Bayesian Entrepreneurship

Ajay Agrawal, University of Toronto and NBER
Arnaldo Camuffo, Bocconi University
Alfonso Gambardella, Bocconi University
Joshua S Gans, University of Toronto and NBER
Erin L Scott, MIT Sloan
Scott Stern, MIT Sloan and NBER

March 24, 2024

Version 0.6

Abstract

How does the fact that entrepreneurs *choose* the opportunity they pursue impact entrepreneurial strategy and performance? Entrepreneurs, while dealing with opportunities whose outcome is inherently uncertain, have choices that must be premised on a belief that the opportunity is worth pursuing. This insight provides an organizing principle for a *Bayesian* approach to entrepreneurial decision-making. A Bayesian approach offers a natural formal framework to assess how entrepreneurs form beliefs about the prospects for a given opportunity, how these beliefs evolve over time through active experimentation and learning, and the consequences of such beliefs for entrepreneurial strategy and performance.

The goal is to shape distinctive implications and empirical approaches to the study of entrepreneurship guided by founding premises. The first premise is that the entrepreneur must be relatively *optimistic* about the opportunity *relative* to others. This involves a distinct theory that translates into a different perspective on the opportunity's prospects. Second, this systematic divergence in beliefs impacts how an entrepreneur will undertake learning about an opportunity. Notably, the demand for "experiments" 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" than "false negatives." Finally, this approach promotes the processes by which entrepreneurs are able to attract resources and capabilities by providing information to other agents. Entrepreneurs are more likely to convince those who share their idiosyncratically optimistic beliefs about an opportunity (with implications for homophily and firm culture), yet will also engage in choosing experiments that cater to those with different (more negative) beliefs than they themselves hold.

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. And 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 non-financial 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 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, D. S., & Jovanovic (1989); Hamilton (2000); Hurst & Lusardi (2004); Lazear (2004); Astebro, Chen & Thompson (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 & 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 & Rhodes-Kropf (2016); Manso (2016)). In other words, while 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) others, Gans et al. (2019); Camuffo et al. (2020); and; importantly, Ries (2011) 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 versus "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 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 entrepreneurial decision-making takes place under a 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 paper is to synthesize the different streams of work in this third approach under an umbrella, which 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. Though 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 & Thiel (2014)). Almost by construction, an entrepreneur choosing a particular opportunity must hold a *relatively optimistic* 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. While there were, of course, many other online retail businesses 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 despite 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: though footwear was a large and growing category (40 billion USD), 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, while 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, while 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 towards 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.

While 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, while 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 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 which 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 optimistic 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 was optimistic 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 also 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.

Finally, this approach yields novel insight into who (and how) entrepreneurs are able to attract resources and capabilities through by providing information to other agents. Notably, entrepreneurs are more likely to persuade those who share their idiosyncratically optimistic 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 not simply reflect a preference for association among similar types but attraction to an opportunity for those that share similar beliefs (Van den Steen (2005); Van den Steen (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 & Gentzkow (2011)) are a natural consequence of a Bayesian entrepreneurship approach).

The remainder of the paper 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 then dig into the role of priors, the nature of learning and experimentation, and the role of informativeness. We 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.

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 through learning (and how that updating informs practical decision-making), the utilization of a Bayesian approach towards 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*.

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 Frank Knight in his classic 1924 volume – in which they 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) allow 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 provides a rigorous framework for managing it. While Knightian uncertainty underscores the limits of predictability in entrepreneurial ventures, the Bayesian approach embraces these limits as a starting point for systematic improvement in decision-making under uncertainty.

Not **Shared Priors**

Traditional equilibrium models in economics and finance often operate under the assumption that all agents, including entrepreneurs and investors like 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 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.

Not **Effectuation**

Entrepreneurial effectuation posits a process where opportunities are not discovered but 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 pre-existing 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 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 co-created with stakeholders, the Bayesian approach views opportunity as something that can be perceived and assessed through rigorous analysis and experimentation.

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 like pattern matching, which are thought to significantly influence entrepreneurial decision-making (Astebro, Herz, Nanda & Weber (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 labelled as biases might, in fact, be the result of varying priors among decision-makers.

Under a Bayesian lens, the apparent over-optimism of entrepreneurs, for instance, may not universally reflect a bias towards 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 is 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. These heuristics may represent efficient, albeit imperfect, strategies for updating beliefs in the face of new evidence, especially under constraints of time and information.

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 the decisions of entrepreneurs. Through this lens, we can begin to see these decision-makers not as inherently flawed in their reasoning, but as individuals navigating uncertainty with the tools and information at their disposal.

3 What is Bayesian Entrepreneurship?

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

3.1 A Quick Primer on Bayesian Learning

With this as a background, it is useful to begin with a quick primer on how 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, $\tilde{\mu}$ is the updated probability that takes into account new information using Bayes' Rule. Specifically,

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

The output of this updating procedure, $\tilde{\mu}$, is called the 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.[1]

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

1. Forming a prior distribution about an entrepreneurial idea;
2. Collect data (e.g. through an experiment) and determine the likelihood function using the information about the parameters available in the data; and

[1]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).

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.

From the 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 & Schmidt, 2022; Felin & Zenger, 2009, 2017), on entrepreneurship as experimentation (Kerr et al., 2014; Koning et al., 2022; Lindholm-Dahlstrand et al. 2019), and on the scientific approach to entrepreneurial decision-making (Camuffo et al., 2020; Camuffo et al., 2024; Zellweger & Zenger, 2023).

From the practitioners' standpoint, the Bayesian approach provides an overarching rationale to understand the widespread adoption, often in combination, of methods like design thinking (Liedtka, , 2018), the business model canvas (Osterwalder & 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 & 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.

3.2 Priors 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 & 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 others' 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; Van den Steen,

2001) who have developed ways to move on from the common priors assumption. Indeed, Harsanyi (1967/68) 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 outlines potential ways forward. His starting point is a game agents, $i \in \{1, ..., N\}$ and a state space $\Theta \in \{\theta_1, \theta_2\}$. One simple approach is to assume that each agent has a prior, μ_i that the state is θ_1 and a possible reference prior, μ_0 that will be useful in order to conduct normative assessments. A more elaborate set-up involves a common prior $\hat{\mu}$ but that agents observe information that is objectively correct with probability p_0 but over which agents have subject beliefs regarding that information being correct of $\{p_i\}_{i \in N}$. 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. This requires each μ_i is continuous with respect to $\tilde{\mu}$ so that there exists a set of p_i such that $\tilde{\mu}_i = \mu_i$ for all $i \in N \cup 0$.

The second approach is cumbersome relative to the first, but Van den Steen (2001) notes that it makes it "easier to judge the 'reasonableness' of priors in the analysis" (p.16) and to link results to psychological biases and behavioural considerations such as over-confidence; it creates a link between beliefs and some fundamental reality that can assist in evaluating efficiency; and that it tackles 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 the theory-driven process described below, which elaborates on how entrepreneurs build the future state spaces (and prior beliefs) underlying their ideas. From this perspective, prior heterogeneity stems from both differences in the envisioned state space and from how entrepreneurs elicit the prior probability distribution. 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). 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 though social processes (e.g. communication, as described below) so that $\phi_i = \phi_j$ ("common knowledge" condition) (Samuelson, 2004).

3.3 Entrepreneurs have Relatively Optimistic Priors

What is optimism? This is likely a nuanced and complex issue. However, when push comes to shove, most agree that, in an innovation context, it is when an individual holds a prior belief that involves a higher likelihood that the innovation would be successful by whatever criteria one may consider. This is the approach taken by Van den Steen (2005). Thus, a working definition is as follows:

Definition 1 (Optimism)**.** *An individual, i, is more optimistic about an idea relative to individual j, if $\mu_i > \mu_j$.*

This definition does not define optimism per se but *relative optimism*. Optimism itself might be a threshold. For instance, it could be a threshold whereby someone would undertake an investment in an idea without additional information, i.e., $\mu_i > \bar{\mu}$ where $\bar{\mu} \equiv \frac{C}{V}$. It could also be defined relative to available experiments whereby optimism is a threshold whereby even a negative signal resulting in $\tilde{\mu}(0)$ would still involve $\tilde{\mu}(0)V \geqslant C$ implying that an individual is so optimistic that no information would dissuade them from pursuing an entrepreneurial opportunity.

The relative definition of optimism, however, does not require any such thresholds to be met. It only requires that (a) individuals hold heterogeneous priors and (b) that they can be ranked along an 'optimisim' dimension. 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.

A central premise of Bayesian Entrepreneurship is that **the entrepreneur is more optimistic about their ideas than all (or most) others**. Relative to optimism as an entrepreneurial 'trait' (Dushnitsky, 2010; Fraser & Greene, 2006; Hmieleski & Baron, 2009; Landier & Thesmar, 2008), Bayesian Entrepreneurship simply requires that the entrepreneur is optimistic about the idea they are pursuing even if they might be a pessimist about realisations of other uncertain variables. While we do not explicitly model here the idea selection mechanism (e.g., one could ground opportunity recognition and initial action in the context of evolutionary game theory), this will be an assumption that we will maintain throughout.[2]

The consequence of this premise is that entrepreneurs and others understand that entrepreneurs are more optimistic, so they agree to disagree along the lines described above. 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 expected posterior equals their prior.

Interestingly, agreement is desirable, but only for instrumental reasons. In a competitive context, entrepreneurs may benefit if others do not share their priors. This relative

[2]See, for instance, Gans (1995).

pessimism from others implies 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 pre-condition for the entrepreneur to find it worthwhile to pursue an opportunity (Robson & Samuelson, 2011).

A key implication of this approach 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.

3.4 Priors Matter for Experimentation Incentives

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 commercialisation 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, the second constituent part of the Bayesian approach is to model entrepreneurs' learning process as a process of Bayesian updating of priors to posterior beliefs that will guide decision-making. Those posterior beliefs will be formed through experimentation.

Here, we derive a fundamental and, we believe, underappreciated result that entrepreneurs' demand and the form of experimentation they choose will depend critically on their priors. In particular, optimistic 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 do that before returning to the issues of experimental choice in Section 6.

Let's consider an example. Recall that an entrepreneur has a prior, μ, that a venture idea will be successful, earning them V, versus not successful, earning them 0. If the cost of launching that venture were C, then the expected return from that venture given the prior of μ is:

$$R(\mu) = \mu V - C$$

If $R(\mu)$ exceeds the return of the entrepreneur's alternative opportunity (something we typically normalise to 0 in theoretical models), then without more information, an entrepreneur might choose to launch the venture.

However, 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. See Table 1. Note, here that, $\Pr[s = 1|\mu] = \lambda_1\mu + (1-\lambda_0)(1-\mu)$ and $\Pr[s = 0|\mu] = (1-\lambda_1)\mu + \lambda_0(1-\mu)$. Note that if the entrepreneur follows the signal in deciding whether to invest or not, the rate of entrepreneurial implementation (i.e., starts) will be $\Pr[s = 1|\mu]$, and the observed rate of successful ventures will be $\lambda_1\mu$. Thus, observed success will increase with μ if the entrepreneur is in control of whether the venture is launched or not following an experiment.

It is useful to orient ourselves by translating the information in Table 1 to some familiar terms.

- Sensitivity or True Positive Rate: $\lambda_1\mu$
- Specificity or True Negative Rate: $\lambda_0(1-\mu)$
- False Positive Probability: $(1-\lambda_1)(1-\mu)$
- False Negative Probability: $(1-\lambda_0)\mu$

Notice, therefore, that an experiment that involves both a higher λ_1 and λ_0 increases 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 will likely factor into such trade-offs.

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 or not, we need to calculate the posterior beliefs contingent on "good news" (i.e., $s = 1$) or "bad

Table 1: **Experiment Maps from Signal to State**

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

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

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

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

Given this, the expected return, contingent on the signal realisations, are:

$$R(\tilde{\mu}(1)) = \tfrac{\lambda_1 \mu}{\lambda_1 \mu + (1-\lambda_0)(1-\mu)} V - C$$

$$R(\tilde{\mu}(0)) = \tfrac{(1-\lambda_1)\mu}{(1-\lambda_1)\mu + \lambda_0(1-\mu)} V - C$$

This allows us to calculate the *necessary* conditions for the signal to be useful (even if c is arbitrarily low). Intuitively, a signal is only potentially useful if it leads to a change in decision. In this case, "good news" should lead to a decision to launch the venture, i.e., $R(\tilde{\mu}(1)) > 0$ and "bad news" should lead to a decision to abandon the venture, i.e., $R(\tilde{\mu}(0)) < 0$. Combining these two conditions implies that:

$$\tfrac{\lambda_1 \mu}{\lambda_1 \mu + (1-\lambda_0)(1-\mu)} V > C > \tfrac{(1-\lambda_1)\mu}{(1-\lambda_1)\mu + \lambda_0(1-\mu)} V$$

which implies that:

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

$$\Leftrightarrow \lambda_1 \lambda_0 (1-\mu) > (1-\lambda_1)(1-\lambda_0)(1-\mu) \Leftrightarrow \lambda_1 + \lambda_0 > 1$$

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.[3] Note that this is independent of prior beliefs, μ. 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. There are two cases of interest that we will describe as 'optimistic' and 'pessimistic', respectively. In the optimistic case, $R(\mu) > 0$, implying that, without information, the entrepreneur will launch the venture. In this case, the entrepreneur only finds it worthwhile to acquire information if:

$$(\lambda_1 \mu + (1-\lambda_0)(1-\mu)) \left(R(\tilde{\mu}(1)) - C\right) - c \geqslant R(\mu)$$

$$\Leftrightarrow \lambda_1 \mu V - (\lambda_1 \mu + (1-\lambda_0)(1-\mu))C - c \geqslant \mu V - C$$

[3] In the information design literature, this condition is known as "obedience"; see Bergemann & Morris (2019).

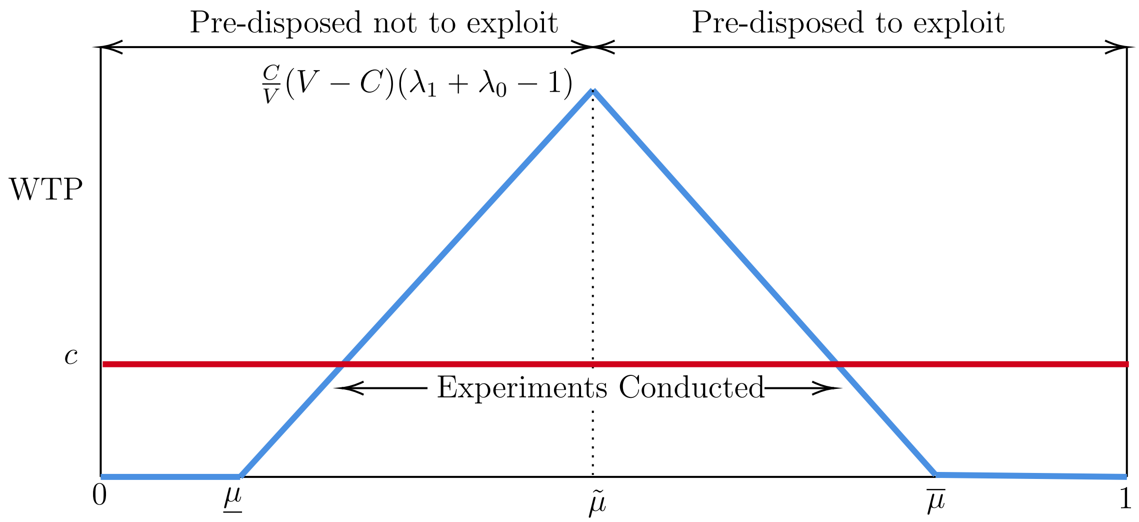

Figure 1: Willingness to Pay for Experiments

$$\Leftrightarrow \bar{\mu}(c) \equiv \mu \leqslant \frac{C\lambda_0 - c}{C\lambda_0 + (1-\lambda_1)(V-C)}$$

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

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

$$\Leftrightarrow \lambda_1\mu V - (\lambda_1\mu + (1-\lambda_0)(1-\mu))C - c \geqslant 0$$

$$\Leftrightarrow \underline{\mu}(c) \equiv \mu \geqslant \frac{c + (1-\lambda_0)C}{(1-\lambda_0)C + \lambda_1(V-C)}$$

Intuitively, a pessimistic entrepreneur is pre-disposed not to launch and can only be persuaded by the signal to launch if they are not too pessimistic.

Given this, Figure 1 shows the willingness to pay for an experiment with parameters $\{\lambda_1, \lambda_0\}$ as a function of the prior μ.[4] Notice that the willingness to pay peaks at $\mu = \tilde{\mu}$. At this point, $\tilde{\mu}V = C$, and so an uninformed entrepreneur is indifferent between pursuing the venture or not. At this point, the entrepreneur is willing to pay $\frac{C}{V}(V-C)(\lambda_1+\lambda_0-1)$, which is the most for information that will determine whether the venture should go ahead or not 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 more experiments are conducted.

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

Thus, we see here that while priors do not play a role in determining what action to take in the face of signals that are informative (i.e., where $\lambda_1 + \lambda_0 > 1$), they play a role in determining whether entrepreneurs will gather information before launching a venture. In particular, **priors will determine entrepreneurial demand for experiments**.

Note that the expected value of the posteriors is:

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

This only equals the prior, μ, if $\frac{\mu}{1-\mu} = \frac{1-\lambda_0}{1-\lambda_1}$.[5] Kamenica & Gentzkow (2011) refer to this condition on experiments where the expected value of the posteriors equals the prior as *Bayes Plausible Experiments* and the literature on information design takes this as an assumed condition (see Bergemann & Morris (2019)).[6] However, if this condition does not hold, the signal is biased. As we will demonstrate in Section 6, 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**.

4 How are Priors formed?

As noted above, priors are heterogeneous and for entrepreneurs, their priors involve the distinction of being more "optimistic" than those of others. This raises the question of where those priors come from. Here, we examine some approaches to the formation of priors. 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 etc., which result in heterogeneous levels of self-efficacy and achievement drive (Frese & Gielnik, 2014).

Second, they derive from differences in entrepreneurs' knowledge, skills and abilities because human capital is unevenly distributed across entrepreneurs, in type and quantity (Lazear, 2004; Unger et al., 2011) whether stemming from unique knowledge of a specific matter (Shane, 2000), education (Martin et al., 2013), experience in a particular knowledge domain, or talent (Eesley & Roberts, 2012), such difference in endowment shape

[5] That is the ratio of false positives to false negatives equals the prior odds.

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

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) and their relevant institutions like universities (Tartari & Stern, 2021) and 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 & Sorenson, 2005).

Fourth, they derive from entrepreneurs' preferences (Roach & Sauerman, 2015), interests (Van de Ven et al., 2007), motivation (Baum & 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 & Covin, 2011), creativity (Amabile, 1997), imagination (Patvardhan & Ramachandran, 2020), as well as logical reasoning and causal thinking (Camuffo et al., 2023a; Felin & Zenger , 2009 and 2017). Logical reasoning and causal thinking (Pearl (2009)) are especially important because, through them, entrepreneurs can intentionally and deliberately form priors through theories (Felin & Zenger, 2017; Ehrig & Schmidt, 2022; Agarwal et al., 2023) and, hence, *choose* (Agrawal et al., 2021; Gans et al., 2019; Gans et al., 2021) which idea is worth pursuing.

4.1 Prior formation through theories

When entrepreneurs build priors using causal reasoning and thinking, they strengthen their beliefs using logical arguments (Carrol & Sorensen, 2021). They are "optimistic" about their ideas in the sense of being "logically persuaded" that what they intend to pursue is plausible and valuable.

They need to learn about both states and probabilities, which are both unknown. For example, an entrepreneur might think that the development of a *new* technology will generate a *new* market. Still, they can't coast on known states and have sparse information. 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 et al. (2023a) to illustrate how prior beliefs can be formed as "theories of value" and then tested and updated through experiments.

4.1.1 Future state spaces

In order 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 identification of attributes and causal relationships among them is 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. Entrepreneurs focus on two attributes: whether they can develop technology and whether the new market will rise. Let these two attributes and their spaces be

$$X_t = \{Y, N\} \qquad X_m = \{Y, N\}$$

where subscripts t and m stand for "technology" and "market", and Y and N for *yes* and *no*. A simple way of thinking about *attributes* is that they are random variables whose realizations are *answers to questions*. In our example, the questions the entrepreneur asks are: "will the technology be available?" and "will that new market emerge?"

Attributes define a state space as the Cartesian product of the individual attributes. In our case, the two attributes define a state space $X = X_t \times X_m$ made of four states:

$$X = \{(Y, Y), (Y, N), (N, Y), (N, N)\} \tag{1}$$

The entrepreneur is interested in defining probability distributions in this space. In particular, they are interested in the probability of the relevant subset of states of interest for their decision. In our example, this is $\Pr(Y)$, which is the sum of the probabilities corresponding to states in which the new market will exist:

$$\Pr(Y) = \Pr(Y, Y) + \Pr(Y, N) \tag{2}$$

Because the state space is *subjective*, the probabilities of the four states are also subjective and rely on sparse data. In Appendix A, we operationalize our approach using Dirichlet probability distributions. The parameters of such distributions derive from the subjective *empirical distributions* of the four states envisioned by the entrepreneur in our example[7].

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 $v(\theta)$. But entrepreneurs *do not know the probability distribution of the parameters*, either. Thus, they also form, in their head, 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 to select a set of probability distributions on which entrepreneurs can concentrate their beliefs. In our example, the

[7] n = number of observations included in the empirical distribution of the decision maker comprising both actual observations (facts and data) and pseudo-observations (opinions, conjectures or beliefs).

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 (2009), which, in our simple example is

$$X_t \longrightarrow X_m$$

and $\Pr(Y \mid Y) > \Pr(Y \mid N)$ or, in terms of parameters, $\theta_{YY} > \theta_{YN}$.

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 $P(Y)$.

4.1.2 Priors and expected values of theories

Let $\Theta \equiv \{\theta : \theta_{YY} > \theta_{YN}\}$ and $\mu_\Theta(\theta)$ be, respectively, the set of parameters consistent with the theory and the subjective probability distribution of these parameters under the state space defined by the theory. $\mu_\Theta(\theta)$ is a prior probability distribution. Each possible $\theta \in \Theta$ is assigned a prior probability that reflects the entrepreneur's subjective beliefs.[8]. Each theory has an expected value, which is *the expected value of the future states of interest under the theory*, that is

$$\mathop{\mathbb{E}}_{\theta \in \Theta}[v(\theta)] \equiv V_\Theta \equiv \int_{\theta \in \Theta} v(\theta)\mu_\Theta(\theta)d\theta \tag{3}$$

where $v(\theta)$ is the expected probability of state θ and μ_Θ is the prior probability distribution subjectively defined by the entrepreneur on the set of parameters Θ of the theory. 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 $\mu_\Theta(\theta)$, makes the entrepreneur more "optimistic" because -thanks to their theory- they are persuaded that the states of interest are more likely to occur. This depends on their reasoning, i.e. 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.[9] In this way, entrepreneurs build, through logical thinking and causal reasoning, prior beliefs about an idea or opportunity.

[8]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 method suggested in Bayesian statistics (see Kass & 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

[9]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.*

Other things equal, larger V_Θ mean that entrepreneurs "believe" in their idea. Given their individual characteristics and the context in which they operate, such belief is grounded on their knowledge, the choice of attributes, causal reasoning and logic.

In our example, in the entrepreneur's mind, the potential development of the new technology raises the probability of the emergence of a new market. 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.

4.1.3 Priors on priors, Null Hypothesis and Expected Value of Theories

Entrepreneurs might be aware that the future state space they build their prior beliefs on might not be true. This equals 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 & 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." 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' attribute and causal link* or, equivalently, to the probability the new market will rise under all other imaginable alternative theories.

In our example, this occurs when $\theta_{YY} = \theta_{YN}$, which corresponds to the case in which the development of the new technology does not affect the emergence of the new market. In other words, the entrepreneur's theory is "false." This yields a different set of parameters $\tilde{\Theta} = \{\theta : \theta_{YY} = \theta_{YN}\}$, and it is easy to see from (A2) that this implies $v(\theta) = \theta_{YN}$. The expected value of the theory under the null hypothesis is then

$$\mathop{\mathbb{E}}_{\theta \in \tilde{\Theta}}[v(\theta)] \equiv V_{\tilde{\Theta}} \equiv \int_{\theta \in \tilde{\Theta}} v(\theta) \mu_{\tilde{\Theta}}(\theta) d\theta \tag{4}$$

Note, again, that $\tilde{\Theta}$ is a different parameter set consistent with any alternative theory that the entrepreneur may have. It is a different set of attributes or causal links, a different Bayesian network and probability distribution of parameters $\mu_{\tilde{\Theta}}$, that may generate the probability of the rise of the new market.

The *expected value of the theory* is then

$$V = \omega V_\Theta + (1 - \omega) V_{\tilde{\Theta}} \tag{5}$$

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.

5 Bayesian Learning Through Experiments

Given a prior belief, an entrepreneur may choose to resolve uncertainty before launching their venture. That is, they might engage in exploration before exploitation. Experiments are the ways signals of the underlying state variable are surfaced (Gans, 2023). Entrepreneurs learn about the value of opportunities or ideas by conducting experiments. Experiments are deliberate attempts to collect information about a state. They update the probability distribution $\mu_\Theta(\theta)$ of the parameters ($[v(\theta))$ under the theory to $\mu'_\Theta(\theta)$, and therefore they update V_Θ, $V_{\tilde{\Theta}}$, and V (their prior belief).

The Bayesian approach will inform how experimental information is translated into updated posterior probabilities of a venture's success. Given this, entrepreneurs will select experiments to perform. As already noted, a key insight from this approach is that the type of experiment undertaken will be determined, in part, by an entrepreneur's prior beliefs.

5.1 Experiments to Test and Update Prior Beliefs

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, i.e. 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 $\|\mu' - \mu\|$ between the two distributions is higher than a threshold μ^*, entrepreneurs also change their prio-on-prior ω to 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 μ^*, not only do V_Θ, $V_{\tilde{\Theta}}$, and V change because of the update on μ but also because entrepreneurs update ω. An entrepreneur who gains extremely evidence negative (beyond a threshold they have in mind) will not only adjust their prior μ, 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 perspective on the idea but also update their perspective 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.

Entrepreneurs can either accommodate these novel states in their theories, for example, just adding attributes and causal links and proportionally adjusting their subjective probability distributions on the parameters $\mu_\Theta(\theta)$. This approach is consistent with Reverse Bayesianism (Karni & Vierø, 2013) and does not imply a change in theory. Alternatively, contingent on moving beyond the above-defined threshold μ^* of *dynamic coherence*, entrepreneurs also update their prior-on-prior ω, questioning the whole set of current attributes and causal links and, hence, possibly changing theory. This approach is consistent with the "Hypothesis-Testing Model" (Ortoleva, 2012).

Entrepreneurs can run experiments on any subset of the parameters of their theory. They can be joint experiments when they focus on more parameters at the same time, or they can be experiments on a specific parameter.

Experiments can be conceptual or real. Conceptual experiments involve using reasoning and hypothetical observations to update and refine their understanding, while real experiments involve collecting real data or observations. Entrepreneurs can produce these observations either from quantitative data analyzed using statistical tools or by drawing qualitative information from phenomena or using informants' opinions.

After the experiment, the *expected value under the theory* (the prior belief on the opportunity or idea) changes to

$$\mathop{\mathbb{E}}_{\theta\in\Theta}[v(\theta) \mid \mu'_\Theta] \equiv V'_\Theta = \int_{\theta\in\Theta} v(\theta)\mu'_\Theta(\theta)d\theta \tag{6}$$

where the argument μ'_Θ in the expectation clarifies that this expected value now depends on the distribution μ'_Θ and not μ_Θ. 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, $V'_{\tilde{\Theta}}$, developed analogously to $V_{\tilde{\Theta}}$, using $\mu'_{\tilde{\Theta}}$ instead of $\mu_{\tilde{\Theta}}$.

After the experiment, the *expected value of the theory* becomes

$$V' = \omega' V'_\Theta + (1-\omega')V'_{\tilde{\Theta}} \tag{7}$$

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

5.2 Experimenting With Alternative Theories

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 – that is $\mathbb{E}[\mu'_\Theta(\theta)] = \mu_\Theta(\theta)$ and $\mathbb{E}[\omega'] = \omega$. 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 is $\mu_\Theta(\theta)$ and the expected updated prior-on-prior is ω.

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 $\mu_\Theta(\theta)$ and ω. 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 it. In turn, this implies that, using (5), 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.[10]

Therefore, entrepreneurs run experiments against alternative theories.[11] If entrepreneurs have an alternative theory with expected value Q (for example, an alternative opportunity, idea or strategy), then important changes in V – such as those induced by updating priors $\mu_\Theta(\theta)$ or priors on priors ω – are more likely to imply $V' < Q$ and thus a switch to the alternative theory Q. Besides, experiments that generate unexpected ("surprising") outcomes – for instance observations that represent unforeseen contingencies (Karni & Vierø, 2013 and 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, not only do entrepreneurs use experiments to test their prior beliefs about a future state space, aiming at higher V through a more favorable $\mu_\Theta(\theta)$ or higher ω, but also as a potential source of "anomalous" observations that can give rise to alternative theories (Mullainathan & Rambachan, 2023).

An alternative theory comprises a different state space characterized by different attributes or causal links. Using the relevant parameters and probabilities, the entrepreneur calculates Q, the unconditional expected value of the alternative theory, in the same way as V.

To the entrepreneur, the value of experimentation increases with the number of theories. A higher V (higher expected value of theory Θ) makes it more valuable either to experiment with it or to commit to it. If V is relatively small compared to Q, the same applies to the alternative theory. As Figure 2 shows (see Appendix A for the formal proposition), there are two *exploration zones*, one for Θ and the other one for the alternative theory.

Camuffo et al. (2023a) 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 ex-

[10]As above highlighted, the condition under which the expected value of the posteriors equals the prior corresponds to *Bayes Plausible Experiments* (Kamenica & Gentzkow, 2011). However, contingent on the level of their priors, entrepreneurs might prefer signals that are biased.

[11]We assume that entrepreneurs can compare and rank the expected value of the theories.

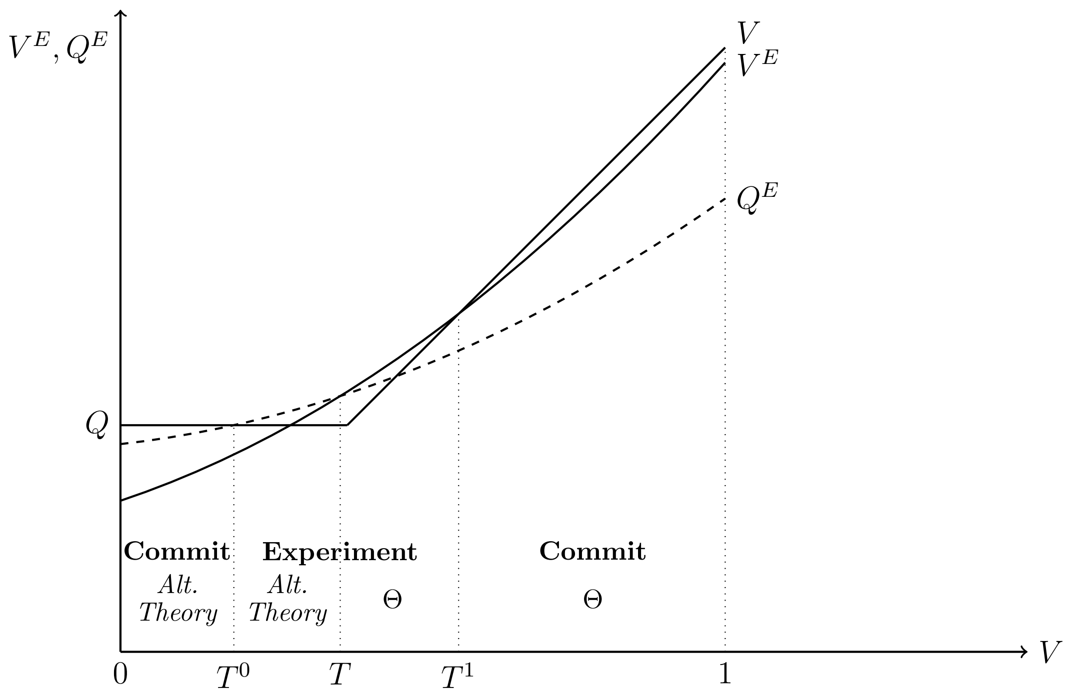

Figure 2: Value of Experimentation and Exploration Zones

perimenting 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 & Zenger's (2017) conjecture about the value of testing "contrarian beliefs". In addition, our framework is compatible with testing weaker priors – that is, theories such that, initially, $V < Q$. Thus, not only can it be optimal to test more uncertain theories but also weaker priors.

Furthermore, uncertain theories are complementary, i.e. 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.

5.3 Prior belief Updating

Bayesian entrepreneurs update their prior beliefs using Bayes theorem, i.e. calculating a posterior probability by interacting their prior with the likelihood function. After specifying the prior and the likelihood and collecting the data, entrepreneurs can obtain the posterior distribution by fitting a model to the data and hence estimate the unknown parameters of the model. As illustrated above, Bayesian entrepreneurs update their priors not only about the success or value of a specific strategy or idea but also about its generating mechanism, i.e. the underlying theory.

Such "upper level" updating can occur in two ways. The first, defined "Reverse Bayesianism" (Karni & Vierø, 2013 and 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 (1) new states that may emerge when they obtain awareness about previously unforeseen actions or consequences and (2) over 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., 2023b). We can think of "Reverse Bayesianism" as "forward-looking" Bayesianism.

The second, defined "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," further developed by Karni (2022), 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 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.

5.4 Illustrative Example

5.4.1 The Original Theory

A founder thinks they could grow a company in an emerging GAI-enhanced service market. they focus on four attributes that they think are causally linked as follows. Attribute $X_d = \{x_d\}$, where x_d is a continuous measure of demand in the targeted market. they think this depends on two attributes: $X_e = \{x_e\}$, a continuous index of the efficiency of a given GAI-related technology, and $X_s = \{x_s\}$, a continuous index of the perceived customer need for the service. The attribute $X_s = \{x_s\}$, is determined by $X_h = \{x_h\}$, which captures the spread of GAI.

We then represent the founder's theory with the following causal structure.

$$
\begin{tikzcd}
& X_e \arrow[dr] & \\
X_h \arrow[dr] & & X_d \\
& X_s \arrow[ur] &
\end{tikzcd}
$$

which rests on the following chain of subjective probabilities

$$p(x_d, x_e, x_s, x_h \mid \theta) = p(x_d \mid \theta_{des}, x_e, x_s)\ p(x_e \mid \theta_e)\ p(x_s \mid \theta_{sh}, x_h)\ p(x_h \mid \theta_h) \tag{8}$$

where $\theta = \{\theta_{des}, \theta_e, \theta_{sh}, \theta_h\}$ is the parameter set of the distributions.

In Appendix A, we work with discrete attributes and a Dirichlet distribution, which sets dichotomous realizations of the attributes. Here, we generalize the framework by working with continuous variables. In Appendix B, we show that a causal structure such as (8) generates a sequence of expected values that produces the following linear approximation of the expected value $v(\theta)$ of the state (x_d, x_e, x_s, x_h) conditional on the parameter set θ of the underlying probability distribution:

$$v(\theta) = \theta_{de}\theta_e + \theta_{ds}\theta_{sh}\theta_h \tag{9}$$

where the subscripts ij of each parameter denote that the parameter accounts for the strength of the causal link from j to i, while θ_e and θ_h denote the beliefs at the top of the causal chain regarding the quality of the encryption technology and the spread of handheld devices.

The theory is the set Θ with θ_{de}, θ_{ds}, $\theta_{sh} > 0$, and relatively high values of θ_e and θ_s. The expected value of this theory is V_Θ, defined by (3), with a probability distribution $\mu(\theta \mid \Theta)$ of these parameters. Given a null hypothesis on the parameters and a prior that the theory is true, (5) defines the unconditional expected value of this theory.

The founder is "optimistic" about this theory, i.e. they believe in it and set a high V_Θ, defined by (3) and high V, defined by (5). The founder's strong prior belief ("optimism") derives from a high $p(x_h \mid \theta_h)$. At the same time, they think they have good technology and are confident that it can be put to good use, so they expect $p(x_e \mid \theta_e)$ also to be high. The conditional probabilities associated with the causal links $p(x_d \mid \theta_{des}, x_e, x_s)$ and $p(x_s \mid \theta_{sh}, x_h)$ are also high because the available information on GAI adoption converges on indicating that the parameters are positive.

In our framework language, this is a "low-variance" theory. Since there is widespread consensus that the targeted market will rise and that there is an opportunity to grow a company providing a GAI-enhanced service. Hence, the contribution of the founder's parameters θ to the conditional expected value of the theory (the founder's prior belief) is small. In other words, other entrepreneurs looking at this opportunity are "optimistic" about it, and our founder's theory only mildly strengthens that widespread belief.

5.4.2 The Alternative Theory

Almost in parallel, the founder also thinks about alternative theories. Based on their experience, they thinks of another emerging GAI-enhanced financial service market. Hence, they comes up with another conceptual structure made of four attributes, logically connected through the following causal structure.

$$\begin{tikzcd}
& X_d \arrow[dr] & \\
X_i \arrow[dr] & & X_p \\
& X_s \arrow[ur] &
\end{tikzcd}$$

The future state of interest is attribute $X_p = \{x_p\}$, which denotes the extent to which people will buy the GAI-enhanced new financial service. This depends on two attributes: $X_d = \{x_d\}$, which measures the extent to which a secure and efficient AI application will be available, and $X_s = \{x_s\}$, which denotes whether people will perceive the need for it. In turn, $X_s = \{x_s\}$ is determined by $X_i = \{x_i\}$, which captures the extent to which GAI will be adopted by the general public. The founder's theory can be formalized as

$$p(x_p, x_d, x_s, x_i \mid \theta) = p(x_p \mid \theta_{pds}, x_d, x_s)\ p(x_i \mid \theta_i)\ p(x_s \mid \theta_{si}, x_i)\ p(x_d \mid \theta_d)$$

Like in the previous section, here we also consider for simplicity the linear approximation $v(\theta) = \theta_{pd}\theta_d + \theta_{ps}\theta_{si}\theta_i$, where again the subscripts ij of each parameter denote that the parameter accounts for the strength of the causal link from j to i, while the single subscripts account for the beliefs at the top of the causal chain.

The theory is the set Θ with θ_{pd}, θ_{ps}, $\theta_{si} > 0$, and relatively high values of θ_d and θ_i. The conditional expected value of this theory is again analogous to V_Θ, defined by (3), with a probability distribution $\mu(\theta \mid \Theta)$ of these parameters, while V, defined by (5), represents the unconditional expected value of the theory given a null hypothesis and a prior that the theory is true.

The founder's prior belief about the original theory is stronger than their prior belief about the alternative theory. they are more optimistic about the original theory, as reflected in its higher unconditional expected value. The alternative theory is more novel and, absent the founder's attributes and causal links, hard to believe. But the attributes and causal links the founder uses greatly increase, in the founder's eyes.

This is a "high-variance" theory. Since there is little information about GAI-enhanced financial services, the founder's subjective probability distributions of the attributes and causal links (parameters θ) are dispersed, and their potential contribution to the expected value of the theory (increase in founder's optimism) is large. Hence, the potential update of V from experiments is large.

As per our framework, the founder explores this theory and runs experiments to test it. For example, they test the parameters θ_{pd} and θ_{ps}, finds supporting evidence and, hence, positively updates $\mu(\theta_{pd})$ and $\mu(\theta_{ps})$. The new values $\mu'(\theta_{pd})$ and $\mu'(\theta_{ps})$ imply an updated, larger value V'_{Θ}, defined by (3), and V', defined by (5). Hence, the entrepreneur decides to explore this theory further.

6 How and Whom to Persuade

The main premise of this research agenda is that entrepreneurs are optimistic. 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 optimism. 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.

6.1 Experimental Choice for Self-Informativeness

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 3.1 by endogenising experimental choice. The experimental choice set is $E \equiv \{\lambda_1, \lambda_0 | \lambda_1 + \lambda_0 \leqslant \Lambda\}$. Here, Λ parameterises 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 3 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.

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

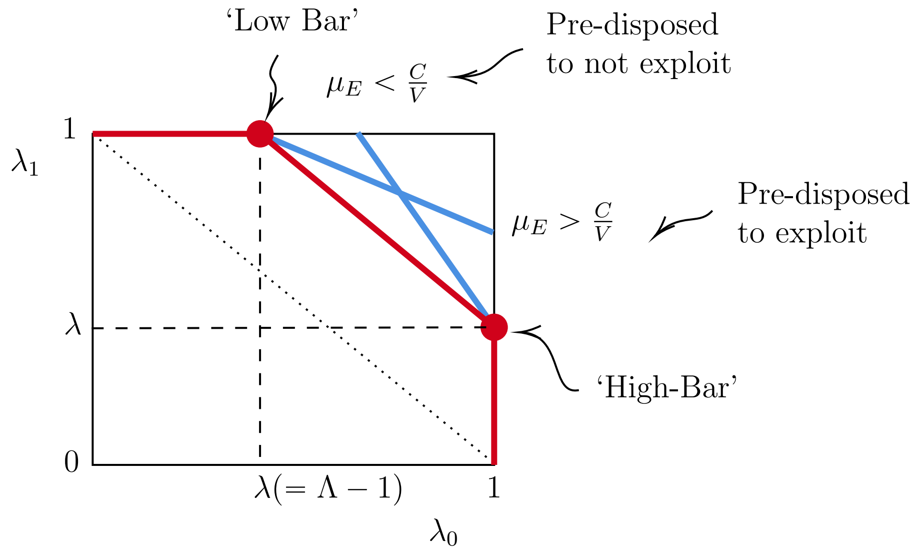

Figure 3: **Choice of Experiment**

Note that depending upon whether the entrepreneur is pre-disposed 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 pre-disposed 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 emphasised by Moore (2014). The idea is to target customer segments where you believe you can gain traction amongst 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 2 for a characterisation).

Table 2: **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
Startup Formation Rate	$(\Lambda - 1)\mu$	$\mu + (2 - \Lambda)(1 - \mu)$
Startup Failure Rate	$(2 - \Lambda)(1 - \mu)$	0
Example	Target Lead Users (Moore, 2014)	Minimum Viable Product (Ries, 2011)

By contrast, when the entrepreneur is pre-disposed 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 considered here to be 'optimistic' and, hence, have a higher prior, μ_E. Thus, we have demonstrated that:

Proposition 1. *An optimistic entrepreneur's (for whom $\mu_E > \frac{C}{V}$) optimal experiment is $\{\lambda_1, \lambda_0\} = \{1, \Lambda - 1\}$.*

The proof is in Gans (2023). The result that decision-makers should 'bias' their information signals towards their priors was most clearly demonstrated in Che & Mierendorff (2019). Their application was a dynamic process; however, it was considered biased in how media outlets presented information to users. Antecedents also focussing on media and politics while not deriving clear own-bias results include Calvert (1985), Suen (2004), Burke (2008) and Damiano, Li & Suen (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.

It is instructive to note how this approach differs from the literature on Bayesian learning. Typically, experiments are restricted to E_{BP} where the expectation of the posteriors equals the prior. However, 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.

Finally, the construction of the experimental choice set in the example analysed here is somewhat arbitrary. In particular, $\Lambda(\lambda_1, \lambda_0)$ could be a concave boundary as depicted in Figure 4. 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.

6.2 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 $\mu_E \geqslant \mu_O$. Given this, we can now define two experimental benchmarks:

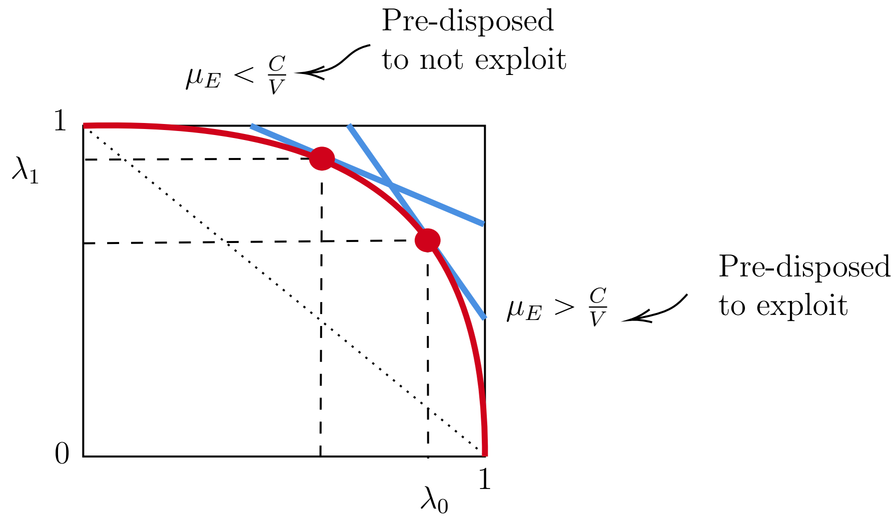

Figure 4: **Strictly Concave Experimental Frontier**

- 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 optimisation 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 $\mu_E \geqslant \frac{C}{V}$ while the investor-optimal experiment involves $\lambda_1 = 1$ if $\mu_O \geqslant \frac{C}{V}$. 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 $\mu_E > \frac{C}{V} > \mu_O$, the entrepreneur will favour an experiment with $\lambda_1 = 1$ while others will favour an experiment with $\lambda_0 = 1$.[12]

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.[13] 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.

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

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

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

Proposition 2. *Suppose that $\mu_E > \frac{C}{V} > \mu_O$. The entrepreneur chooses an experiment $\{\lambda_1, \lambda_0\} = \{\Lambda - 1, 1\}$.*

Proof. See Appendix C. ■

Intuitively, while 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, μ^* of success is. There are two interesting cases: (i) $\mu^* = \mu_E$ or (ii) $\mu^* = \mu_O$.

1. If the entrepreneur is correct ($\mu^* = \mu_E$), then following Proposition 2, the venture is launched with probability $\mu_E + (2 - \Lambda)(1 - \mu_E)$, and it never fails.
2. If the venture capitalist is correct ($\mu^* = \mu_O$), then following Proposition 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 ($\mu^* = \mu_E$), then following Proposition 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 ($\mu^* = \mu_O$), then following Proposition 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. 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 realisations. Specifically, they 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.

6.3 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 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 signalling, this newer literature allows the sender far more commitment power.

The new literature was founded by Kamenica & Gentzkow (2011) (see Kamenica (2019) for a review).[14] The basic model assumes that there is a receiver (for instance, a venture capital investor) who has a payoff function $u(a, \theta)$, which is a function of the action they take, a (say, invest or not invest) and a state of the world, θ (say, the success of the venture). The sender (or entrepreneur) has their own payoff, $\pi(a, \theta)$. 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.[15] In the entire Bayesian persuasion literature, it is always assumed that the experimental set is Bayes plausible, meaning, in our context here, that experiments are chosen from E_{BP} 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. However, a key feature

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

[15]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 & Espín-Sánchez (2022) argues, this is closer to applications and also leads to different outcomes than the full Bayesian persuasion approach.

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?[16] How does this impact other-informativeness and also on incentives to distort persuasive signals? The work of Andrews & Shapiro (2021) makes a start on this interesting question.[17]

7 Specific Research Topics

As a final section, here we outline two areas where we believe that a Bayesian approach can improve specific research areas in entrepreneurship.

7.1 A Bayesian Learning 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:

- 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

[16] One paper that has explored these differing understandings of signals is Chavda, Gans & Stern (2024).

[17] 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 & Câmara, 2016), work on persuasion with anecdotes, i.e., small samples (Haghtalab et al., 2021), non-Bayesian persuasion (de Clippel & Zhang, 2022) and price theoretic approaches (Dworczak & Martini, 2018).

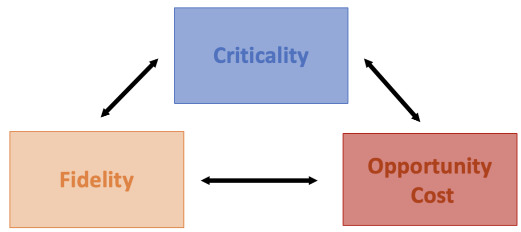

Figure 5: A Strategic Approach to Entrepreneurial Experimentation

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 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 neither yields 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, rather it is 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 allow founders to avoid leaping too quickly to a conclusion on the basis of very preliminary evidence.

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 (Murray & Tripsas (2004); Nanda & Rhodes-Kropf (2016); Cohen et al. (2019); Koning et al. (2022); among others).

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. While 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 speed is at the heart of the lean startup 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, though 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., 3-D 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 that 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, while 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 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.

7.2 The Scientific Approach to Entrepreneurial Decision Making as 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 & Zenger, 2023). Prior formation through theorizing is the first step of the Bayesian learning cycle and involves engaging in deliberate cognitive efforts to formulate decision problems as causal conceptual structures (Camuffo et al., 2023a; Ehrig & Schmidt, 2022; Felin & Zenger, 2017). Prior testing occurs through experiments which, 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 continue with the current project, pivot to a modified or different project, or terminate the venture.

The evidence coming from 15 randomized control trials involving approximately 3,000 startups around the world show that the Bayesian learning embodied in a scientific approach to entrepreneurial decision-making increases the probability of terminating projects, of efficient pivoting and, conditional on survival, leads to higher performance (revenues) (Camuffo et al., 2020 and 2024). The evidence also shows that theories and experiments have separate, identifiable effects (Agarwal et al., 2023), and that "scientific" entrepreneurs better navigate uncertainty accounting for unknown events (Camuffo et al., 2023b). As previously highlighted, this occurs in ways compatible with Reverse Bayesianism (Karni & Vierø, 2013 and 2017) and theory-based decision making (Ortoleva, 2012; Karni, 2022).

The birth and early developments of ØSense –one of the startups 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 originally 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, 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 or not. Hence, he set out to run an experiment in the form of field interviews with potential customers to test his priors. He ran several interviews at vintage fairs – a higher likelihood of individuals inclined towards reuse/rental of used goods– and set ex ante a threshold of 60% 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 AI 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 well 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.

8 Conclusions and Next Steps

TBD

References

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

Agarwal, R., Bacco, F., Camuffo, A., Coali, A., Gambardella, A., Msangi, H., ... & Wormald, A. (2023). Does a Theory-of-Value Add Value? Evidence from a Randomized Control Trial with Tanzanian Entrepreneurs. Evidence from a Randomized Control Trial with Tanzanian Entrepreneurs Available at SSRN: http://dx.doi.org/10.2139/ssrn.4412041.

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

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

Alvarez, S A. & Porac, J. (2020). Imagination, Indeterminacy, and Managerial Choice at the Limit of Knowledge. *Academy of Management Review*, 45(4), 735-744.

Amabile, T.M. (1997). Entrepreneurial creativity through motivational synergy. *The journal of creative behavior*, 31(1), 18-26.

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

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

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

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

Aunlann, R. (1976), Agreeing to disagree, *Annals of Statistics*, 4, 1236- 1239.

Aunnann, R. and A. Brandenburger (1995), Epistemic conditions for Nash equilibrium, *Econometrica*, 63, 1 16 1 - 1 180.

Ball, I. & Espín-Sánchez, J.A., 2022. Experimental persuasion. *mimeo.*, MIT.

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

Baum, J.R., & Locke, E.A. (2004). The relationship of entrepreneurial traits, skill, and motivation to subsequent venture growth. *Journal of applied psychology*, 89(4), 587.

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

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), 03.

Blackwell, D., 1953. Equivalent comparisons of experiments. *The annals of mathematical statistics*, pp.265-272.

Blume, B.D., & Covin, J.G. (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), 137-151.]

Burke, J. (2008). Primetime Spin: Media Bias and Belief Confirming Information. *Journal of Economics and Management Strategy*, 17(3), 633–665.

Calvert, R. L. (1985). The Value of Biased Information: A Rational Choice Model of Political Advice. *The Journal of Politics*, 47(2), 530–555.

Camuffo, A., Cordova, A., Gambardella, A. & Spina, C. (2020). A Scientific Approach to Entrepreneurial Decision-Making: Evidence from a Randomized Control Trial. *Management Science*, 66(2), 564-586.

Camuffo, A., Gambardella, A. & Pignataro, A. (2023a). Theory-Driven Strategic Management Decisions. *CEPR Discussion paper n. 17664-3.*

Camuffo, A., Gambardella, A. & Messinese, D. (2023b). "Bayesianism and Unforeseen Events: Empirical Evidence from Two Field Experiments". *Academy of Management Proceedings* 2023 (1).

Camuffo, A., Gambardella, A., Messinese, D., Novelli, E., Paolucci, E. & Spina, C. (2024). A Scientific Approach to Entrepreneurial Decision Making: Large Scale Replication and Extension. *Strategic Management Journal.*

Carroll, G. R. & Sorensen, J.B. (2021) *Making Great Strategy: Arguing for Organizational Advantage*. Columbia University Press, New York.

Cassiman, B. and Ueda, M., 2006. Optimal project rejection and new firm start-ups. *Management science*, 52(2), pp.262-275.

Chavda, Ankur and Gans, Joshua S. and Stern, Scott (2024) Theory-Driven Entrepreneurial Search. Available at SSRN: https://ssrn.com/abstract=4706860.

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

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

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

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

Delgado, M., Porter, M.E., & Stern, S. (2010). Clusters and entrepreneurship. *Journal of economic geography*, 10(4), 495-518.

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

Dworczak, P. and Martini, G. (2018). The Simple Economics of Optimal Persuasion. *Journal of Political Economy*. https://doi.org/10.1086/701813

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

Ehrig, T. & Schmidt, J. (2022). Theory-Based Learning and Experimentation: How Strategists Can Systematically Generate Knowledge at the Edge Between the Known and the Unknown. *Strategic Management Journal*, 43(7), 1287-1318.

Eisenmann, T.R., Ries, E., & Dillard, S. (2012). Hypothesis-driven entrepreneurship: The lean startup. Harvard business school entrepreneurial management case, (812-095).

Evans, D. S., & Jovanovic, B. (1989). An estimated model of entrepreneurial choice under liquidity constraints. *Journal of political economy*, *97*(4), 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. & Zenger, T.R. (2009). Entrepreneurs as theorists: on the origins of collective beliefs and novel strategies. *Strategic Entrepreneurship Journal*, 3(2), 127-146.

Felin, T. & Zenger, T.R. (2017). The Theory-Based View: Economic Actors as Theorists. *Strategy Science*, 2(4), 258-271.

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

Frese, M., & Gielnik, M.M. (2014). The psychology of entrepreneurship. *Annu. Rev. Organ. Psychol. Organ. Behav.*, 1(1), 413-438.

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

Gans, J.S. (2022). A Theory of Visionary Disruption. *Working Paper*, NBER.

Gans, J.S. (2023). Experimental Choice and Disruptive Technologies. *Management Science*, 69(11), pp.7044-7058..

Gans, J.S. (2025) *The Microeconomics of Artificial Intelligence*, MIT Press: Cambridge, forthcoming.

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

Gans, J.S., Scott, E.L. & Stern. S. (2024), *Entrepreneurship: Choice and Strategy*, Norton: New York.

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

Geanakoplos, J. & M. Polemarchakis (1982), We can't disagree forever, *Journal of Economic Theory*, 26,363-390.

Gimeno, J., Folta, T.B., Cooper, A.C., & Woo, C.Y. (1997). Survival of the fittest? Entrepreneurial human capital and the persistence of underperforming firms. *Administrative Science Quarterly*, 750-783.

Guzman, J., Oh, J.J., & Sen, A. (2020). What motivates innovative entrepreneurs? Evidence from a global field experiment. *Management science*, 66(10), 4808-4819.

Haghtalab, N., Immorlica, N., Lucier, B., Mobius, M. & Mohan, D., (2021). Persuading with anecdotes, *Working Paper*, (No. w28661). National Bureau of Economic Research.

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

Harsanyi, J.C. (1967/68). 'Games with Incomplete Information Played by 'Bayesian' Players', Parts I, II and III, *Management Science*, 14: 159- 182

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

Hsieh, T. (2010). Zappos's CEO on Going to Extremes for Customers. *Harvard business review*, *88*(7-8), 41-45.

Hurst, E., & Lusardi, A. (2004). Liquidity constraints, household wealth, and entrepreneurship. *Journal of political Economy*, *112*(2), 319-347.]

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

Jakobsen, A.M. (2023) Coarse Bayesian Updating, *mimeo.*, Northwestern.

Jaynes, E.T. (1968). Prior probabilities. *IEEE Transactions on systems science and cybernetics*, *4*(3), 227-241.

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

Kamenica, E. & Gentzkow, M. (2011). Bayesian persuasion. *American Economic Review* 101(6), 2590-2615.

Karni, E. (2022). A Theory-Based Decision Model. *Journal of Economic Theory*, 201, 105444.

Karni, E. & Vierø, M.L. (2013). "Reverse Bayesianism": A Choice-Based Theory of Growing Awareness. *American Economic Review*, 103(7), 2790-2810.

Karni, E. & Vierø, M.L. (2017). Awareness of Unawareness: A Theory of Decision Making in the Face of Ignorance *Journal of Economic Theory*, 168(C), 301-328.

Kartik, N., F. Xu Lee, and W. Suen (2021). Information Validates the Prior: A Theorem on Bayesian Updating and Applications. *American Economic Review: Insights*, 3 (2): 165-82.

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

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

Koning, R., Hasan, S. & Chatterji, A. (2022). Experimentation and Startup Performance: Evidence from A/B Testing. *Management Science*, 68(9): 6434-6453.

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

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

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

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

Manso, G. (2016). Experimentation and the Returns to Entrepreneurship. *The Review of Financial Studies*, *29*(9), 2319-2340.

Marinacci, M. (2015). Model Uncertainty. *Journal of the European Economic Association*, 13(6), 1022-1100.

Martin, B.C., McNally, J.J., & Kay, M.J. (2013). Examining the formation of human capital in entrepreneurship: A meta-analysis of entrepreneurship education outcomes. *Journal of business venturing*, 28(2), 211-224.

Masters, B., & 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, Inc.".

McClelland, D. C. (1961). *Achieving society.* New York, Simon and Schuster.

Milgrom P. & Roberts P. (1990). Rationalizability, Learning, and Equilibrium in Games with Strategic Complementarities. *Econometrica*, 58(6), 1255-1277.

Minniti, M., & Bygrave, W. (2001). A dynamic model of entrepreneurial learning. *Entrepreneurship theory and practice*, 25(3), 5-16.

Moore, G.A. 2014. *Crossing the Chasm*, Harper Business.

Moscarini G. & Smith L. (2001). The Optimal Level of Experimentation. *Econometrica*, 69(6), 1629-1644.

Mullainathan, S., & Rambachan, A. (2023). From Predictive Algorithms to Automatic Generation of Anomalies. Available at SSRN 4443738.

Murray, F. and Tripsas, M., 2004. The exploratory processes of entrepreneurial firms: The role of purposeful experimentation. In *Business strategy over the industry lifecycle* Emerald Group Publishing Limited, pp.45-75.

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

Nanda, R. & 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), 2410–2436.

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

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

Pearl, J. (2009). *Causality.* Cambridge University Press, Cambridge UK.

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

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

Robson, Arthur J., and Larry Samuelson (2011). The evolutionary foundations of preferences. *Handbook of Social Economics* 1, North-Holland: 221-310.

Rosenberg, N. (1992). Economic experiments. *Industrial and Corporate change*, *1*(1), 181-203.

Samuelson, L. (2004). Modeling knowledge in economic analysis. *Journal of Economic Literature*, 42(2), 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), 243-263.

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

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

Stevenson, H. H. (1983). *A perspective on entrepreneurship* (Vol. 13). Cambridge, MA: Harvard Business School.

Stuart, T.E., & Sorenson, O. (2005). Social networks and entrepreneurship. *Handbook of entrepreneurship research: Interdisciplinary perspectives*, 233-252.

Suen, W. (2004). The Self-Perpetuation of Biased Beliefs. *The Economic Journal*, 114(495), 377–396.

Sull, D.N. (2004) Disciplined entrepreneurship. *MIT Sloan Management Rev.* 46(1):71.

Tartari, V., & Stern, S. (2021). More than an ivory tower: The impact of research institutions on the quantity and quality of entrepreneurship (No. w28846). National Bureau of Economic Research.

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

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

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

Van den Steen, E. (2004). Rational overoptimism (and other biases). *American Economic Review*, 94(4), 1141-1151.

Van den Steen, E. (2005). Organizational beliefs and managerial vision. *Journal of Law, Economics, and organization*, 21(1), 256-283.]

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

Yildiz, M., (2000), *Essays on Sequential Bargaining*, PhD Dissertation, Stanford.

Zellweger, T.M. & Zenger, T.R. (2023). Entrepreneurs as Scientists: A Pragmatist Approach to Producing Value Out of Uncertainty. *Academy of Management Review*, 48(3): 379-408.

Appendix A

Models, Causal Links, and Theories

To operationalize our approach, we need to work with specific probability distributions. Let's assume that entrepreneurs' probabilities are distributed as Dirichlet with parameters $n_{IJ} > 0$, $I = Y, N$; $J = Y, N$:

$$\propto P(Y,Y)^{n_{YY}-1} \cdot P(Y,N)^{n_{YN}-1} \cdot P(N,Y)^{n_{NY}-1} \cdot P(N,N)^{n_{NN}-1} \tag{A1}$$

In a Dirichlet distribution $\mathbb{E}P(I,J) = \dfrac{n_{IJ}}{n}$, 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[18].

It helps to redefine these parameters as $\theta_{YY} \equiv \frac{n_{YY}}{n_{YY}+n_{NY}}$, $\theta_{YN} \equiv \frac{n_{YN}}{n_{YN}+\eta_{NN}}$, and $\theta_Y \equiv \frac{n_{YY}+n_{NY}}{n}$, 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: 1037), we can then write the *statistical model*

$$\mathbb{E}P(Y) \equiv v(\theta) = \theta_{YY}\theta_Y + \theta_{YN}\,(1-\theta_Y) = \theta_{YN} + \theta_Y(\theta_{YY} - \theta_{YN}) \tag{A2}$$

Entrepreneurs acknowledge that their entrepreneurial decisions (pursue or not the opportunity) are not affected only by the randomness of their subjective probabilities $P(I,J)$, but also by the uncertainty about the parameters of their probability distributions (A1). 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 (A1) identified by the different θ in the space of all the possible realizations of the three parameters in the example.

Extent of Experimentation

We model the extent of entrepreneurial experimentation using models of optimal information acquisition (Moscarini & Smith, 2001; Camuffo et al., 2023a). 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 $H(V')$ the cumulative probability distribution of the update V' of V from the experiment on Θ. Specifically, V' is the expected value of committing to Θ after the experiment or of running a new experiment on Θ.

[18] n = number of observations included in the empirical distribution of the decision maker comprising both actual observations (facts and data) and pseudo-observations (opinions, conjectures or beliefs).

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

$$V^E \equiv \rho\left[\int_{Q^*}^{1} V' dH(V') + Q^* H(Q^*)\right] - c > \Pi_V \tag{A3}$$

where $Q^* \equiv max(Q, Q^E)$, $\Pi_V \equiv max(V, Q^*)$, and Q^E is the equivalent value of the experiment on the alternative theory. This condition says that: 1) if after the experiment entrepreneurs observe $V' > Q^*$, they will either commit to Θ or run a new experiment on Θ obtaining V'; 2) otherwise, they either commit to or run an experiment on the alternative theory and obtain Q^*[19]. The key insight is that entrepreneurs keep experimenting till the present value of the experiment is higher than all the other options V , Q, and Q^E.

Proposition 3. *There are three thresholds such that: a) If $V \leqslant T^0$, entrepreneurs commit to the alternative theory. If $V \in (T^0, T)$, they run an experiment on the alternative theory. If $V \in (T, T^1)$, they run an experiment on Θ. If $V \geqslant T^1$, they commit to Θ and obtain V*

Proof. Using (A3), and the fact that $\mathbb{E}V' = V$, rewrite V^E as

$$V^E \equiv \rho\left[V + Q^* H(Q^*) - \int_{0}^{Q^*} V' dH(V')\right] - c = \rho\left[V + \int_{0}^{Q^*} H(V') dV'\right] - c \tag{A4}$$

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

We first establish that $\dfrac{\partial V^E}{\partial V} > \dfrac{\partial Q^*}{\partial V} \geqslant 0$. Let $D = 1 - \rho^2 H(Q^E) K(V^E)$ where K is the cumulative distribution of the update Q' of Q if decision-makers run an experiment on the alternative theory. Take the differentials dV^E, dQ^*, and dV in (A4) and, if $Q^* = Q^E$, in the equivalent expression for Q^E. Solving the system, or setting $dQ^* = 0$ if $Q^* = Q$, obtain $\dfrac{\partial V^E}{\partial V} = D^{-1}\rho > \dfrac{\partial Q^*}{\partial V}$, which is equal to 0 or $D^{-1}\rho^2 K(V^E)$ depending on whether $Q^* = Q$ or $Q^* = Q^E$.

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

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

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

Appendix B

A causal structure such as (8) generates the following sequence of expected values

$$\mathbb{E}(x_d \mid \theta_{des}, x_e, x_s) \equiv v_d(\theta_{des}, x_e, x_s) = \int_{X_d} x_d \; p(x_d \mid \theta_{des}, x_e, x_s) \; dx_d$$

$$\mathbb{E}\left[v_d(\theta_{des}, x_e, x_s) \mid \theta_e, \theta_{sh}, x_h)\right] \equiv v_{des}(\theta_{des}, \theta_e, \theta_{sh}, x_h) =$$

$$\int_{X_s}\int_{X_e} v_d(\theta_{des}, x_e, x_s) \; p(x_e \mid \theta_e) \; p(x_s \mid \theta_{sh}, x_h) \; dx_e \; dx_s$$

$$\mathbb{E}\left[v_{des}(\theta_{des}, \theta_e, \theta_{sh}, x_h) \mid \theta_h\right] \equiv v(\theta) = \int_{X_h} v_{des}(\theta_{des}, \theta_e, \theta_{sh}, x_h) \; p(x_h \mid \theta_h) \; dx_a$$

Consider for simplicity the following linear approximation.

$$v_d(\theta_{des}, x_e, x_s) = \theta_{de} x_e + \theta_{ds} x_s, \qquad \text{with} \quad \mathbb{E}(x_e \mid \theta_e) = \theta_e, \quad \mathbb{E}(x_s \mid \theta_{sh}, x_h) = \theta_{sh} x_h$$

where we distinguish between the two elements θ_{de} and θ_{ds} of the vector of parameters θ_{des} that represent, respectively, the correlations between x_e and x_d and x_s and x_d. By replacing the two expected values in $v(\cdot)$, we obtain

$$v_{des}(\theta_{des}, \theta_e, \theta_{sh}, x_h) = \theta_{de}\theta_e + \theta_{ds}\theta_{sh} x_h$$

and

$$v(\theta) = \theta_{de}\theta_e + \theta_{ds}\theta_{sh}\theta_h$$

which is equation (9) in the text.

Appendix C

Proof of Proposition 1

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

$$(1-\alpha)\tfrac{\lambda_1\mu_O}{\lambda_1\mu_O+(1-\lambda_O)(1-\mu_O)}V \geqslant C$$

$$\implies \alpha \leqslant \tfrac{\lambda_1\mu_O V-(\lambda_1\mu_O+(1-\lambda_O)(1-\mu_O))C}{\lambda_1\mu_O V}$$

The entrepreneur chooses an experiment to maximise:

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

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

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

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:

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

For $\lambda_1 = \Lambda - 1$ we have:

$$V - C$$

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

$$\mu_O(V-C) \geqslant \tfrac{\mu_E}{\mu_E+(2-\Lambda)(1-\mu_E)}(\mu_O(V-C)-(2-\Lambda)(1-\mu_O)C)$$

which can be easily shown to hold.