcomes—for instance, observations that represent unforeseen contingencies (Karni and Vierø, 2013, 2017), or regarding low-probabilities states (Ortoleva, 2012)—may nudge the entrepreneur to consider novel attributes and causal links and hence be conducive of novel theories. Therefore, entrepreneurs use experiments not only to test their prior beliefs about a future state space, aiming at higher V_1, but also as a potential source of anomalous observations that can give rise to alternative theories (Mullainathan and Rambachan, 2023).

To the entrepreneur, the value of experimentation increases with the number of theories. A higher V_1 (higher expected value of theory Θ) makes it more valuable either to experiment with it or to commit to it. If V_1 is relatively small compared to V_2, the same applies to the alternative theory. As figure 1.4 shows (see the appendix for the formal proposition), there are two exploration zones: one for Θ and the other one for the alternative theory.

Downloaded from direct.mit.edu/books/oa-edited-volume/chapter-pdf/2595206/c000600_9780262052160.pdf by Università Bocconi user on 04 November 2026

This is a portion of the eBook at doi:10.7551/mitpress/15918.001.0001

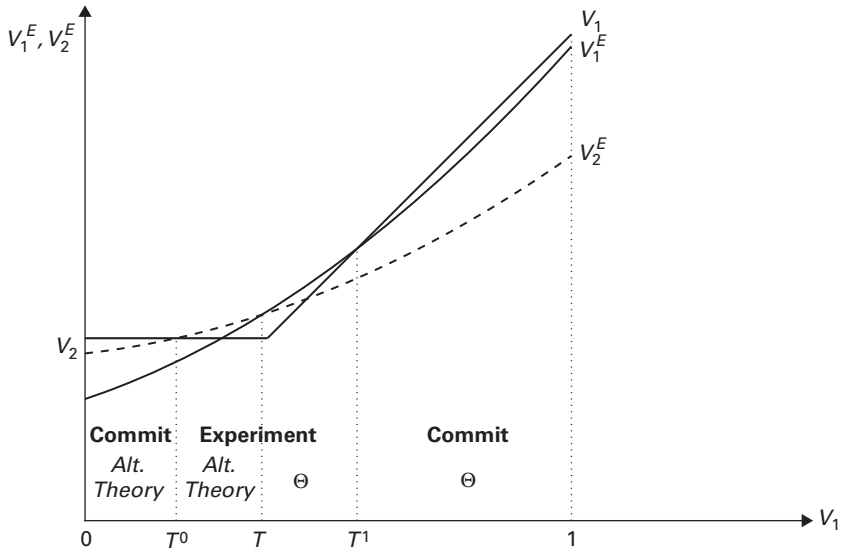


Figure 1.4
Value of experimentation and exploration zones.

Camuffo, Gambardella and Pignataro (2024) show under which conditions it is optimal for entrepreneurs to choose a given theory or to experiment. They also show that more uncertain theories (i.e., theories with larger spreads in parameters or with higher entropy priors) increase their relative space of exploration, the overall space of exploration, and the value of experimenting with it. These experiments are more informative because they might yield a higher information upside, while the alternative theory protects against the potential negative outcomes of the experiment. This greater value of more uncertain theories supports Felin and Zenger's (2017) conjecture about the value of testing contrarian beliefs. In addition, this framework is compatible with testing weaker priors—that is, theories such that, initially, $V_1 < V_2$. Thus, it can be optimal to test not only more uncertain theories but also weaker priors.

Furthermore, uncertain theories are complementary—that is, experimenting with a more uncertain theory increases the benefits of experimenting with other more uncertain theories. The intuition is that a more uncertain alternative theory raises the value of the outside option of an experiment on theory Θ, which raises the value of this experiment. This widens the exploration zone and the expected value of theories.

Downloaded from direct.mit.edu/books/oa-edited-volume/chapter-pdf/2595206/c000600_9780262052160.pdf by Università Bocconi user on 04 November 2026

This is a portion of the eBook at doi:10.7551/mitpress/15918.001.0001

1.6 Specific Research Topics

In this final section, we outline two areas where we believe a Bayesian approach can improve specific research areas in entrepreneurship.

1.6.1 A Bayesian Approach to Entrepreneurial Strategy

Entrepreneurial strategy is the sequence of choices a founding team makes to test specific value creation and capture hypotheses when entrepreneurial experimentation requires partial commitment (Gans, Scott and Stern, 2018; Gans et al., 2019; Gans, Scott, and Stern, 2024). Specifically, as entrepreneurs explore a given idea, they face many alternatives that cannot be pursued at once, and so must adopt (implicitly or explicitly) a process for testing their priors about, and ultimately choosing among, entrepreneurial strategies for moving forward their chosen entrepreneurial opportunity. The presence of deep uncertainty and limitations on commitment-free learning has the implication that entrepreneurial experimentation provides only partial insight into the relative value of alternative strategies. As a result, most early-stage entrepreneurial experiments are not singularly decisive, with founders facing trade-offs between three competing elements of experimental design (see figure 1.5):

- Criticality—the relative importance of the hypothesis being tested in the overall entrepreneurial strategy
- Fidelity—the degree to which a test of a given level of criticality provides meaningful and informative feedback to the founders about the hypothesis being tested
- Opportunity cost—the overall cost, including resources, time, and strategic commitments, required to conduct a test of a given fidelity with a given level of criticality

When choosing to conduct an entrepreneurial experiment, the single most important element to consider is whether a successful test would meaningfully reduce the most significant drivers of uncertainty associated with the entrepreneurial opportunity. Put simply, entrepreneurial experimentation is only valuable if the experiment is testing something that actually matters. Concretely, for any given idea and strategy, there is a set of hypotheses and sub-hypotheses that must turn out to be true in order for the idea and strategy to succeed. These include, among other factors, the presence of a paying customer, feasible technology and organization, and a commercialization path that allows for value creation and capture on a sustainable basis. The first step in the process of testing an entrepreneurial strategy is to delineate the most critical elements of the

Downloaded from direct.mit.edu/books/oa-edited-volume/chapter-pdf/2595206/c000600_9780262052160.pdf by Università Bocconi user on 04 November 2026

This is a portion of the eBook at doi:10.7551/mitpress/15918.001.0001

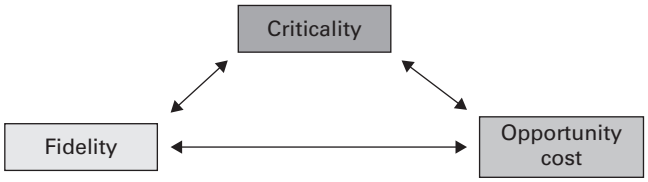


Figure 1.5
A strategic approach to entrepreneurial experimentation.

strategy to be tested. The most concrete consequence of stating these hypotheses clearly and crisply is to help identify both the underlying opportunity and most critical risks facing a venture. Being clear and sharp about the most critical hypotheses facing the venture allows the founding team to effectively test their entrepreneurial strategy and focus their time and attention on ensuring that those critical hypotheses turn out to be true.

Of course, simply knowing that a hypothesis is critical is not enough. Instead, one must be able to test that hypothesis in a way that is informative. For a test to have fidelity, it has to have the potential to meaningfully shape the future choices of the founders. To do so, the test, therefore, needs to meaningfully reduce uncertainty, providing validation (or not) to the overall idea or to a particular strategy. Entrepreneurial experiments must drive the entrepreneurs' choice of whether to pursue a given idea and choice of the entrepreneurial strategy to realize that opportunity. But, as with criticality, simply exhorting entrepreneurs to engage in high-fidelity, informative experiments is much simpler than actually designing and conducting those experiments. Two key obstacles—noise and bias—stand in the way.

First, even if one expends significant time and resources, the outcome of an entrepreneurial experiment might simply be noisy. Without a sufficiently high signal-to-noise ratio, it will be difficult to draw out meaningful insights that inform entrepreneurial choice. Particularly at the earliest stages of a venture, it is possible to engage in a wide range of "experiments" that nonetheless leave founders with data that yields neither holistic themes nor a consistent body of evidence. At the conclusion of such experiments, despite often considerable effort, founders have gained little to no clarity on the validity of their hypotheses or the path ahead. The challenge is often not the lack of centrality of the hypotheses but rather the lack of fidelity of the experiment. When exploring a new idea, any data—qualitative or quantitative—is likely to have significant variation, and being cognizant of how noisy data can be allow founders to avoid leaping too quickly to a conclusion on the basis of very preliminary evidence.

Downloaded from direct.mit.edu/books/oa-edited-volume/chapter-pdf/2595206/c000600_9780262052160.pdf by Università Bocconi user on 04 November 2026

This is a portion of the eBook at doi:10.7551/mitpress/15918.001.0001

The second potential obstacle to experimental learning is experimental bias. Bias occurs when the data or feedback from an experiment is distorted by the ways in which data are gathered or analyzed. For example, if founders first test their product or service idea on their friends and family (a sensible, low-cost, low-commitment approach), they are likely to simultaneously receive too much positive feedback from acquaintances who are simply trying to be polite and fail to gain accurate feedback from potential beachhead customers if there is limited overlap between that beachhead and their social network or other accessible settings for entrepreneurial experimentation (Cao et al., 2023). As a consequence, bias has the potential to encourage founders to believe in an assumption that is, in fact, false or dismiss a critical assumption that, in fact, holds.

The impact of noise and bias does not negate the value of systematic experimentation and learning. Instead, the very fact that many experiments are noisy and biased simply reinforces how little entrepreneurs actually know about their most critical hypotheses when they are choosing a given entrepreneurial opportunity and the entrepreneurial strategy to pursue that idea. The key is to take a proactive and disciplined approach to generate accurate and actionable information about the most critical hypotheses. The field continues to contribute a range of tools and frameworks that improve the fidelity of entrepreneurial experimentation, given the natural constraints entrepreneurs face (see, among others, Murray and Tripsas, 2004; Nanda and Rhodes-Kropf, 2016; Cohen et al., 2019; Koning et al., 2022).

The final critical factor impacting the value of entrepreneurial experimentation and learning is the opportunity cost of a given experiment. All else equal, the cheaper and faster one can make an entrepreneurial experiment, the more valuable it will be in terms of being able to inform the core choices of a new venture. Although this point might seem obvious, it is also easy to underestimate the value that arises from less costly or more timely testing. Most notably, when the uncertainty around an opportunity is significant, the simple process of clarifying the most critical assumptions or how to design a high-fidelity test is both challenging and time-consuming. However, identifying low-cost and speedy ways to explore the underlying idea, potential beachhead markets, or potential commercialization paths can sharply reduce uncertainty.

The relationship between cost, speed, and commitment is subtle. On the one hand, when the main cost of experimentation, is the time cost of the founders, lower cost, rapid experimentation, and low commitment all go together; interviewing more potential customers or running more experiments per day simply accelerates the process of learning (and does not necessarily commit the venture to one path over the other). This natural inverse relationship between cost and

Downloaded from direct.mit.edu/books/oa-edited-volume/chapter-pdf/2595206/c000600_9780262052160.pdf by Università Bocconi user on 04 November 2026

This is a portion of the eBook at doi:10.7551/mitpress/15918.001.0001

speed is at the heart of the lean start-up philosophy: "The only way to win is to learn faster than anyone else" (Ries, 2011). When entrepreneurs can focus their efforts and scarce resources on maximizing learning, rapid learning can serve as a potential source of competitive advantage. On the other hand, although rapid learning and low cost go together when the main resource being used is founder time, the relationship between speed and cost can be reversed if one is willing to, for example, enhance speed by incurring a higher upfront cost. For example, doubling the rate of customer interviews might involve hiring (perhaps temporary) employees who would need to be paid in cost or equity; more rapid product design iterations might involve investing in costly machinery (e.g., 3D printers). As such, it is important for entrepreneurs to clarify the relationship between cost and speed for the experiments they undertake and, in particular, how they can best make use of their own time (i.e., enhance the virtuous cycle of low-cost and rapid learning) while also being judicious in expending resources in order to allow learning to go faster (i.e., by expending money for more rapid feedback cycles).

However, in many cases, the most important "cost" of an experiment is neither time nor resources but instead a strategic commitment cost. Two particular types of strategic commitments are particularly salient in shaping the opportunity cost of an entrepreneurial experiment. First, many potentially informative entrepreneurial strategy tests involve committing to particular stakeholders, including potential customers, collaboration partners, or early employees. For example, launching a crowdfunding campaign, such as on Kickstarter or Indiegogo, is an excellent way to gauge the preliminary level of demand for an early-stage product idea, but such a test comes with the commitment to actually build and deliver the product for those who contribute to the campaign. However, as documented by Ethan Mollick, the vast majority of successful crowd funding projects are unable to meet their promised timelines, resulting in significant negative feedback from early users. As a result, the use of crowdfunding as a tool for entrepreneurial experimentation brings with it a high likelihood of ending up with a negative reputation with precisely those customers who initially expressed the greatest interest. Said otherwise, although a minimal viable product allows the entrepreneur to learn a great deal about customers, it also allows customers to learn about the product and venture, which should be pursued with awareness as to the commitments the venture is making to key stakeholders. Second, many entrepreneurial experiments involve the disclosure of information—either knowledge about the underlying idea or statements of future strategic intention—that reduce the scope for future strategic choice. Perhaps most notably, if one reveals information publicly about a novel technology, it is possible that the potential scope for formal intellectual property rights

Downloaded from direct.mit.edu/books/oa-edited-volume/chapter-pdf/2595206/c000600_9780262052160.pdf by Università Bocconi user on 04 November 2026

This is a portion of the eBook at doi:10.7551/mitpress/15918.001.0001

such as patents might be reduced. The strategic opportunity costs involved in an entrepreneurial experiment may limit the ability of an entrepreneur to effectively pursue that tested path, even if that path would have otherwise been viable based on the learning that occurred.

In other words, the costs of an experiment are not simply time but also the new resources and potential strategic commitments that come along with a given experiment. Rather than simply weighing the value of learning versus the value of time, entrepreneurs will have to consider more carefully whether that learning outweighs the overall opportunity cost of the proposed experiment. Calculating the potential value of alternative experiments is central to the effective design and sequencing of experiments in the process of entrepreneurial learning. Understanding the interplay between criticality, fidelity, and opportunity cost is central to this process. In an ideal case, it would be possible to quickly identify and implement experiments and approaches to entrepreneurial learning that simultaneously tested the most critical hypotheses with a high level of fidelity and at minimal opportunity cost. Most tests of entrepreneurial strategy involve significant trade-offs between the three dimensions—testing the most critical assumptions may require the highest cost or the best available data to inform hypotheses that are of only modest importance. Figuring out how to navigate the natural and inherent trade-offs between criticality, fidelity, and opportunity cost is among the most challenging yet rewarding elements of taking a strategic approach to entrepreneurship. A Bayesian approach to entrepreneurship will allow the field to further guide entrepreneurs in designing their entrepreneurial experiments and information on their choices as to whether to pursue a given idea and the entrepreneurial strategy to realize that opportunity.

1.6.2 The Scientific Approach for Deep Bayesian Learning

Another example of an entrepreneurial method that embodies Bayesian learning is the scientific approach (Camuffo et al., 2020). This approach encourages entrepreneurs to form prior beliefs on future state spaces using theories, test such beliefs through experiments, and use the associated evidence to update beliefs (Zellweger and Zenger, 2023). Prior formation through theorizing is the first step of the Bayesian learning cycle and involves engaging in deliberate cognitive efforts to formulate entrepreneurial problems, ideas, or opportunities as causal conceptual structures (Camuffo, Gambardella, Messinese et al., 2024; Ehrig and Schmidt, 2022; Felin and Zenger, 2017). Prior testing occurs through experiments that, as previously highlighted, are shaped by priors. Experiments elicit signals that are incorporated, through Bayes' rule, into updated beliefs. This allows entrepreneurs to make more informed decisions about whether to

Downloaded from direct.mit.edu/books/oa-edited-volume/chapter-pdf/2595206/c000600_9780262052160.pdf by Università Bocconi user on 04 November 2026

This is a portion of the eBook at doi:10.7551/mitpress/15918.001.0001

continue with the current project, pivot to a modified or different project, or terminate the venture.

As discussed above, experiments can either test prior beliefs within a given future state space (simple Bayesian learning) or test "priors on priors"—that is, question which future state space entrepreneurs should attend to when forming priors (deep Bayesian learning on theories of value). Likewise, the information produced by experiments can prompt straightforward Bayesian updating, or—if the evidence is sufficiently surprising—lead entrepreneurs to challenge their theories of value and revise their priors on priors (deep Bayesian learning). Notably, viewing the scientific process as deep Bayesian learning provides a theoretical foundation for entrepreneurial pivoting and helps explain the various typologies proposed in the literature (Kirtley and O'Mahony, 2023).

Below, we explain how the experimental evidence on the adoption of the scientific approach to entrepreneurial decision-making points to two mechanisms at work in generating more quitting, more focused pivoting, and higher performance: search efficiency, which corresponds to simple Bayesian learning, and methodic doubt, which corresponds to deep Bayesian learning.

The evidence coming from fourteen randomized control trials (RCTs) involving approximately three thousand start-ups around the world shows that the Bayesian learning embodied in a scientific approach to entrepreneurial decision-making increases the probability of terminating projects and of efficient pivoting and, conditional on survival, leads to higher performance (revenues) (Camuffo et al., 2020; Camuffo, Gambardella, Messinese et al., 2024). The evidence also shows that theories and experiments have separate, identifiable effects (Agarwal et al., 2025), and that scientific entrepreneurs better navigate uncertainty accounting for unknown events (Camuffo, Gambardella and Messinese, 2023). As previously highlighted, this occurs in ways compatible with reverse Bayesianism (Karni and Vierø, 2013, 2017) and theory-based decision-making (Ortoleva, 2012; Karni, 2022).

The scientific approach affects entrepreneurial decision-making in two ways (Camuffo, Gambardella, Messinese et al., 2024). On the one hand, entrepreneurs who adopt it use more clearly defined theory and validation procedures when they iteratively explore their ideas. This makes their search process more (*efficient search*). This raises the probability of success in earlier stages because it combines a higher probability of success stemming from theories and experiments and the ability to identify successful ideas earlier. This is because theories enable entrepreneurs to envision opportunities, form priors, and rank them better at the outset of the process. Entrepreneurs will then work first on priors that they expect to have a higher probability of success. When one of the ideas

Downloaded from direct.mit.edu/books/oa-edited-volume/chapter-pdf/2595206/c000600_9780262052160.pdf by Università Bocconi user on 04 November 2026

This is a portion of the eBook at doi:10.7551/mitpress/15918.001.0001

they are working on proves unsuccessful, they pivot to other ideas, but these will have a lower expected probability of success. Nonscientific entrepreneurs, absent theories, are instead more likely to learn about the value of their ideas as they work on them rather than ex ante in a trial-and-error fashion. Thus, they start with an idea/prior, and if it proves to be unsuccessful, the new idea they pivot to will not have an expected probability of success as low as in the case of scientific entrepreneurs because ideas are considered sequentially as they emerge in subsequent stages of exploration and not ranked beforehand based on theories.

As a result, efficient search allows entrepreneurs to work on more solid priors earlier in their exploration process. However, if efficient search was the only mechanism associated with the scientific approach, entrepreneurs who adopted it would simply be more likely to commit to their ideas rightaway, less likely to terminate, and more likely to pivot.

However, as described above, the scientific approach also instills in entrepreneurs a methodic doubt. The very fact that entrepreneurs who adopt this approach develop a theory for why their ideas may be successful makes them aware that this might be imperfect, false, or less plausible than others. As a result, they are more cautious (less optimistic about their ideas). Again, if methodic doubt was the only mechanism associated with the scientific approach, entrepreneurs who adopted it would be more likly to terminate and less likely to pivot or commit to their ideas.

It can be easily seen that efficient search and methodic doubt go in opposite directions, and it is an empirical matter to establish which of the two mechanisms dominates the other. The birth and early developments of ØSense—one of the start-ups accelerated in the above-mentioned RCTs—help to illustrate how the scientific approach embodies a process of Bayesian learning. The founder's initial idea was to develop existing technology (matching algorithms) to build a peer-to-peer rental service for household items, inspired by the success of Airbnb in the accommodation sector. The company was named YouRent, and its theory was that the concept underlying the sharing economy could be extended, by analogy, to everything, including everyday products and small items, allowing people—through technology—to rent items instead of buying them, and that this would create a new market with the value proposition of saving money while promoting sustainability and reducing waste. The theory summary was "Imagine Airbnb for everyday items." The founder was optimistic about the project, but the scientific approach learned during the RCT's intervention led him to doubt and wonder whether his beliefs were strong for good, logically grounded reasons. Hence, he set out to run an experiment in the form of field interviews

Downloaded from direct.mit.edu/books/oa-edited-volume/chapter-pdf/2595206/c000600_9780262052160.pdf by Università Bocconi user on 04 November 2026

This is a portion of the eBook at doi:10.7551/mitpress/15918.001.0001

with potential customers to test his priors. He ran several interviews at vintage fairs—a higher likelihood of individuals inclined toward reuse/rental of used goods—and set ex ante a threshold of 60 percent to accept their hypotheses. This was how he operationalized running an efficient experiment and getting the likelihood function necessary to apply Bayes' rule. The interviews did not support the necessity or willingness to engage in a rental platform for low-value items as hypothesized for YouRent. The founder negatively updated his prior and decided to terminate the initial project.

However, from those very interviews, a recurring theme of sustainability emerged, spurring a new potential future state space (a new potential theory and prior). The founder conceptually developed the insights from the interviews into WERI, a platform aimed at enhancing corporate car fleet sharing with a gamified experience that tracks and rewards sustainability efforts. Here, the founder believed that incentivizing sustainable behavior would also encourage the sharing of underutilized corporate fleet vehicles, leading to cost savings and reduced environmental impact.

Another Bayesian cycle started. Field interviews were conducted with fleet managers and sustainability officers to validate WERI. Also in this case, the theory underlying WERI did not get corroborated. However, the interviews pointed to a strong and generalized concern for sustainability and the need for reliable data to measure the carbon footprint of corporate activities, mainly driven by the pressure coming from top management teams (pushed by EU policy constraints). The founder refined this conceptual causal structure and crafted a new theory in which he envisioned his technology as an artificial intelligence platform to collect and reduce carbon emissions across the entire value chain. This led to Øsense (https://osense.ai/), a startup that helps firms "activate a real carbon reduction strategy by leveraging Artificial Intelligence to measure, simulate, reduce, and report your environmental footprint at scale." The pivot from WERI to Øsense was a total change of the theory. In the founder's words, "I reevaluated my entire theory. This new piece of information seemed initially disconnected and 'far' from my original framework. Through the process of questioning and engaging in cognitive efforts to explore logical connections, I managed to recognize the value of this opportunity and incorporate it into my understanding."

The case illustrates how Bayesian learning occurs for entrepreneurs adopting the scientific approach at two levels. First, scientific entrepreneurs form, test, and update their beliefs about the success or value of their ideas. Second, they form, test, and update their beliefs about the generating mechanisms underlying such ideas.

Downloaded from direct.mit.edu/books/oa-edited-volume/chapter-pdf/2595206/c000600_9780262052160.pdf by Università Bocconi user on 04 November 2026

This is a portion of the eBook at doi:10.7551/mitpress/15918.001.0001

1.6.3 A Bayesian Approach to the Market for Judgment

A Bayesian approach to the market for judgment builds on the foundational understanding that both entrepreneurs and investors operate under potentially different priors about the viability and potential of new ventures. Judgment, a critical component of decision-making, is the application of experiential knowledge to prioritize the priors that most urgently require updating and how to do so most effectively (Agrawal et al., 2021). Experienced entrepreneurs possess well-honed judgment derived from past successes and failures, enabling them to prioritize the actions that maximize the likelihood of venture success.

In a Bayesian framework, better judgment lowers the cost of experimentation by better sequencing and designing experiments. Better judgment enables better assessing the return on investment associated with updating which priors for which actors. In addition, better judgment increases the return on investment of the experiment itself by enhancing its design to generate a more effective signal at a lower cost. Experienced entrepreneurs, with their superior judgment, can devise experiments that provide effective signals that update the targeted actor(s). For instance, they might recommend conducting pilot studies or market tests that yield high-quality data on customer acceptance, technological feasibility, or competitive positioning. By leveraging their judgment, they increase the likelihood that each experiment is designed to minimize ambiguity and maximize the information gained to address the priors that are in most urgent need of updating.

However, the market for judgment faces significant frictions due to the difficulty of transferring experiential knowledge. Unlike tangible assets, judgment cannot be easily bought or sold. Programs such as Creative Destruction Lab (CDL) aim to mitigate these frictions by facilitating interactions between experienced and novice entrepreneurs (Lakhani et al., 2019; Luo, 2022). By creating structured environments where judgment can be shared, programs such as the CDL help new entrepreneurs access the insights and guidance of those who have successfully navigated the commercialization process before (Gans, 2018; Sariri, 2024). This exchange of judgment allows new entrepreneurs to update their priors and the priors of potential investors, customers, and partners more effectively. The Bayesian approach to the market for judgment may be implemented through such institutions, enabling the transfer of experiential knowledge in order to design experiments that more efficiently generate signals to address the highest-priority priors.

When introducing a new product based on a recent scientific invention, different actors—such as the entrepreneur, investors, and potential customers—are likely to have varying priors about the product's viability. Entrepreneurs, driven by their vision and understanding of the underlying science, might have a more

Downloaded from direct.mit.edu/books/oa-edited-volume/chapter-pdf/2595206/c000600_9780262052160.pdf by Università Bocconi user on 04 November 2026

This is a portion of the eBook at doi:10.7551/mitpress/15918.001.0001

optimistic prior about the product's potential. Investors, who assess the product's marketability and financial prospects, may hold more skeptical priors, balancing enthusiasm with caution. Potential customers, who are crucial for market validation, often base their priors on perceived utility and credibility of the product. These divergent perspectives necessitate targeted strategies to align beliefs and facilitate the product's successful commercialization.

Consider the case of a new medical device designed to detect early-stage cancer using a novel biomarker discovered in academic research. The entrepreneur, typically a scientist involved in the discovery, believes strongly in the device's efficacy and potential to revolutionize cancer diagnostics. However, investors might be wary due to the high failure rate of medical start-ups and the lengthy, costly regulatory approval process. Potential customers, such as hospitals and clinics, may doubt the device's accuracy and reliability compared with established methods. These differing priors highlight the challenge of securing investment and market adoption.

The urgency to update priors depends on the stage of product development and the most urgent barriers to overcome. In the early stages, it may be more critical to update the entrepreneur's prior, especially if their optimism leads to overestimation of the product's readiness or market fit. Accurate self-assessment is crucial to avoid costly missteps and to design feasible development plans. Conversely, once a prototype is available, updating investors' priors becomes more urgent, as their funding is essential for further development and scaling. Investors' confidence can be bolstered by credible demonstrations of the product's potential, reducing their perceived risk.

In our medical device example, in the initial phase, the entrepreneur might rigorously test the device within controlled environments to ensure that their confidence in its efficacy is justified. Suppose initial laboratory tests show promising results, indicating that the device accurately detects cancer biomarkers. At this point, updating the investors' priors becomes the priority. A different signal is now required. In other words, there may be an optimal sequence for generating different signals for updating the priors of different actors.

Once the sequence for updating priors is established, the entrepreneur must design experiments to provide the necessary signals as effectively and efficiently as possible. The experiment must be tailored to address the specific doubts and questions of the actor whose prior is most in need of adjustment. For entrepreneurs, this might involve extensive prototype testing and iterative refinement. For investors, it may involve pilot studies, market analyses, and validation from respected third-party entities. The goal is to produce robust, credible data that either confirm or challenge the prevailing beliefs, facilitating informed decision-making.

Downloaded from direct.mit.edu/books/oa-edited-volume/chapter-pdf/2595206/c000600_9780262052160.pdf by Università Bocconi user on 04 November 2026

This is a portion of the eBook at doi:10.7551/mitpress/15918.001.0001

In our example, after initial laboratory validation, the next step might involve a small-scale clinical trial to update investors' priors. This trial should be designed to provide clear, reliable data on the device's performance in a real-world setting. For instance, the trial could involve multiple hospitals and a diverse patient population to demonstrate the device's accuracy and usability across different conditions. Positive results from such a trial would significantly reduce investor skepticism, providing strong evidence that the device is ready for further development and scaling. This approach ensures that the most critical priors are addressed efficiently, paving the way for successful commercialization.

Thus, a Bayesian approach to the market for judgment provides a structured framework for aligning the diverse priors of different actors involved in the commercialization of scientific innovations. By leveraging the judgment of experienced entrepreneurs, this approach enhances the efficiency of experiments aimed at updating critical priors. Programs such as the CDL facilitate the transfer of experiential knowledge, helping novice entrepreneurs design and execute experiments that yield clear, actionable signals. Through this iterative process, entrepreneurs can systematically address the doubts of investors, potential customers, and other stakeholders, thereby increasing the likelihood of securing the necessary resources and market acceptance. Ultimately, the Bayesian framework for the market for judgment increases the likelihood that the commercialization process is guided by well-informed decisions, maximizing the potential for success in bringing scientific advancements to market.

1.7 Conclusions and Next Steps

This paper has outlined the key tenets of a Bayesian approach to entrepreneurship, emphasizing the central role of heterogeneous priors, the relatively strong priors held by entrepreneurs, and the process of belief updating through Bayesian learning. We argue that embracing this Bayesian perspective provides a unifying framework to connect and enhance diverse streams of research on entrepreneurial decision-making, including the formation of entrepreneurial beliefs, purposeful entrepreneurial experimentation, and the persuasion of key stakeholders.

By focusing attention on the interplay between an entrepreneur's prior beliefs and the potential for learning through experimentation to shape those beliefs, a Bayesian approach provides insight into the choices that entrepreneurs make regarding which opportunities to pursue and how to design experiments to reduce uncertainty and ultimately improve entrepreneurial performance. Critically, the presence of heterogeneous priors provides a natural explanation for

Downloaded from direct.mit.edu/books/oa-edited-volume/chapter-pdf/2595206/c000600_9780262052160.pdf by Università Bocconi user on 04 November 2026

This is a portion of the eBook at doi:10.7551/mitpress/15918.001.0001

why entrepreneurs agree to disagree about the value of a given opportunity, and the systematic stronger priors of entrepreneurs choosing to pursue a given idea provides a micro-foundation for observed entrepreneurial action.

Moreover, by distinguishing between simple versus deep Bayesian learning approaches, we highlight how entrepreneurial theories about value creation and capture can themselves evolve as experiments yield outcomes that challenge fundamental assumptions. This provides a more nuanced understanding of entrepreneurial pivoting behavior, grounded in how entrepreneurs interpret potentially surprising signals.

Building upon this foundation, we see several promising avenues for future research. First, empirical studies can leverage the Bayesian framework to develop sharper predictions and measures related to entrepreneurial belief formation, experimental choices, and pivoting behavior. This could involve both qualitative work examining how entrepreneurs articulate and update their beliefs, as well as quantitative analyses linking measures of prior beliefs to specific strategic choices and venture outcomes.

Second, we believe there is significant potential to develop richer formal models of entrepreneurial decision-making within a Bayesian framework. This could include game-theoretic models of entrepreneurial competition and cooperation under heterogeneous priors, optimal stopping models of entrepreneurial experimentation and abandonment, and persuasion models of entrepreneurial communication with key stakeholders. Such models could yield new insights and predictions to guide empirical work.

Third, the Bayesian approach suggests opportunities to reexamine and potentially reframe findings from behavioral approaches to entrepreneurship. Although behavioral models often emphasize cognitive biases and heuristics as deviations from rationality, a Bayesian perspective suggests some of these patterns may, in fact, reflect entrepreneurs' priors and how those priors shape their interpretation of new information. Empirical studies could seek to disentangle these mechanisms.

Finally, from a practical perspective, the Bayesian framework offers a structured approach to guide entrepreneurial decision-making. Pedagogical research could explore how to most effectively teach a Bayesian approach to entrepreneurship students and practitioners, potentially building on research related to the scientific approach. Tools and frameworks to help entrepreneurs articulate key assumptions, design informative experiments, and pivot based on surprising information could have significant impact.

In conclusion, we believe that the Bayesian approach to entrepreneurship provides a promising foundation for integrating key strands of entrepreneurship

Downloaded from direct.mit.edu/books/oa-edited-volume/chapter-pdf/2595206/c000600_9780262052160.pdf by Università Bocconi user on 04 November 2026

This is a portion of the eBook at doi:10.7551/mitpress/15918.001.0001

research and opening up productive new avenues for scholarly inquiry and practical impact. By placing the dynamic interplay of prior beliefs and experiential learning at the center of the entrepreneurial process, this perspective offers a compelling lens to understand the entrepreneurial judgments and choices that drive the creation of new value in the economy.

1.8 Appendix: Details of Deep Bayesian Learning

1.8.1 Models, Causal Links, and Theories

When entrepreneurs build theories, they envision future state spaces for which they (and, to the extent to which the theory is novel, others) have limited information. Hence, about this state space, they have sparse empirical observations and pseudo-observations they subjectively develop, for example using analogies or "bootstrapping" from abductive reasoning. Therefore, they ultimately fit preliminary probability distributions on the state space stemming from "soft" rather than "hard" empirical frequencies that occur in worlds that decision-makers envision in their mind.[20] To operationalize our approach, let's assume that entrepreneurs' probabilities are distributed as Dirichlet with parameters $n_{IJ} > 0, I = Y, N; J = Y, N$:

$$\propto Pr(Y,Y)^{n_{YY}-1} \cdot Pr(Y,N)^{n_{YN}-1} \cdot Pr(N,Y)^{n_{NY}-1} \cdot Pr(N,N)^{n_{NN}-1}. \quad (1.6)$$

In a Dirichlet distribution:

$$\mathbb{E}Pr(I,J) = \frac{n_{IJ}}{n},$$

with $n \equiv n_{YY} + n_{YN} + n_{NY} + n_{NN}$. These parameters represent subjective empirical distributions of the four states envisioned by the entrepreneurs in the example.[21]

It helps to redefine these parameters as:

$$\theta_{YY} \equiv \frac{n_{YY}}{n_{YY} + n_{NY}}, \quad \theta_{YN} \equiv \frac{n_{YN}}{n_{YN} + \eta_{NN}}, \quad \text{and} \quad \theta_{Y} \equiv \frac{n_{YY} + n_{NY}}{n},$$

and we define $\theta \equiv \{\theta_{YY}, \theta_{YN}, \theta_{Y}\}$ to be the set of parameters of our distribution. As discussed by Marinacci (2015, p. 1037), we can then write the statistical model:

$$\mathbb{E}Pr(Y) \equiv \nu(\theta) = \theta_{YY}\theta_{Y} + \theta_{YN}\,(1 - \theta_{Y}) = \theta_{YN} + \theta_{Y}(\theta_{YY} - \theta_{YN}). \quad (1.7)$$

Entrepreneurs acknowledge that their entrepreneurial decisions (pursue or not the opportunity) are affected not only by the randomness of their subjective probabilities $Pr(I,J)$ but also by the uncertainty about the parameters of

Downloaded from direct.mit.edu/books/oa-edited-volume/chapter-pdf/2595206/c000600_9780262052160.pdf by Università Bocconi user on 04 November 2026

This is a portion of the eBook at doi:10.7551/mitpress/15918.001.0001

their probability distributions (equation 1.6). A specific set of parameters, θ, identifies one Dirichlet distribution and, therefore, in our example, one expected probability of new market rise, $v(\theta)$. But entrepreneurs do not know the probability distribution of the parameters either. They realize that their empirical distributions are also random variables because they know that n_{IJ} is not an objective frequency of the state (I, J). Thus, they also form, in their head, a subjective probability distribution $\mu(\theta)$ of the parameters. This defines a family of probability distributions (equation 1.6) identified by the different θ in the space of all the possible realizations of the three parameters in the example.

1.8.2 Extent of Experimentation

We model the extent of entrepreneurial experimentation using models of optimal information acquisition (Moscarini and Smith, 2001; Camuffo, Gambardella and Pignataro, 2024). Consider an entrepreneur who has two theories, Θ and Γ, with expected values, respectively, V_1 and V_2. Without loss of generality, we focus on an experiment on theory Θ. Let $c > 0$ be a fixed cost of running an experiment, $\rho \in (0, 1)$ a discount factor accounting for the fact that the experiment delays the final decision, and $H(V_1')$ the cumulative probability distribution of the update V_1' of V from the experiment on Θ. Specifically, V_1' is the expected value of committing to Θ after the experiment or of running a new experiment on Θ.

Entrepreneurs can take a stream of sequential experiments, and at the beginning of each stage of experimentation, the condition for running the experiment is:

$$V_1^E \equiv \rho \left[\int_{V_2^*}^{1} V_1' dH(V_1') + V_2^* H(V_2^*) \right] - c > \Pi_{V1} \tag{1.8}$$

where $V_2^* \equiv max(V_2, V_2^E)$, $\Pi_{V1} \equiv max(V_1, V_2^*)$, and V_2^E is the equivalent value of the experiment on the alternative theory. This condition says that: (1) if after the experiment entrepreneurs observe $V_1' > V_2^*$, they will either commit to Θ or run a new experiment on Θ obtaining V_1'; (2) otherwise, they either commit to or run an experiment on the alternative theory Γ and obtain V_2^*.[22] The key insight is that entrepreneurs keep experimenting until the present value of the experiment is higher than all the other options, V_1 , V_2, and V_2^E.

Proposition 1.3. There are three thresholds such that: (a) If $V_1 \leq T^0$, entrepreneurs commit to the alternative theory Γ. If $V_1 \in (T^0, T)$, they run an experiment on the alternative theory Γ. If $V_1 \in (T, T^1)$, they run an experiment on theory Θ. If $V_1 \geq T^1$, they commit to theory Θ and obtain V.

Downloaded from direct.mit.edu/books/oa-edited-volume/chapter-pdf/2595206/c000600_9780262052160.pdf by Università Bocconi user on 04 November 2026

This is a portion of the eBook at doi:10.7551/mitpress/15918.001.0001

Proof. Using equation 1.8, and the fact that $\mathbb{E}V_1' = V_1$, rewrite V_1^E as:

$$V_1^E \equiv \rho\left[V_1 + V_2^* H(V_2^*) - \int_0^{V_2^*} V_1' dH(V_1')\right] - c = \rho\left[V_1 + \int_0^{V_2^*} H(V_1')dV_1'\right] - c. \tag{1.9}$$

where the second equality stems from integration by parts. This expressions establishes that V_1^E increases with V_1.

We first establish that:

$$\frac{\partial V_1^E}{\partial V_1} > \frac{\partial V_2^*}{\partial V_1} \geq 0.$$

Let $D = 1 - \rho^2 H(V_2^E)K(V_1^E)$, where K is the cumulative distribution of the update V_2' of V_2 if decision-makers run an experiment on the alternative theory. Take the differentials dV_1^E, dV_2^*, and dV_1 in (equation 1.9) and, if $V_2^* = V_2^E$, in the equivalent expression for V_2^E. Solving the system, or setting $dV_2^* = 0$ if $V_2^* = V_2$, obtain:

$$\frac{\partial V_1^E}{\partial V_1} = D^{-1}\rho > \frac{\partial V_2^*}{\partial V_1},$$

which is equal to 0 or $D^{-1}\rho^2 K(V_1^E)$, depending on whether $V_2^* = V_2$ or $Q^* = Q^E$.

When $V_1 \leq V_2^*$, $\Pi = V_2^*$, and the fact that V_1^E increases with V_1 faster than V_2^* is a necessary condition to state that there is a threshold T such that V_1^E switches from smaller to higher than V_2^*, this condition is sufficient for appropriate values of ρ or c. Since V_2^E grows with V_1, unlike V_2, V_2^* could switch from V_2 to V_2^E at a threshold $T^0 < T$.

When $V_1 > V_2^*$, $\Pi = V_1$, and V_1^E increases at a lower rate than V_1 if $D^{-1}\rho < 1$. Since $V_1, V_1^E < 1$, the V_1^E curve will cut the V_1 curve from above at a threshold T^1 such that if $V_1 > T^1$, entrepreneurs commit to theory Θ. The slower growth of V_2^E than V_1^E with respect to V_1 is a necessary condition for $V_2^E < V_1^E$ at $V_1 = T^1$. □

Notes

1. For example, Melnikoff and Strohminger (2024).

2. This requires each μ_i is continuous with respect to $\tilde{\mu}$ so that there exists a set of subjective beliefs that the information is correct such that $\tilde{\mu}_i = \mu_i$ for all $i \in N \cup 0$.

3. See, for instance, Gans (1995).

4. In the information design literature, this condition, where a signal determines the action, is known as "obedience"; see Bergemann and Morris (2019).

5. Arora and Fosfuri (2005) derive this willingness to pay function and explore different types of pricing structures for that information.

Downloaded from direct.mit.edu/books/oa-edited-volume/chapter-pdf/2595206/c000600_9780262052160.pdf by Università Bocconi user on 04 November 2026

This is a portion of the eBook at doi:10.7551/mitpress/15918.001.0001

6. It has also been implicitly assumed in work on entrepreneurial experimentation in finance; e.g., Nanda and Rhodes-Kropf (2016).

7. Gans (2025, ch. 10) shows that the value of an experiment for a decision-maker with a given prior is always maximized at an extremal choice, implying that an extremal choice always maximizes the aggregate willingness to pay of a population of decision-makers where the priors are distributed uniformly on [0, 1].

8. The notion of experimentation prior to larger finance commitments has been explored extensively by Nanda and Rhodes-Kropf (2017), Kerr et al. (2014), and Ewens et al. (2018).

9. The other investors' participation constraint for a given experiment is:

$$(1-\alpha)\frac{\lambda_1\mu_O}{\lambda_1\mu_O+(1-\lambda_O)(1-\mu_O)}V \geq C,$$

which implies that

$$\alpha \leq \frac{\lambda_1\mu_O V-(\lambda_1\mu_O+(1-\lambda_O)(1-\mu_O))C}{\lambda_1\mu_O V}.$$

The entrepreneur chooses an experiment to maximize:

$$\alpha\frac{\lambda_1\mu_E}{\lambda_1\mu_E+(1-\lambda_O)(1-\mu_E)}V$$

subject to the investor's participation constraint. Substituting this constraint and the experiment boundary gives:

$$\frac{\mu_E}{\mu_O}\frac{\lambda_1\mu_O V-(\lambda_1\mu_O+(1-\Lambda+\lambda_1)(1-\mu_O))C}{\lambda_1\mu_E+(1-\Lambda+\lambda_1)(1-\mu_E)}.$$

Note that the second derivative of this with respect to λ_1 is positive, so there is a corner solution. For $\lambda_1 = 1$, we have:

$$\frac{\mu_E}{\mu_O}\frac{\mu_O V-(\mu_O+(2-\Lambda)(1-\mu_O))C}{\mu_E+(2-\Lambda)(1-\mu_E)}.$$

For $\lambda_1 = \Lambda - 1$, we have $V - C$.

Note that the expected entrepreneurial profit, with $\lambda_1 = \Lambda - 1$ is greater than the profit with $\lambda_1 = 1$ if and only if:

$$\mu_O(V-C) \geq \frac{\mu_E}{\mu_E+(2-\Lambda)(1-\mu_E)}(\mu_O(V-C)-(2-\Lambda)(1-\mu_O)C),$$

which can be easily shown to hold.

10. This new literature also has deep connections linking persuasion to the behavioral concept of multiple selves over time; see Jakobsen (2021).

11. There is an important difference between the approach here and the Bayesian persuasion literature. That literature supposes that senders can commit to arbitrary signals. Here, it is assumed that the sender can only commit to an experiment, and the receiver can see the outcome. As Ball and Espín-Sánchez (2022) argue, this is closer to applications and also leads to different outcomes than the full Bayesian persuasion approach.

12. One paper that has explored these differing understandings of signals is Chavda et al. (2024).

13. There is a voluminous literature on Bayesian persuasion now. Some papers that are of relevance to the Bayesian entrepreneurship agenda include work that takes into account heterogeneous priors (Alonso and Câmara, 2016), work on persuasion with anecdotes (i.e., small samples; Haghtalab et al., 2021), non-Bayesian persuasion (de Clippel and Zhang, 2022), and price theoretic approaches (Dworczak and Martini, 2018).

14. More generally, an entrepreneur's goal is to concentrate the subjective probability distribution on $x_m \in \{Y, N\}$—that is, to become more persuaded whether the state they are interested in will occur. Alternatively, they could also reason counterfactually (Pearl and Mackenzie, 2018) and build a counter-theory that increases $\Pr[N]$ using a falsificationist approach.

15. There are often several $\mu_\Theta(\theta)$ compatible with the entrepreneur's partial knowledge and information. How entrepreneurs elicit a prior probability distribution from empirical distributions—their partial information and knowledge—is another driver of priors' heterogeneity. Among the various

Downloaded from direct.mit.edu/books/oa-edited-volume/chapter-pdf/2595206/c000600_9780262052160.pdf by Università Bocconi user on 04 November 2026

This is a portion of the eBook at doi:10.7551/mitpress/15918.001.0001

methods suggested in Bayesian statistics (see Kass and Wasserman, 1996, for a review), the principle of maximum entropy—that is, a prior probability assignment that incorporates the fewest assumptions on the data (Jaynes, 1968)—is particularly appealing to Bayesian entrepreneurs, as it approximates objectivity, ensures the highest learning potential, and reduces potential bias.

16. A more plausible theory is correlated but does not strictly imply that the theory is more valuable. This depends on the decision-maker's goals and the actions they choose, *given the theory that they choose*.

17. Chavda et al. (2024) consider a related notion of "theory-based" experimentation based on comparing experimental choices when one type of entrepreneur has a deeper theoretical understanding of the environment than other entrepreneurs, whereas Agrawal et al. (2021) examine the role of conflated signals and how a theory regarding the relationship between random variables determines experimental sequencing.

18. As highlighted above, the condition under which the expected value of the posteriors equals the prior corresponds to Bayes plausible experiments (Kamenica and Gentzkow, 2011). However, contingent on the level of their priors, entrepreneurs might prefer signals that are biased.

19. We assume that entrepreneurs can compare and rank the expected value of the theories.

20. Different approaches have been used to model the subjective process of building distributions for novel state spaces out of sparse information. For example, Gilboa et al. (2020) propose a model that couples case-based reasoning (i.e., sparse past data) and theories to form prior beliefs. Our approach is similar, although we do not limit the role of theories to analogical reasoning.

21. n = number of observations included in the empirical distribution of the decision-maker comprising both actual observations (facts and data) and pseudo-observations.

22. We model a general exploration stage of a dynamic problem. At every exploration stage, entrepreneurs redefine V_1, V_1^E, V_2, V_2^E . Since we also assumed decreasing information gains over time from experimentation, V_1^E and V_2^E will get smaller over time so that entrepreneurs will eventually commit to one of the two theories.

References

Acs, Z. J., Stam, E., Audretsch, D. B., and O'Connor, A. (2017). The lineages of the entrepreneurial ecosystem approach. *Small Business Economics*, 49, pp. 1–10.

Agarwal, R., Bacco, F., Camuffo, A., Coali, A., Gambardella, A., Msangi, H., Sonka, S., Temu, A., Waized, B., and Wormald, A. (2025). Does a theory-of-value add value? Evidence from a randomized control trial with Tanzanian entrepreneurs. *Organization Science*, 36(2), pp. 601–625.

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

Alonso, R., and Câmara, O. (2016). Bayesian persuasion with heterogeneous priors. *Journal of Economic Theory*, 165, pp. 672–706.

Amabile, T. M. (1997). Entrepreneurial creativity through motivational synergy. *The Journal of Creative Behavior*, 31(1), pp. 18–26.

Andrews, I., and Shapiro, J. M. (2021). A model of scientific communication. *Econometrica*, 89(5), pp. 2117–2142.

Arora, A., and Fosfuri, A. (2005). Pricing diagnostic information. *Management Science*, 51(7), pp. 1092–1100.

Åstebro, T., Chen, J., and Thompson, P. (2011). Stars and misfits: Self-employment and labor market frictions. *Management Science*, 57(11), pp. 1999–2017.

Åstebro, T., Herz, H., Nanda, R., and Weber, R. A. (2014). Seeking the roots of entrepreneurship: Insights from behavioral economics. *Journal of Economic Perspectives*, 28(3), pp. 49–70.

Aumann, R. (1976). Agreeing to disagree. *Annals of Statistics*, 4, pp. 1236–1239.

Ball, I., and Espín-Sánchez, J. (2022). *Experimental persuasion*. Cowles Foundation Discussion Paper No. 2298. https://ssrn.com/abstract=3906869

Downloaded from direct.mit.edu/books/oa-edited-volume/chapter-pdf/2595206/c000600_9780262052160.pdf by Università Bocconi user on 04 November 2026

This is a portion of the eBook at doi:10.7551/mitpress/15918.001.0001

Baron, R. (1998). Cognitive mechanisms in entrepreneurship: Why and when entrepreneurs think differently than other people. *Journal of Business Venturing*, 13(4), pp. 275–294.

Baum, J., Frese, M., and Baron, R. (Eds). (2014). *The psychology of entrepreneurship*. Psychology Press.

Baum, J., and Locke, E. (2004). The relationship of entrepreneurial traits, skill, and motivation to subsequent venture growth. *Journal of Applied Psychology*, 89(4), p. 587.

Bergemann, D., and Morris, S. (2019). Information design: A unified perspective. *Journal of Economic Literature*, 57(1), pp. 44–95.

Bhide, A. (1994). How entrepreneurs craft strategies that work. *Harvard Business Review*, 72(1994), p. 3.

Blackwell, D. (1953). Equivalent comparisons of experiments. *The Annals of Mathematical Statistics*, 24(3), pp. 265–272.

Blume, B., and Covin, J. (2011). Attributions to intuition in the venture founding process: Do entrepreneurs actually use intuition or just say that they do? *Journal of Business Venturing*, 26(1), pp. 137–151.

Burke, J. (2008). Primetime spin: Media bias and belief confirming information. *Journal of Economics and Management Strategy*, 17(3), pp. 633–665.

Calvert, R. L. (1985). The value of biased information: A rational choice model of political advice. *The Journal of Politics*, 47(2), pp. 530–555.

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), pp. 564–586.

Camuffo, A., Gambardella, A., and Messinese, D. (2023). Bayesianism and unforeseen events: Empirical evidence from two field experiments. *Academy of Management Proceedings*, Vol. 2023, 1.

Camuffo, A., Gambardella, A., Messinese, D., Novelli, E., Paolucci, E., and Spina, C. (2024). A scientific approach to entrepreneurial decision-making: Large-scale replication and extension. *Strategic Management Journal*, 45(6), pp. 1209–1237.

Camuffo, A., Gambardella, A., and Pignataro, A. (2024). Theory-driven strategic management decisions. *Strategy Science*, 9(4), pp. 382–396.

Carroll, G. R., and Sorensen, J. (2021). *Making great strategy: Arguing for organizational advantage*. Columbia University Press.

Chavda, A., Gans, J. S., and Stern, S. (2024). Theory-based entrepreneurial search. *Strategy Science*, 9(4), pp. 397–415.

Che, Y., and Mierendorff, K. (2019). Optimal dynamic allocation of attention. *American Economic Review*, 109(8), pp. 2993–3029.

Clippel, G. de, and Zhang, X. (2022). Non-Bayesian persuasion. *Journal of Political Economy*, 130(10), pp. 2594–2642.

Cohen, S., Bingham, C., and Hallen, B. (2019). The role of accelerator designs in mitigating bounded rationality in new ventures. *Administrative Science Quarterly*, 64(4), pp. 810–854.

Damiano, E., Li, H., and Suen, W. (2020). Learning while experimenting. *The Economic Journal*, 130(625), pp. 65–92.

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

Dushnitsky, G. (2010). Entrepreneurial optimism in the market for technological inventions. *Organization Science*, 21(1), pp. 150–167.

Dworczak, P., and Martini, G. (2018). The simple economics of optimal persuasion. *Journal of Political Economy*, 127(5), pp. 1993–2048.

Eesley, C., and Roberts, E. (2012). Are you experienced or are you talented?: When does innate talent versus experience explain entrepreneurial performance? *Strategic Entrepreneurship Journal*, 6(3), pp. 207–219.

Downloaded from direct.mit.edu/books/oa-edited-volume/chapter-pdf/2595206/c000600_9780262052160.pdf by Università Bocconi user on 04 November 2026

This is a portion of the eBook at doi:10.7551/mitpress/15918.001.0001

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), pp. 1287–1318.

Eisenmann, T., Ries, E., and Dillard, S. (2012). *Hypothesis-driven entrepreneurship: The lean startup*. Harvard Business School Entrepreneurial Management Case (812-095).

Evans, D. S., and Jovanovic, B. (1989). An estimated model of entrepreneurial choice under liquidity constraints. *Journal of Political Economy*, 97(4), pp. 808–827.

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

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

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

Fraser, S., and Greene, F. (2006). The effects of experience on entrepreneurial optimism and uncertainty. *Economica*, 73(290), pp. 169–192.

Frese, M., and Gielnik, M. (2014). The psychology of entrepreneurship. *Annual Review of Organizational Psychology and Organizational Behavior*, 1(1), pp. 413–438.

Gans, J. (1995). Evolutionary selection of beliefs. *Economics Letters*, 49(1), pp. 13–17.

Gans, J. (2018). *A better way to bring science to market*. MIT Sloan Management Review.

Gans, J. (2022). *A theory of visionary disruption*. NBER Working Paper (No. 30091).

Gans, J. (2023). Experimental choice and disruptive technologies. *Management Science*, 69(11), pp. 7044–7058.

Gans, J. (2025). *The microeconomics of artificial intelligence*. MIT Press.

Gans, J., Kearney, M., Scott, E., and Stern, S. (2021). Choosing technology: An entrepreneurial strategy approach. *Strategy Science*, 6(1), pp. 39–53.

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

Gans, J., Scott, E., and Stern, S. (2024). *Entrepreneurship: Choice and strategy*. W. W. Norton & Company.

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

Geanakoplos, J., and Polemarchakis, M. (1982). We can't disagree forever. *Journal of Economic Theory*, 28(1), pp. 192–200.

Gilboa, I., Minardi, S., and Samuelson, L. (2020). Theories and cases in decisions under uncertainty. *Games and Economic Behavior*, 123, pp. 22–40.

Gimeno, J., Folta, T., Cooper, A., and Woo, C. (1997). Survival of the fittest? Entrepreneurial human capital and the persistence of underperforming firms. *Administrative Science Quarterly*, 40(4), pp. 750–783.

Guzman, J., Oh, J., and Sen, A. (2020). What motivates innovative entrepreneurs? Evidence from a global field experiment. *Management Science*, 66(10), pp. 4808–4819.

Haghtalab, N., Immorlica, N., Lucier, B., Mobius, M., and Mohan, D. (2021). *Persuading with anecdotes*. NBER Working Paper (No. 28661).

Hamilton, B. (2000). Does entrepreneurship pay? An empirical analysis of the returns to self-employment. *Journal of Political Economy*, 108(3), pp. 604–631.

Harsanyi, J. (1967). Games with incomplete information played by "Bayesian" players, parts I, II and III. *Management Science*, 14(5), pp. 159–182.

Hmieleski, K., and Baron, R. (2009). Entrepreneurs' optimism and new venture performance: A social cognitive perspective. *Academy of Management Journal*, 52(3), pp. 473–488.

Hsieh, T. (2010). Zappos's CEO on going to extremes for customers. *Harvard Business Review*, 88(7–8), pp. 41–45.

Downloaded from direct.mit.edu/books/oa-edited-volume/chapter-pdf/2595206/c000600_9780262052160.pdf by Università Bocconi user on 04 November 2026

This is a portion of the eBook at doi:10.7551/mitpress/15918.001.0001

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

Jakobsen, A. (2021). An axiomatic model of persuasion. *Econometrica*, 89(5), pp. 2081–2116.

Jakobsen, A. (2023). Coarse Bayesian updating. Working Paper, Northwestern University, Evanston, IL.

Jaynes, E. (1968). Prior probabilities. *IEEE Transactions on Systems Science and Cybernetics*, 4(3), pp. 227–241.

Kamenica, E. (2019). Bayesian persuasion and information design. *Annual Review of Economics*, 11, pp. 249–272.

Kamenica, E., and Gentzkow, M. (2011). Bayesian persuasion. *American Economic Review*, 101(6), pp. 2590–2615.

Karni, E. (2022). A theory-based decision model. *Journal of Economic Theory*, 201, p. 105444.

Karni, E., and Vierø, M. (2013). Reverse Bayesianism: A choice-based theory of growing awareness. *American Economic Review*, 103(7), pp. 2790–2810.

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

Kartik, N., Xu Lee, F., and Suen, W. (2021). Information validates the prior: A theorem on bayesian updating and applications. *American Economic Review: Insights*, 3(2), pp. 165–182.

Kass, R. E., and Wasserman, L. (1996). The selection of prior distributions by formal rules. *Journal of the American Statistical Association*, 91(435), pp. 1343–1370.

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

Kirtley, J., and O'Mahony, S. (2023). What is a pivot? explaining when and how entrepreneurial firms decide to make strategic change and pivot. *Strategic Management Journal*, 44(1), pp. 197–230.

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

Koning, R., Hasan, S., and Chatterji, A. (2022). Experimentation and startup performance: Evidence from A/B testing. *Management Science*, 68(9), pp. 6434–6453.

Lakhani, K., Luo, H., and Katsnelson, L. (2019). *Market for judgement: The Creative Destruction Lab*. Harvard Business School Case (No. 719-479).

Landier, A., and Thesmar, D. (2008). Financial contracting with optimistic entrepreneurs. *The Review of Financial Studies*, 22(1), pp. 117–150.

Lazear, E. (2004). Balanced skills and entrepreneurship. *American Economic Review*, 94(2), pp. 208–211.

Liedtka, J. (2018). Why design thinking works. *Harvard Business Review*, 96(5), pp. 72–79.

Lindholm-Dahlstrand, Å., Andersson, M., and Carlsson, B. (2019). Entrepreneurial experimentation: A key function in systems of innovation. *Small Business Economics*, 53, pp. 591–610.

Luo, H. (2022). *Market for judgement: The Creative Destruction Lab*. Harvard Business School Teaching Note (No. 722-419).

Manso, G. (2016). Experimentation and the returns to entrepreneurship. *The Review of Financial Studies*, 29(9), pp. 2319–2340.

Marinacci, M. (2015). Model uncertainty. *Journal of the European Economic Association*, 13(6), pp. 1022–1100.

Martin, B., McNally, J., and Kay, M. (2013). Examining the formation of human capital in entrepreneurship: A meta-analysis of entrepreneurship education outcomes. *Journal of Business Venturing*, 28(2), pp. 211–224.

Masters, B., and Thiel, P. (2014). *Zero to one: Notes on start ups, or how to build the future*. Random House.

Maurya, A. (2022). *Running lean*. O'Reilly Media.

Downloaded from direct.mit.edu/books/oa-edited-volume/chapter-pdf/2595206/c000600_9780262052160.pdf by Università Bocconi user on 04 November 2026

This is a portion of the eBook at doi:10.7551/mitpress/15918.001.0001

Melnikoff, D. E., and Strohminger, N. (2024). Bayesianism and wishful thinking are compatible. *Nature Human Behaviour*, 8, pp. 692–701.

Minniti, M., and Bygrave, W. (2001). A dynamic model of entrepreneurial learning. *Entrepreneurship Theory and Practice*, 25(3), pp. 5–16.

Moore, G. (2014). *Crossing the chasm*. Harper Business.

Moscarini, G., and Smith, L. (2001). The optimal level of experimentation. *Econometrica*, 69(6), pp. 1629–1644.

Mullainathan, S., and Rambachan, A. (2023). *From predictive algorithms to automatic generation of anomalies*. NBER Working Paper (No. 32422).

Murray, F., and Tripsas, M. (2004). The exploratory processes of entrepreneurial firms: The role of purposeful experimentation. *Business strategy over the industry lifecycle*. Emerald, pp. 45–75.

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

Nanda, R., and Rhodes-Kropf, M. (2017). Financing risk and innovation. *Management Science*, 63(4), pp. 901–918.

Ortoleva, P. (2012). Modeling the change of paradigm: Non-Bayesian reactions to unexpected news. *American Economic Review*, 102(6), pp. 2410–2436.

Osterwalder, A., and Pigneur, Y. (2010). *Business model generation: A handbook for visionaries, game changers, and challengers*, volume 1. John Wiley & Sons.

Patvardhan, S., and Ramachandran, J. (2020). Shaping the future: Strategy making as artificial evolution. *Organization Science*, 31(3), pp. 671–697.

Pearl, J. (2000). *Causality: Models, reasoning, and inference*. Cambridge University Press.

Pearl, J., and Mackenzie, D. (2018). *The book of why: The new science of cause and effect*. Hackhette UK.

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

Roach, M., and Sauermann, H. (2015). Founder or joiner? The role of preferences and context in shaping different entrepreneurial interests. *Management Science*, 61(9), pp. 2160–2184.

Robson, A. J., and Samuelson, L. (2011). The evolutionary foundations of preferences. *Handbook of social economics 1*, North-Holland, pp. 221–310.

Rosenberg, N. (1992). Economic experiments. *Industrial and Corporate Change*, 1(1), pp. 181–203.

Samuelson, L. (2004). Modeling knowledge in economic analysis. *Journal of Economic Literature*, 42(2), pp. 367–403.

Sarasvathy, S. D. (2001). Causation and effectuation: Toward a theoretical shift from economic inevitability to entrepreneurial contingency. *Academy of Management Review*, 26(2), pp. 243–263.

Sariri, A. (2024). *Are angel investors more likely than venture capitalists to drive entrepreneurial experimentation?* Purdue University Working Paper.

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

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

Stevenson, H. H. (1983). *A perspective on entrepreneurship*, Vol. 13. Harvard Business School Working Paper #9-384-131.

Stuart, T., and Sorenson, O. (2005). Social networks and entrepreneurship. *Handbook of entrepreneurship research: Interdisciplinary perspectives*, Springer, pp. 233–252.

Suen, W. (2004). The self-perpetuation of biased beliefs. *The Economic Journal*, 114(495), pp. 377–396.

Sull, D. (2004). Disciplined entrepreneurship. *MIT Sloan Management Review*, 46(1), p. 71.

Downloaded from direct.mit.edu/books/oa-edited-volume/chapter-pdf/2595206/c000600_9780262052160.pdf by Università Bocconi user on 04 November 2026

This is a portion of the eBook at doi:10.7551/mitpress/15918.001.0001

Tartari, V., and Stern, S. (2021). *More than an ivory tower: The impact of research institutions on the quantity and quality of entrepreneurship*. NBER Working Paper (No. w28846).

Unger, J. M., Rauch, A., Frese, M., and Rosenbusch, N. (2011). Human capital and entrepreneurial success: A meta-analytical review. *Journal of Business Venturing*, 26(3), pp. 341–358.

Van de Ven, A., Sapienza, H., and Villanueva, J. (2007). Entrepreneurial pursuits of self- and collective interests. *Strategic Entrepreneurship Journal*, 1(3–4), pp. 353–370.

Van den Steen, E. (2001). Essays on the Managerial Implications of Differing Priors. PhD thesis, Stanford University.

Van den Steen, E. (2005). Organizational beliefs and managerial vision. *Journal of Law, Economics, and Organization*, 21(1), pp. 256–283.

Van den Steen, E. (2011). Overconfidence by Bayesian-rational agents. *Management Science*, 57(5), pp. 884–896.

Yildiz, M. (2000). Essays on Sequential Bargaining. PhD thesis, Stanford University.

Zellweger, T., and Zenger, T. (2023). Entrepreneurs as scientists: A pragmatist approach to producing value out of uncertainty. *Academy of Management Review*, 48(3), pp. 379–408.

Downloaded from direct.mit.edu/books/oa-edited-volume/chapter-pdf/2595206/c000600_9780262052160.pdf by Università Bocconi user on 04 November 2026

This is a section of doi:10.7551/mitpress/15918.001.0001

Bayesian Entrepreneurship

Edited by: Ajay Agrawal, Arnaldo Camuffo, Alfonso Gambardella, Joshua Gans, Erin L Scott, Scott Stern

Citation:
Bayesian Entrepreneurship
Edited by: Ajay Agrawal, Arnaldo Camuffo, Alfonso Gambardella, Joshua Gans, Erin L Scott, Scott Stern
DOI: 10.7551/mitpress/15918.001.0001
ISBN (electronic): 9780262052160
Publisher: The MIT Press
Published: 2026

The open access edition of this book was made possible by generous funding and support from MIT Press Direct to Open



The MIT Press

Downloaded from direct.mit.edu/books/oa-edited-volume/chapter-pdf/2595206/c000600_9780262052160.pdf by Università Bocconi user on 04 November 2026

MIT Press Direct

The MIT Press
Massachusetts Institute of Technology
77 Massachusetts Avenue
Cambridge, MA 02139
mitpress.mit.edu

© 2026 Massachusetts Institute of Technology

This work is subject to a Creative Commons CC-BY-NC-ND license.

This license applies only to the work in full and not to any components included with permission. Subject to such license, all rights are reserved. No part of this book may be used to train artificial intelligence systems without permission in writing from the MIT Press.



The MIT Press would like to thank the anonymous peer reviewers who provided comments on drafts of this book. The generous work of academic experts is essential for establishing the authority and quality of our publications. We acknowledge with gratitude the contributions of these otherwise uncredited readers.

This book was set in Times New Roman by Westchester Publishing Services. Printed and bound in the United States of America.

Library of Congress Cataloging-in-Publication Data is available.

ISBN: 978-0-262-05215-3

EU Authorised Representative: Easy Access System Europe, Mustamäe tee 50, 10621 Tallinn, Estonia | Email: gpsr.requests@easproject.com

Downloaded from direct.mit.edu/books/oa-edited-volume/chapter-pdf/2595206/c000600_9780262052160.pdf by Università Bocconi user on 04 November 2026