NBER WORKING PAPER SERIES

KNIGHTIAN UNCERTAINTY AND BAYESIAN ENTREPRENEURSHIP

Joshua S. Gans

Working Paper 33507
http://www.nber.org/papers/w33507

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
February 2025

Joshua Gans has drawn on the findings of his research for both compensated speaking engagements and consulting engagements. He is also chief economist of the Creative Destruction Lab, a University of Toronto-based program that helps seed-stage companies, from which he receives compensation. He conducts consulting on cryptocurrency and payments matters with an association with Keystone and his ownership of Core Economic Research Ltd. He also has equity and advisory relationships with a number of startup firms. The views expressed herein are those of the author and do not necessarily reflect the views of the National Bureau of Economic Research.

The author has disclosed additional relationships of potential relevance for this research. Further information is available online at http://www.nber.org/papers/w33507

Knightian Uncertainty and Bayesian Entrepreneurship
Joshua S. Gans
NBER Working Paper No. 33507
February 2025
JEL No. D81, O30

ABSTRACT

This paper examines the relationship between Knightian uncertainty and Bayesian approaches to entrepreneurship. Using Bewley's formal model of uncertainty and incomplete preferences, it demonstrates that key predictions from Bayesian entrepreneurship remain robust when accounting for Knightian uncertainty, particularly regarding experimentation, venture financing, and strategic choice. The analysis shows that while Knightian uncertainty creates a more challenging decision environment, it maintains consistency with the three pillars of Bayesian entrepreneurship: heterogeneous beliefs, stronger entrepreneurial priors, and Bayesian updating. The paper also explores connections to effectuation theory, finding that formal uncertainty models can bridge different entrepreneurial methodologies.

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

1 Introduction

New ideas are a domain where Knightian uncertainty naturally applies. Ideas are often contained within the mind of a single person and have yet to be explored or tested in any way. The newness of ideas encompasses a new event or state that was not conceived by others, let alone a probability assessment of their reasonableness or future prospects. One recent example was the discovery of generative AI by Vaswani et al. (2017). This new methodology led directly to ChatGPT, one of the most successful products in terms of adoption speed in history. ChatGPT, when it was released in November 2022, was a surprise advance even to those familiar with the scientific advance of five years earlier. Thus, for those working in this area, it is reasonable to propose that they faced Knightian uncertainty regarding whether large language models (or LLMs) would be of practical use prior to November 2022, which was resolved substantially shortly afterwards.

Knightian uncertainty is, thus, intimately tied to entrepreneurial research, with Knight (1921) himself proposing the famous distinction between risk and uncertainty as part of a theory to explain the existence of entrepreneurial profit. At the same time, Knightian uncertainty has presented a challenge to models of choice. Unlike the clear axiomatic foundations that underlie Bayesian decision-making, the presence of Knightian uncertainty undermines those axioms and, with it, the ability of such theories to generate empirical predictions, let alone robust predictions. In its wake, many entrepreneurial researchers forgo deductive formal reasoning and instead base analyses on inductive examinations of entrepreneurial choice processes.[1]

This paper challenges the notion that taking Knightian uncertainty seriously will undermine the ability to use formal models in entrepreneurial research. The context here is based on the agenda laid out by Agrawal et al. (2025) called Bayesian Entrepreneurship that seeks to motivate and ground research questions in a Bayesian decision-making approach. Specifically, that research agenda rests on three premises: (i) that important agents in entrepreneurial environments have heterogeneous beliefs; (ii) that entrepreneurs hold stronger (some might say, more optimistic) beliefs regarding venture success than other agents; and (iii) that all agents use Bayes rule to update their priors upon receipt of new information. They argue that this can lead to a number of fruitful, positive examinations of entrepreneurial choice as well as provide normative prescriptions that will improve entrepreneurial success.

[1] Bhidé (2024) argues that Knightian uncertainty has been neglected in formal economic research, although I would contend that this is, in part, caused by the difficulty of integrating it into formal models of any kind. That said, many formal decision theorists have examined these issues, particularly ambiguity; see the review by Gilboa and Marinacci (2013). As Packard et al. (2017) note, by launching a venture, uncertainty is eventually resolved and thus, it is fruitful to examine the changing nature of uncertainty as part of the process of studying entrepreneurial decision-making.

What this paper shows is that the approaches and important empirical conclusions from the Bayesian entrepreneurship agenda are, and will continue to be, robust to a consideration of Knightian uncertainty. This is done by adopting the formal approach to Knightian uncertainty proposed by Bewley (2002) and studied by numerous economists to understand the impact of such uncertainty on financial markets (Rigotti and Shannon, 2005) and incentive contracts (Lopomo et al., 2011). That approach adopts a strong form of uncertainty aversion but has not hitherto been explored in the published literature on entrepreneurship.

Section 2 describes the Bewley (2002) approach and establishes its operation in a simple entrepreneurial context. Section 3 then examines three aspects of the Bayesian entrepreneurship agenda – the role of experiments, venture financing and entrepreneurial strategy – to demonstrate that the empirical predictions of the Bayesian approach are robust to the inclusion of Knightian uncertainty. The reason for this is that while Knightian uncertainty generates a more challenging decision environment, it is still amendable to a consideration of the three pillars of the Bayesian approach listed above.

Section 4 then turns to consider the methodology in entrepreneurial research that, in recent years, has been most grounded on a foundation of Knightian uncertainty; effectuation theory as developed by Sarasvathy (2001). While not providing a comprehensive examination of that theory, the discussion demonstrates that it is difficult, although not impossible, to generate two key predictions of effectuation theory – that entrepreneurs will make choices based on resources and relationships at hand and that they will base decisions on affordable loss rather than a full rate of return analysis – to be generated from the same set of underlying assumptions. That said, the analysis shows that using a formal model of Knightian uncertainty can act as a bridge to help clarify some of the distinctions between the methodologies of effectuation and those of Bayesian entrepreneurship.

A final section concludes.

2 Formal Representation of Knightian Uncertainty

It is with some trepidation that anyone purports to consider a formal treatment of Knightian uncertainty, especially in the field of entrepreneurship. The term itself is the subject of much controversy. My intention here is not to wade into that controversy but, instead, to adopt what I regard as a well-grounded, explored and meaningful formal definition of uncertainty and explore its implications.

2.1 Origins of the Concept

Nonetheless, it is useful to understand the origins of the formal treatment I will adopt here. The term "Knightian uncertainty" has its origins with the distinction of Frank Knight (1921) between risk (where probabilities have an objective foundation, such as might arise in gambling games or from data-rich statistical estimations) and uncertainty (where that isn't the case).

> Uncertainty must be taken in a sense radically distinct from the familiar notion of risk, from which it has never been properly separated.... The essential fact is that 'risk' means in some cases a quantity susceptible of measurement, while at other times it is something distinctly not of this character; and there are farreaching and crucial differences in the bearings of the phenomena depending on which of the two is really present and operating.... It will appear that a measurable uncertainty, or 'risk' proper, as we shall use the term, is so far different from an unmeasurable one that it is not in effect an uncertainty at all. Knight (1921), p.19.

Keynes (1921) made a similar distinction.[2] Knight's purpose, however, was directly relevant to the study of entrepreneurship. He related his distinction to the existence of markets that would allow entrepreneurs and businesses to purchase insurance. He contended that they could do so for risks but not for uncertainty.[3] As a consequence, Knight viewed entrepreneurs as being agents who bore uncertainty and that the supra-normal profits (if any) generated were as compensation that bearing.

2.2 A Formal Definition

Formal decision theory emerged after the work of Knight, mostly notably with Von Neumann and Morgenstern (1944) and Savage (1954), who gave expected utility theory an axiomatic foundation. Anscombe and Aumann (1963) then extended this research program to encompass subjective utility, which, in effect, rendered some of Knight's concerns moot. That said, Aumann (1962) noted that if a decision-maker was unable to form even subjective utilities

[2]See Packard et al. (2021) for a discussion. Recently, Kay and King (2020) took this further and focused on "radical uncertainty," which involves situations where not all outcomes can be described and consequently, probabilities cannot even be subjectively formed. This is also akin to Sarasvathy and Kotha (2001)'s third urn.

[3]The reasons Knight gave were arguably related to the standard reasons for the non-existence of insurance markets grounded in adverse selection and moral hazard; see LeRoy and Jr. (1987). That said, others, notably Bhidé (2021), disagree. My purpose here is not to enter the fray as to what Knight meant but to use his discussion as a springboard to the examination of uncertainty for entrepreneurial choice.

over states, their preference ordering would be incomplete and, therefore, incapable of being represented by a utility function or something similar.[4] This meant that when faced with Knightian uncertainty, there was an important gap in decision theory and its ability to generate behavioural hypotheses, including those that could be tested with data using revealed preference.

Meanwhile, work empirically testing various forms of decision theory revealed that some of these gaps may arise in common situations, notably, Allais (1953) and Ellsberg (1961). This led to various new theories of choice that explored how to adjust the axioms of expected utility theory to explain or rationalise the empirical challenges. My purpose here is not to review these theories because, apart from the decision-making under an ambiguity approach (e.g., Gilboa and Schmeidler (1989)), few dealt head-on with Knightian uncertainty or something related.

As noted in the Introduction, my purpose here is to find a way of formally representing uncertainty, as opposed to risk, and to do so in a way that makes it capable of generating hypotheses regarding entrepreneurial choices. Given these criteria, rather than evaluate the broader merits of alternative decision-theoretic models, here I explore the so-called "unanimity" approach developed by Bewley (2002).[5] In contrast to other developments, this approach directly allowed for preferences to be incomplete in the Aumann (1962) sense. That is, when faced with uncertainty, a decision-maker was assumed not to be indifferent between uncertain outcomes but to have no preference ranking at all. In effect, the decision-maker was permitted to claim and/or believe that they "simply do not know" what to do.[6]

Bewley operationalised this in a natural way. To see this, suppose there is a decision-maker who can make a choice, x, from a set of actions, X. The decision-maker faces uncertainty with respect to a random variable, θ, that is a state realisation from a set Θ. Let $F(\theta)$ be a probability distribution function with respect to θ drawn from a set, $\mathcal{F}$. I assume the decision-maker's preference relation, $\succ$, for actions, is represented by a unique closed, convex set of probability distributions $\mathcal{F}$ and a continuous, strictly increasing, concave function $u : \mathbb{R}_+ \to \mathbb{R}$, unique up to positive affine transformations, such that:

$$x \succ y \text{ if and only if } \mathbb{E}_F[u(x, \theta)] > \mathbb{E}_F[u(y, \theta)] \text{ for all } F \in \mathcal{F}$$

Here $\mathbb{E}_F[u(x)]$ denotes the expected value with respect to the probability distribution F, and $u(x, \theta)$ is the decision-maker's payoff from choosing x if the state is θ. This definition of

[4] See also Gans (1996).

[5] This theory was originally written in 1986 but was unpublished for many years and its companion piece Bewley (1987) was never published.

[6] This is paraphrased from Keynes (1937), p.214: "About these matters there is no scientific basis on which to form any calculable probability whatever. We simply do not know."

preferences says that the decision-maker will only strictly prefer x over y if their expected payoff is strictly higher in choosing x versus y for each possible probability distribution in $\mathcal{F}$. Thus, if there is one probability distribution for which the ranking using expected payoffs is reversed, preferences between the two options will be incomplete. This also means that if we observed such a decision-maker choosing x over y, we could not say x was revealed preferred to y, but only that y was not revealed preferred to x.

This is and should be interpreted as a strong condition regarding the decision-maker's aversion to uncertainty. Uncertainty is captured by the fact that there is a set of possible probability distributions and the decision-maker cannot form preferences over actions unless it is completely clear that, regardless of the actual distribution, the ranking would be unchanged.[7] At the same time, this preference condition does not say, as can sometimes be assumed in the informal literature, that Knightian uncertainty (as seen as the absence of the ability to pin down probabilities for various states) implies that there is no room for choice. This definition allows for such choice when an option is undominated in the sense defined by the preference condition. It is this that allows us to consider the impact of uncertainty itself.

The size of $\mathcal{F}$ is itself a measure of the degree of uncertainty (and the decision-maker's uncertainty aversion). To see this, note that if $\mathcal{F}$ is a singleton, then the preference condition collapses to the standard expected utility theory preference relation where there is no uncertainty. Moreover, Rigotti and Shannon (2005) shows that as the size of $\mathcal{F}$ becomes smaller, the preferences of the decision-maker converge to the expected utility theory outcome. Thus, the size of $\mathcal{F}$ is itself a measure of preference incompleteness; that is, if $\mathcal{P}_{\succ} = \{(x, y) \in \mathbb{R}_+ \times \mathbb{R}_+ | x \succ y\}$, then if $\succ'$ is the preference order under $\mathcal{F}'$, then if $\mathcal{F}' \subset \mathcal{F}$, $\mathcal{P}_{\succ'} \subset \mathcal{P}_{\succ}$ implying that $\succ'$ is more complete than $\succ$.[8]

There is an interesting issue as to whether there are any constraints on $\mathcal{F}$? Bewley (2002) allows for it to differ between agents. But the question is whether it can constitute uncertainty unless each agent places some non-zero probability of each state over which each element of $\mathcal{F}$ is defined. That is, could a distribution that assigns a probability to a state of 1 or, indeed, of 0 be permissible? To do so would allow for an agent to act as if they had no doubt as to what might or might not occur. This would invalidate the contention of Bhidé (2024), for example, that operationally, Knightian uncertainty was a situation where agents faced some doubt as to what might occur. In what follows, I will be forced to adopt some

[7]Bewley (1987) proposed the preference condition in this strict form, but subsequent researchers have shown that it could be relaxed to include weak preferences; see the discussion in Rigotti and Shannon (2005). For example, the decision-maker could be indifferent to two actions if the expected payoff was the same under each distribution. See also Ok et al. (2012) and Gerasimou (2018).

[8]Proposition 1 of Rigotti and Shannon (2005).

practical bounds to $\mathcal{F}$. For instance, I could surely not permit entrepreneurial returns to be greater than all of the monies and wealth of Earth (although equal to it would be fine). And I would surely have to rule out those returns being 0 as an agent who acted on such a distribution or could only act as if that distributional outcome could possibly be true would be acting as if they had no doubt. The precise, reasonable bounds and implications of this are left for future research. Conveniently, as we will see, this allows me to use Bayes rule with abandon in this environment.

It is instructive to consider all of this in a (very basic) entrepreneurial example. Suppose that a potential venture can be a success or a failure; that is, $\Theta = \{\text{success}, \text{failure}\}$. Suppose that success generates a payoff of V and failure a payoff of 0 if a venture is launched (i.e., $x = 1$); i.e., $u(1, \text{success}) = V$ and $u(1, \text{failure}) = 0$. Finally, suppose that the elements of $\mathcal{F}$ can be expressed by a probability, $p = \Pr[\text{success}]$ where $p \in [\underline{p}, \bar{p}] \subset (0, 1)$; where $\underline{p} \equiv \arg\min_{p\in\mathcal{F}}$ and $\bar{p} \equiv \arg\max_{p\in\mathcal{F}}$.[9] Finally, suppose that it costs C to launch a venture; that is, $u(0, \text{success}) = u(0, \text{failure}) = -C$. Then, the preference ordering between launching ($x = 1$) and not launching ($x = 0$) is:

$$\text{launch} \succ \text{not launch}, \text{ if and only if } \underline{p}V > C$$

$$\text{not launch} \succ \text{launch}, \text{ if and only if } \bar{p}V < C$$

Note that if $\bar{p} > \underline{p}$, then it could easily be the case that neither of these inequalities hold; that is $\underline{p}V \leq C \leq \bar{p}V$. In this case, the model cannot tell us what will be chosen.

This outcome is illustrated in Figure 1. Observe that as $\bar{p} > \underline{p}$, there is a domain in (C, V) space where the entrepreneur cannot rank the actions to launch and not launch. This preference incompleteness is a consequence of uncertainty. As $\bar{p} \to \underline{p}$, this set shrinks ultimately to a simple indifference condition defined by $p^* = \frac{C}{V}$. What this also highlights is that the analyst faces an issue: since the entrepreneur will either launch or not launch the venture, given the uncertainty, what theoretical predictions can be made about which choice will be made? It is this issue that I will turn to next.

2.3 The Importance of Defaults

Allowing for preference incompleteness involves a cost; preferences alone cannot pin down likely decisions. To generate a theoretical prediction, we need an additional assumption. A natural way to achieve this is to posit that the uncertainty that emerges is the result of something "new." In the domain of entrepreneurship, this would be a new idea with the

[9]This bears some relationship to the ϵ-contamination approach of Nishimura and Ozaki (2004) although their application is markedly different.

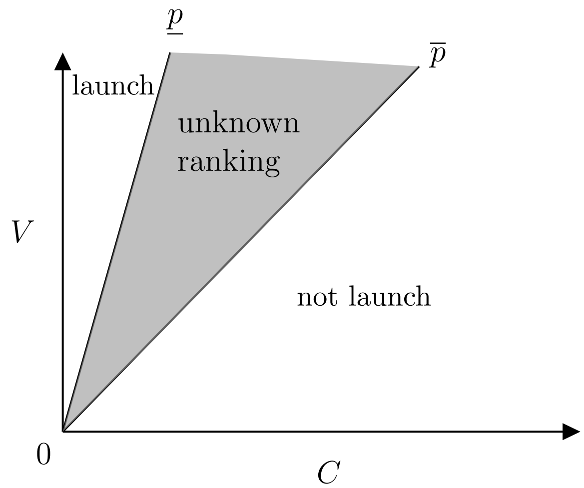

Figure 1: Entrepreneurial Preferences

status quo being whatever the decision-maker was doing prior to that emergence. This was the approach taken by Bewley (2002) when he proposed an *inertia* assumption.[10] He argued that the natural starting point was the decision-maker's choice "pre-newness." For instance, previously, the states might have had a probability of zero or the actions being considered were not available, making at least the payoff that the decision-maker was previously earning a natural benchmark from which to rank new alternatives.[11]

From this perspective, if the idea for the venture arose at a given date, then because the venture had, by definition, not been launched at that point, the appropriate default could be the payoff the entrepreneur receives if they did not launch a venture; that is, 0. Thus, in this case, the entrepreneur would only choose to launch a venture if:

$$\underline{p}V > C$$

This creates a high bar in terms of V relative to C for a venture to be launched and is consistent with the sometimes expressed notion that uncertainty is a barrier to entrepreneurial activity. This can be termed the *no-launch* (N) default.

It can immediately be seen that this is not a very satisfying argument in the domain of entrepreneurship. Entrepreneurs are often imagined to be those who are more optimistic, less risk averse and more committed to action than not (Gans et al., 2024). Indeed, as already

[10]Sutton (2006) who adopted a framework based on the unmeasurability of a subset of states used an evolutionary framework that allowed the entry of agents with a fixed action preference so long as it could be justified by one of the possible distributions.

[11]Bewley (2002) also considered bounded rationality with the status quo defined by a previous heuristic choice that can be reconsidered in light of a changed decision environment. Other researchers examining this model in light of trade viewed the status quo as involving no trade (Rigotti and Shannon, 2005) or, in a contracting environment, the individual rationality payoff of the decision-maker (Lopomo et al., 2011).

noted, the entire premise of the Bayesian Entrepreneurship agenda is that entrepreneurs have stronger priors in their assessment of the likelihood of venture success. This suggests that when a new entrepreneurial opportunity emerges, the appropriate default is a *launch* (L) default. In this case, the entrepreneur would choose to launch a venture if:

$$\bar{p}V \geq C$$

An interpretation of this is that the entrepreneur is committed to pursuing the venture in the face of uncertainty and will only be deterred from this if V is low relative to C.

This discussion implies that we need to be both explicit in our default assumptions and circumspect to the default assumption we make in analysing entrepreneurial choices. Thus, in what follows, I will often examine outcomes from both a L-default and N-default perspective.[12]

It is worth mentioning at this point that there may be alternatives to simply assuming that one action or outcome is a default. From a modelling perspective, it might be possible to parameterise the decision-maker's preferences with regard to uncertainty. Following a suggestion by Arrow and Hurwicz (1972), suppose that, in the case of the entrepreneurial launch decision, weights were placed on the best and worst distributions such that it was assumed that:

$$L \succ N \text{ if and only if } \left(\omega\underline{p} + (1-\omega)\bar{p}\right)V > C$$

for some $\omega \in [0,1]$ where $\omega = 1$ ($\omega = 0$) corresponds to the $N-$default (L-default) entrepreneurial cases. ω is a parameter that can be used to characterise uncertainty aversion on the part of the decision-maker. While this parameterisation is not used below, in some applications, it may prove fruitful for more precise empirical predictions.

2.4 Comparison with Theory-Based Decision-Making

One dimension of Bayesian entrepreneurship is where prior beliefs come from. Agrawal et al. (2025) consider developments in theory-based decision-making whereby entrepreneurs are presumed to use theoretical constructs and narratives to place more structure on probability distributions that impact entrepreneurial choice. This approach has been pushed most recently by Felin and Zenger (2017), who argue that such theories can be causal structures on stochastic processes. This examination has its formal roots in Ryall (2009) and has been

[12]This also implies that defaults will also play a role in the assessment of other innovation domains. For instance, while startups might be viewed as having an L-default, incumbents could be viewed in the opposite manner, not simply because of their preferences but also because they earn profits in the status quo; see, Gans (2023).

developed further by Camuffo et al. (2023) and Ehrig et al. (2025). Montiel Olea et al. (2022) provide an alternative but potentially important approach making use of model specification.

Theory-based decision-making differs from Knightian uncertainty in that, ultimately, the probability distribution is knowable even if its parameters are pinned down by theoretical thought processes rather than empirical evidence per se. This is useful in modelling the decision-making sophistication of agents and how this might reveal itself through observational data. For instance, Chavda et al. (2024) considers an environment where there are two potential probability distributions for venture valuation, with some agents not considering that fact and acting as if signals are from a single distribution while others understand those signals can reveal the precise distribution itself. This leads to different search behaviour and pivoting behaviour.

The Knightian uncertainty approach stands apart here in that there is no avenue for theory to identify the correct distribution or partition of distributions from the set $\mathcal{F}$. Nonetheless, that does not mean that signals cannot refine uncertainty. It is just that the process requires close attention to what is known and unknown, as I will explore next.

3 Implications for Bayesian Entrepreneurship

Having set up a formal approach to Knightian uncertainty, I now turn to consider its implications for the Bayesian Entrepreneurship research agenda, specifically, the part of that agenda that focuses on how agents learn for the purpose of dealing with uncertainty over their own venture prospects and also as a means of informing others as to a venture's prospects.

3.1 Experimentation

Experiments and information gathering are at the core of Bayesian decision-making. It is sometimes suggested, however, that when Knightian uncertainty is present, the value of such experiments and information is limited or potentially of no value at all.[13] The logic behind this view is that Bayesian updating procedures cannot be applied to events that have zero probability or where probabilities cannot be established. On the other hand, it seems unlikely that there is no ability to resolve uncertainty, especially given that Knightian uncertainty does, eventually, get resolved to some degree. For instance, an entrepreneur facing Knightian uncertainty might resolve that uncertainty by launching a venture and discovering whether it is successful or not ex-post and why.

[13] This seems implied by, for instance, Shackle (1961) and Wiltbank et al. (2006).

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

Table 1: **Experiment Maps from Signal to State,** $\Pr[s|\theta]$

This section demonstrates that, based on the formal definition of Knightian uncertainty used here, the role of experimentation and information gathering is undiminished. Indeed, it is shown that many of the predictions regarding experimental choice used in Bayesian entrepreneurship are robust and potentially sharper when Knightian uncertainty is present. The reason for this is that the structure of that uncertainty still posits a set (potentially large) of possible distributions of the state variable, and each of those distributions can be updated and modified based on experiments run. This can have the effect of changing and even reducing that set and so enable more choices by reducing inertia.

Consider again the entrepreneurial launch decision. Suppose that prior to making that decision, the entrepreneur can engage in an experiment that generates a signal, s that takes on a value of 1 with probability λ_1 if the venture will be successful and a value of 0 with probability λ_0 if the venture is unlikely to be successful; that is, $\Pr[s = 1|x = V] = \lambda_1$ and $\Pr[s = 0|x = 0] = \lambda_0$; see Table 1. This is the same signal structure that is used by Gans (2023) in a setting without Knightian uncertainty.

As Bewley (1989) noted, a signal will update every distribution in $\mathcal{F}$. This updating can be derived using Bayes' rule. In particular, for a distribution with $\Pr[\text{success}] = p_i$, the posterior distribution characterised by posterior beliefs, $p_i(s)$, will be:

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

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

for the case of "good news" (i.e., $s = 1$) and "bad news" (i.e., $s = 0$), respectively. A necessary condition for $p_i(1) \geq p_i \geq p_i(0)$ is that $\lambda_1 + \lambda_0 \geq 1$; which will be assumed to be the case.

Agrawal et al. (2025) examined how the demand for experiments changed as entrepreneurial priors regarding the success of the venture changed. Here, the question of interest is how the demand for experiments changes with the degree of Knightian uncertainty. To examine this, consider first an N-default entrepreneur who will, in the absence of an experiment, only launch a venture if $\underline{p}V > C$. In this case, an experiment will be of value if it causes the entrepreneur to launch a venture. In other words, it is the impact of the experiment on $\underline{p}(1)$ that matters.[14]

The realisation of $\Pr[s = 1]$ depends on the ground truth distribution. However, since this is unknown to the entrepreneur when considering the value of the experiment, they will use $\Pr[s = 1]$ evaluated according to the distribution $\underline{p}$. Moreover, since no launch is their default, they will compare the expected payoff if an experiment is deployed with a default payoff of 0. Then, the value of the experiment $\{\lambda_1, \lambda_0\}$ to the entrepreneur is:

$$\underline{p} \leq \tfrac{C}{N} : \Pr[s = 1]\big(\underline{p}(1)V - C\big) = \lambda_1 \underline{p}(V - C) - (1 - \lambda_0)(1 - \underline{p})C$$

$$\underline{p} > \tfrac{C}{N} : \Pr[s = 0]\big(C - \underline{p}(0)V\big) = \lambda_0(1 - \underline{p})C - (1 - \lambda_1)\underline{p}(V - C)$$

This is positive if and only if:

$$\underline{p}(0) \leq \tfrac{C}{V} < \underline{p}(1)$$

Intuitively, an experiment is valuable under Knightian uncertainty only where it potentially leads to a change in the entrepreneur's choice. When $\underline{p} \geq \frac{C}{N}$, the experiment shifts the 'unknown ranking' set to potentially allow a ranking if $\frac{C}{V}$ is below the new lower bound of the unknown set, $\underline{p}(1)$. Thus, it is in that range where an experiment is valuable. As Figure 2 shows, the experiment also remains valuable where the entrepreneur would have chosen to launch (at low $\frac{C}{N}$) even under Knightian uncertainty as it can signal not to launch in that case.

The outcome of the experiment for an L-default entrepreneur is the same but for two differences. First, the experiment can resolve uncertainty and provide a ranking where $\bar{p} \geq \frac{C}{N}$ and cause the venture to not launch where it would have launched in the absence of an experiment. Second, the experiment is grounded in the prior, $\bar{p}$, and so it's set where the experiment is of positive value is at a correspondingly higher level than for a N-default entrepreneur (see Figure 2).

What impact does the degree of Knightian uncertainty have on the demand for experiments? The analysis, thus far, has demonstrated that experiments can play a role in 'resolving uncertainty' in the sense of providing a clear ranking between alternatives where none previously existed. However, depending on the defaults of the entrepreneurs, this resolution is undirectional, with experiments playing their usual role in simply updating priors that,

[14]Given the setup of the space of distributions in this entrepreneurial context, the distribution of interest that is impacted on by the experiment is the same before and after the experiment. This is because it defines the threshold and, as Bewley (2002), showed there always is a distribution that defines the 'better than set' for the decision-maker.

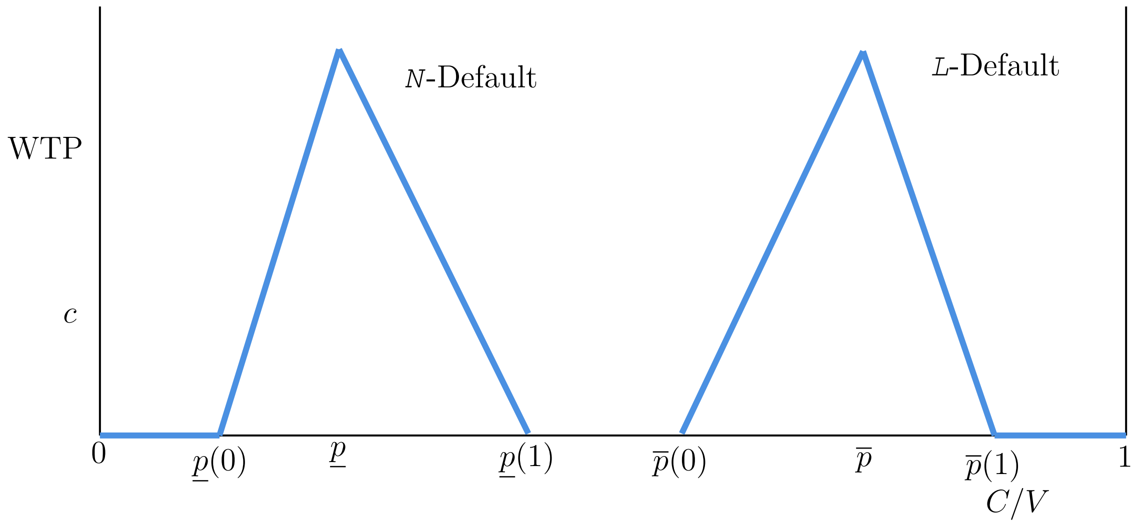

Figure 2: Willingness to Pay for Experiments

in turn, may change entrepreneurial choices. In this setup, a natural measure of the degree of Knightian uncertainty is the difference, $\Delta \equiv \bar{p} - \underline{p}$. Figure 3 shows what happens as Δ becomes smaller. In that depiction, $\bar{p}$ and $\underline{p}$ are brought closer together. As you can see, this shifts the position of each triangular WTP line towards each other, crossing over when $\bar{p}(0) < \underline{p}(1)$ and eventually fully overlapping when $\Delta = 0$. In that final case, there is no uncertainty, and $\mathcal{F}$ is a singleton yielding the usual single willingness to pay for experiments that arise in the usual Bayesian case (see, for example, Arora and Fosfuri (2005)).

While the demand for an experiment appears to be driven by the same fundamentals when Knightian uncertainty is present, it also highlights that Knightian uncertainty means that agents with different defaults in the face of that uncertainty will have positive demand for experiments along distinct ranges of (V, C). L-default entrepreneurs will choose to experiment over lower V and higher C levels than N-default ones. Those differences could be observed empirically by a researcher in the area.

However, it is also possible that the type of experiments chosen might differ between agents with different defaults. As Gans (2023) demonstrated, agents who are pre-disposed not to launch a venture (due to low priors or less favourable (V, C) to their specific situation) will choose experiments with a relatively high λ_0 and low λ_1 (termed a high-bar or "raising the bar" experiment) in order to provide a clearer signal that the venture is worth pursuing. By contrast, agents who are pre-disposed to launch a venture will choose experiments with a relatively low λ_0 and high λ_1 (termed a low-bar or "best foot forward" experiment) that provides a clear signal when it is not worthwhile to launch a venture. When there is Knightian uncertainty, then it is possible that different agents with different defaults but otherwise facing the same fundamentals, (V, C), would choose to pursue distinct experiment types.

To explore this, suppose that each agent can choose a single experiment from a set $\{\lambda_1, \lambda_0 | \lambda_1 + \lambda_0 \in [1, \Lambda)\}$ with $\Lambda \leq 2$. The following proposition characterises the outcomes when there is Knightian uncertainty.

Proposition 1 *Suppose that $\frac{C}{V} \in [\underline{p}, \bar{p}]$. Then an N-default (L-default) entrepreneur will choose an experiment $\{\lambda_1, \lambda_0\} = \{\Lambda - 1, 1\}$ ($= \{1, \Lambda - 1\}$).*

The proof of this proposition follows the same logic as Proposition 1 in Gans (2023) and is omitted. Intuitively, in the range of Knightian uncertainty, an N-default agent will choose not to launch a venture in the absence of an experiment. By setting $\lambda_0 = 1$, if the agent observes $s = 1$, then with probability 1, the venture will be successful, that is $\Pr[\theta = \text{success} | s = 1] = 1$. The 'cost' associated with this choice is that $\Pr[s = 1]$ is low. Thus, a signal pushing that agent to launch a venture will be generated with a relatively low likelihood, but when it arrives, it will be a very clear signal. The opposite holds for an L-default entrepreneur who chooses an experiment that generates a clear signal that a venture should not be pursued. What this demonstrates is that (a) if there is Knightian uncertainty, the defaults of different agents might be identified by their experimental choices and (b) the very observation that different experiments are chosen when there are the same or similar fundamentals facing the agents is itself an indicator that Knightian uncertainty is present. Thus, consistent with the implications of the Bayesian Entrepreneurship research agenda, experiments themselves explain the broader context associated with entrepreneurial choice.

In summary, Knightian uncertainty does not change the fundamental predictions of Bayesian entrepreneurship regarding the demand for experiments and the type of experiments chosen. It does add nuance to those choices, which may help identify the degree of Knightian uncertainty facing entrepreneurs.

3.2 Venture Financing

Knight (1921) did not posit what became known as Knightian uncertainty out of mere intellectual curiosity. Instead, he was interested in entrepreneurial risk-bearing – that is, why did they have to bear risks and could not diversify in some way or access some kind of insurance products? He argued that when probabilities could not be established, such markets would be impossible. Bewley (2002) and then Rigotti and Shannon (2005) explored this formally and concluded that uncertainty could prevent what might otherwise be gains to trade in insurance from being realised and as noted further by Bewley (1989), this could mean that entrepreneurs would not be able to diversify.

To examine this within the context of this model, suppose that the entrepreneur is liquidity-constrained and so must raise C from an investor to launch. They will offer an

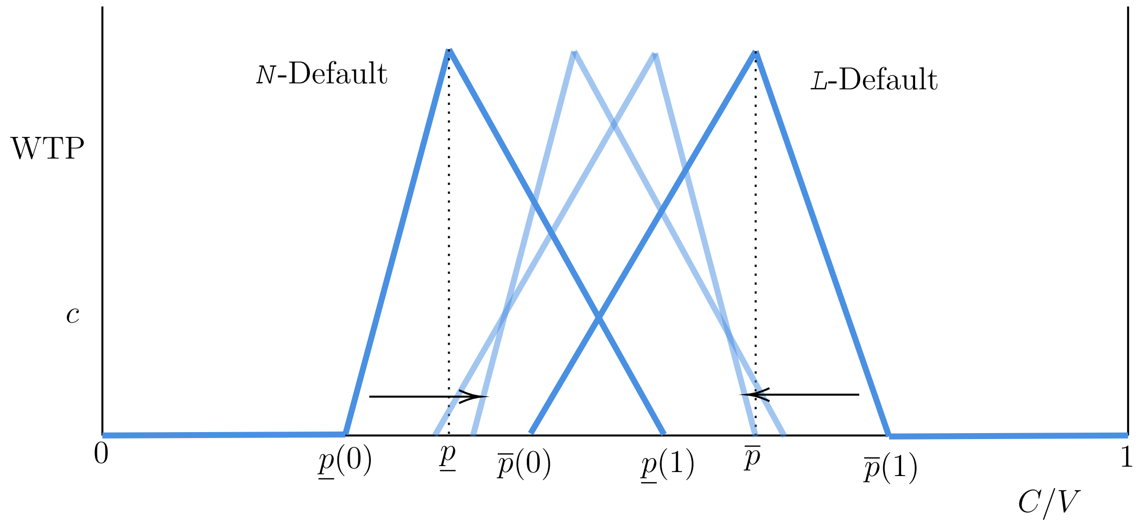

Figure 3: Impact of a Reduction in Knightian Uncertainty

investor an equity share, $1-\alpha$ with $\alpha \in (0,1]$, in return from a C investment. If V is realised, the investor earns $(1-\alpha)V - C$, and the entrepreneur earns αV. What α will the entrepreneur offer?

Assume that there are many investors who each have an outside investment opportunity for C normalised to 0. This is the investor's default. Therefore, an uncertainty-averse investor will only invest if:

$$\underline{p}(1-\alpha)V \geq C$$

This means that the upper bound on α would be:

$$\alpha \leq \hat{\alpha} \equiv \frac{\underline{p}V - C}{\underline{p}V}$$

Note that if they can secure an investor, an entrepreneur will always want to launch so long as $\hat{\alpha}V \geq 0$. Note that:

$$\hat{\alpha} \geq 0 \implies \frac{\underline{p}V - C}{\underline{p}V} \geq 0$$

That is, an investment contract will be accepted if and only if $\underline{p}(1-\alpha)V \geq C$; that is, an N-default entrepreneur would choose to invest.

There are several important things to note here. First, investor preferences drive whether financing occurs. Thus, the existence of Knightian uncertainty for an investor whose alternative options are safe assets with a payoff normalised to 0 will constrain their willingness to participate in funding the venture. Second, the entrepreneur's preferences, optimism or decisiveness do not play a role. Even if the entrepreneur were N-default, the investor's attitude towards Knightian uncertainty will drive the outcome. This reinforces the notion that Knightian uncertainty will constrain entrepreneurial activity.

Perhaps the most striking feature of this analysis is that the broad conclusions that motivate the entrepreneurial finance literature carry over to an environment where there is Knightian uncertainty.

- Staged investment: as Nanda and Rhodes-Kropf (2016) demonstrate, ventures often have staged investment that releases funds condition on reaching certain milestones that themselves involve the resolution of some uncertainty. This is akin to investing a smaller amount than C in return for a signal of venture success. As already demonstrated, this will cause investors to update their priors, specifically, $\underline{p}$, and, thereby, create conditions that allow them to rank a wider set of possibilities and make corresponding funding choices.

- Experimental choice: Gans (2022) shows that an entrepreneur who places a high prior probability on venture success will find it optimal to choose a 'high bar' experiment (that sends a clear signal of venture success) to inform potential investors and, should that experiment be successful, allow them to accept a lower equity share in return for that investment. Here, the value of such experiments would be to allow investors to rank opportunities where otherwise, $\underline{p}V < C$. Thus, the presence of Knightian uncertainty reinforces the prediction that entrepreneurs looking to secure funds from more cautious investors will choose experiments that set a high bar for success. As noted by Gans (2022), one example of such an experiment would be a minimum viable product (as recommended by Ries (2011)) that has a low probability of success but, if successful, sends a clear signal that the venture would be successful.

- Venture debt: In an unpublished paper, Rigotti (2004) considered an environment where Knightian uncertainty was over whether a venture would generate value of V or some lower amount v that might be sufficient to generate a positive net present value for investors. He demonstrated that, in that case, rather than offering pure equity to an investor, a more debt-like contract with higher investor equity if v is realised than if V is realised. This is because any equity for the outcome V is of little value in persuading investors to invest. This outcome might be achieved either by an entrepreneur using their own wealth as collateral or by offering an equity contract with a put that allows the entrepreneur to repurchase equity at a fixed price should they so choose. Again, this is the type of structure that might arise if the investor and entrepreneur held heterogeneous beliefs regarding venture success as in Agrawal et al. (2025).

In summary, many of the predictions of models with heterogeneous beliefs apply to a model where there is Knightian uncertainty. Importantly, the theoretical predictions of what happens to financial contracts and outcomes as beliefs become less heterogeneous are identical to the predictions when there is a lower degree of Knightian uncertainty in the sense that this increases $\underline{p}$. Whether there are settings in which the predictions might diverge is left for future research.

3.3 What Drives Strategic Choice?

Thus far, the uncertainty explored here has been with regard to the value of the underlying idea or entrepreneurial opportunity. However, it is sometimes the case that while the underlying idea of the venture can be evaluated with more certain probability distributions, the strategic options facing the entrepreneur are harder to assess. For instance, an entrepreneur may have an opportunity to offer two product types. One product type may be 'safe' with an expected return that is easy to identify, while another product type may be 'risky' or, in the spirit of this paper, be subject to Knightian uncertainty. This implies that when choosing between distinct strategies, the choice itself will be subject to Knightian uncertainty as the return on one strategy is generated by a distribution from a large set. The question I examine here is how this impacts entrepreneurial choices.

To formalise this choice, suppose that the strategy space for the entrepreneur is to launch the venture with strategy A, launch the venture with strategy B or not launch the venture at all (N). The final choice generates a return of 0. For strategy A, the probability of success is known at p_A and has a cost of C_A and a potential return of V_A. For strategy B, the probability of success, p_B, comes from a set of probability distributions over the range $p_B \in [\underline{p}_B, \bar{p}_B]$ with cost C_B and potential return of V_B. It is assumed that $V_A < V_B$ and $C_A < C_B$. Thus, B is more costly but also has a higher potential upside than A. To focus on the interesting case, it will be assumed that $p_A V_A > C_A$ and $p_A V_A - C_A \in (\underline{p}_B V_B - C_B, \bar{p}_B V_B - C_B)$.[15]

Recall that for two options to be ranked, one must have a higher expected payoff than the other for all distributions, $f \in \mathcal{F}$. Therefore, it is possible to specify two better-than sets:

1. $A \succ B$ if $p_A V_A - C_A > \bar{p}_B V_B - C_B$
2. $B \succ A$ if $\underline{p}_B V_B - C_B > p_A V_A - C_A$

[15] Such settings with strategic choice are explored by Gans et al. (2019), Agrawal et al. (2021) and Chavda et al. (2024). Their focus is on the conflation of signals regarding strategy and the value of the underlying idea and how this impacts the type of search processes in lieu of opportunities for strategic pivots.

Under the assumptions given here, both of these inequalities can be satisfied. Effectively, neither A nor B can be ranked due to the uncertainty over the outcomes associated with B, but each could be preferred if it was the default option.

Interestingly, this does not mean that the model is devoid of empirical predictions. Note that A is more likely to be chosen if $p_A V_A$ is higher and/or if $C_B - C_A$ is higher.[16] Thus, some comparative statics predictions are possible. In addition, as uncertainty is reduced, the prediction of the model becomes clearer in that variations in V_B become more correlated with choice outcomes compared to other 'known' parameters. Therefore, there is a sense in which the uncertainty of the underlying environment can be inferred from behavioural observations.

4 Effectuation

In recent decades, the area of entrepreneurship research that has been most built squarely on a foundation of Knightian uncertainty is effectuation theory. Proposed by Sarasvathy (2001), this theory argues that the entrepreneurial process of decision-making does not proceed according to the approaches grounded in traditional economics. The traditional approach involves the entrepreneur making clear predictions about future variables of interest and then conducting cost-benefit analyses to choose actions with the highest rate of return. By contrast, Sarasvathy (2001) argued that entrepreneurs face Knightian uncertainty that makes such prediction and, hence, cost-benefit analysis difficult. In lieu of that, entrepreneurs take actions that build on resources and relationships under their control with the intention of adjusting as uncertainty is resolved. This theory leads to a number of implications that can be distilled by an in-depth analysis of entrepreneurial decision-making.

The methodology used by those examining effectuation processes is distinct from those who use an economically-oriented methodology.[17] This makes it difficult to clearly understand and delineate when effectuation-type outcomes will result as compared with other outcomes. Having proposed a formal approach to modelling Knightian uncertainty in entrepreneurship research, my goal here is to examine whether a greater degree of Knightian uncertainty will likely lead entrepreneurs to behave in a manner more consistent with effec-

[16] Ching et al. (2019) informally consider this using what they term 'Rosenbergian uncertainty,' where the returns from different strategies facing entrepreneurs are not ranked. They then posit that in such cases, choice will be driven by known factors such as costs. This, however, is imprecise as it effectively treats the entrepreneur as indifferent with regard to future returns rather than unable to rank them and having incomplete preferences. The formal model here suggests that the empirical predictions associated with incomplete preferences are more subtle.

[17] The terminology is also different, and I will set aside there the often quite persuasive terms that have been deployed in effectuation theory to describe its core ideas.

tuation theory than other approaches.

While effectuation theory has a number of predictions, I will focus here on two broad implications.

1. Entrepreneurs following an effectuation process will utilise existing resources and networks rather than attract additional resources.[18]

2. Entrepreneurs will focus on affordable loss rather than maximising the rate of return in selecting options.

Both of these processes contrast with the Bayesian entrepreneurship processes, which focus on maximising the rate of return for an idea by optimising all potential resources and capabilities and actively experimenting before choosing actions.

To examine this, consider our earlier model where there are two options for the entrepreneur, A and B, should they launch a venture. Several adjustments are made to that model. First, both paths involve Knightian uncertainty with the range of distributions for option A being $[\underline{p}_A, \bar{p}_A]$ with $\bar{p}_A > \underline{p}_A$. Second, the entrepreneur has their own resources that are sufficient to pursue option A; specifically, the entrepreneur has wealth, W, with $C_A < W$. Wealth has alternative value to the entrepreneur of $u(W)$. Third, option B requires resources. It is assumed here that these resources cannot come from the entrepreneur's wealth; that is, their wealth is made up of resources and relationships that can only be used in pursuit of option A.[19] Consequently, the entrepreneur must secure outside funding to pursue B.

In this situation, consuming their own wealth is a natural default option for entrepreneurs, earning them $u(W)$. Importantly, this default payoff can be achieved if the entrepreneur secures outside resources for the venture, but as will be argued below, this comes at the cost of ceding equity to others.

Should the entrepreneur pursue option A, given the Knightian uncertainty and the default option, the entrepreneur would only choose A over not launching if:

$$A \succ N \text{ if and only if } \underline{p}_A V_A + u(W - C_A) > u(W)$$

$u(W) - u(W - C_A)$ can be regarded as a measure of the entrepreneur's *affordable loss*. Thus, it can be seen that as Knightian uncertainty is reduced, to the extent this increases $\underline{p}_A$, the

[18] This is sometimes referred to as a 'bird in the hand' principle but is also related to a 'pilot in the plane' principle.

[19] This is not an essential assumption, and it can be relaxed to assume that C_B can be mitigated by the entrepreneur's wealth but that this still requires outside funding.

entrepreneur's decision calculus shifts from being solely driven by affordable loss to being driven by potential returns, V_A, as well.

Now consider option B. The entrepreneur must raise C_B from outsiders. As shown earlier, outsiders will base their contributions of resources on $\underline{p}_B$. Thus, $\hat{\alpha} = \frac{\underline{p}_B V_B - C_B}{\underline{p}_B V_B}$. Note that, so long as $\underline{p}_B > \frac{C_B}{V_B}$, the entrepreneur will prefer option B to N; that is:

$$B \succ N \text{ if and only if } \underline{p}_B \frac{\underline{p}_B V_B - C_B}{\underline{p}_B V_B} V_B + u(W) > u(W) \implies \underline{p}_B V_B > C_B$$

Intuitively, unlike option A, the entrepreneur does not forgo the option of consuming their wealth.

This shows, interestingly, that the entrepreneur's affordable loss is *lower* for option B (attracting new resources) rather than option A (using existing resources). Nonetheless, there is a sense in which it could be argued that the existing resources are considered too fungible in this instance. Suppose instead that W represents the entrepreneur's time and effort. For option A, we assumed that W could be used to mitigate C_A while for B, suppose that W must be foregone to some degree, leaving $w < W$ for consumption but it remains that the entirety of C_B must be procured from outsiders. In that case, we have:

$$B \succ N \text{ if and only if } \underline{p}_B \frac{\underline{p}_B V_B - C_B}{\underline{p}_B V_B} V_B + u(w) > u(W) \implies \underline{p}_B V_B > C_B > u(W) - u(w)$$

Thus, the affordable loss from B is $u(W) - u(w)$. With this condition, the affordable loss for A can be lower than that for B if $u(W - C_A) > u(w)$. The point here is that affordable loss in of itself does not lead us to conclude that using resources on hand is preferable to procuring resources.

We have shown now that (a) as there is a greater degree of Knightian uncertainty, entrepreneurial choices will be driven more by what might be considered affordable loss (how much of their own utility will they be willing to bet on the venture's success) than other aspects of the rate of return (such as V_A or V_B); and (b) that affordable loss itself is not a clear driver of a bias towards using resources at hand rather than procuring resources. The question remains whether the degree of Knightian uncertainty can drive such a bias.

One way such a bias could arise is through defaults. The model here starts from the assumption that the entrepreneur's default option is not to launch the venture. This is consistent with the emphasis on affordable loss as a concept guiding decision-making. The entrepreneur faces Knightian uncertainty regarding both options and, hence, will only choose either option if it is preferred not to launch for all possible distributions.

This was assumed because, as emphasised by Sarasvathy and Kotha (2001), such uncertainty is a constraint for the entrepreneur in considering how to deploy their own resources. But what if we relaxed this assumption and assumed, as in our earlier model, that the entrepreneur does not face Knightian uncertainty on option A and, instead, assesses the probability of success as p_A. Assume also that, at this probability $p_A V_A + u(W - C_A) > u(W)$. Then, A becomes a natural default option for the entrepreneur. In this case, the entrepreneur will stick with that default if:

$$A \succ B \implies p_A V_A + u(W - C_A) \geq \underline{p}_B V_B - C_B + u(w)$$

In this case, the impact of Knightian uncertainty reduces to the constraint it places on others in choosing to devote resources to the venture. By contrast, option A does not require the involvement of others to provide resources other than those that already exist. As Knightian uncertainty is reduced, the bias towards using resources at hand is diminished. As Gans (2022) demonstrated a similar bias can arise in a Bayesian entrepreneurship setting when outsiders hold weaker beliefs on success than the entrepreneur. Ultimately, generating this prediction from effectuation theory using this formal approach to Knightian uncertainty requires additional behavioural assumptions regarding defaults.

In summary, we are left with the conclusion that the formal approach to Knightian uncertainty supports some principles emphasised by effectuation theorists but does not establish a clear relationship between Knightian uncertainty and any bias towards taking advantage of resources at hand rather than procuring resources externality. Indeed, to the extent that Bayesian Entrepreneurship assumes that entrepreneurs hold stronger beliefs than others, the very same constraints on procuring external resources and capital arise in each approach with and without Knightian uncertainty.

5 Conclusion

This paper has demonstrated that incorporating Knightian uncertainty into models of entrepreneurial decision-making need not invalidate the key insights and predictions generated by Bayesian approaches to entrepreneurship. Rather than creating an insurmountable barrier to formal analysis, Knightian uncertainty - when properly modelled using frameworks like Bewley (2002) - provides additional richness to our understanding of entrepreneurial behaviour while preserving important comparative static predictions.

Several key findings emerge from this analysis. First, the value of experimentation remains robust even in the face of Knightian uncertainty. While uncertainty may change the specific experiments entrepreneurs choose to run, it does not diminish their fundamental importance in entrepreneurial decision-making. Second, venture financing patterns predicted by Bayesian models continue to hold when accounting for Knightian uncertainty, particularly regarding staged investment and the use of debt-like instruments. Third, strategic choices under Knightian uncertainty exhibit patterns that are empirically distinguishable from choices made under pure risk, providing potential avenues for empirical research.

The analysis also helps bridge some of the methodological divides between formal economic approaches and other entrepreneurship research traditions, particularly effectuation theory. While the formal model does not perfectly replicate all predictions of effectuation theory, it demonstrates how Knightian uncertainty can generate behaviour patterns that align with some effectuation principles while remaining grounded in economic theory.

Here, I have presented one approach to modelling Knightian uncertainty. It was useful because it allowed a way of considering the degree of Knightian uncertainty and its impact on entrepreneurial choices. It is important to note that there are other ways of considering the types of issues arising from non-Bayesian approaches to entrepreneurial decision-making.[20] One fruitful approach would be to consider games with unawareness. These are situations where some players are unaware of states that other players are aware of. This has a natural home in entrepreneurship research as the entrepreneur can be considered as the agent who is aware, while others (including resource holders, competitors, and complementors) are unaware. Bryan et al. (2022) have been instrumental in bringing these considerations into strategy research, and there is a rich vein from that basis that entrepreneurial researchers can surely tap.

Ultimately, this paper suggests that treating Knightian uncertainty as a formal constraint on decision-making rather than an insurmountable barrier to analysis can enrich our understanding of entrepreneurship while maintaining the benefits of formal economic modelling. The challenge for future research is to continue developing frameworks that capture the unique aspects of entrepreneurial uncertainty while generating clear, testable predictions about entrepreneurial behaviour.

[20] This is demonstrated clearly in the work of Ehrig et al. (2022) and Ehrig and Foss (2022). Another approach is to consider ambiguity more directly, as was done by Amarante et al. (2012a) and Amarante et al. (2012b).

References

Agrawal, Ajay, Arnaldo Camuffo, Alfonso Gambardella, Joshua S Gans, Erin L Scott, and Scott Stern, "The Foundations of Bayesian Entrepreneurship," in "Bayesian Entrepreneurship," MIT Press, 2025.

_ , **Joshua S. Gans, and Scott Stern**, "Enabling entrepreneurial choice," *Management Science*, 2021, *67* (9), 5510–5524.

Allais, Maurice, "Le comportement de l'homme rationnel devant le risque: critique des postulats et axiomes de l'école Américaine," *Econometrica*, 1953, *21* (4), 503–546.

Amarante, Massimiliano, Mario Ghossoub, and Edmund S. Phelps, "Contracting for Innovation under Ambiguity," *Center on Capitalism and Society Working Papers*, 2012, (75).

_ , _ , **and** _ , "Entrepreneurship, Ambiguity, and the Shape of Innovation Contracts," *Center on Capitalism and Society Working Papers*, 2012, (76).

Anscombe, F. J. and Robert J. Aumann, "A Definition of Subjective Probability," *Annals of Mathematical Statistics*, 1963, *34* (1), 199–205.

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

Arrow, Kenneth J. and Leonid Hurwicz, "An Optimality Criterion for Decision-Making under Ignorance," *Uncertainty and Expectations in Economics*, 1972, *1.*

Aumann, Robert J., "Utility Theory Without the Completeness Axiom," *Econometrica*, 1962, *30* (3), 445–462.

Bewley, Truman F., "Knightian Decision Theory, Part II: Intertemporal problems," Technical Report, Yale 1987.

_ , "Market Innovation and Entrepreneurship: A Knightian View," 1989. mimeo.

_ , "Knightian Decision Theory, Part I," *Decisions in Economics and Finance*, 2002, *25* (2), 79–110.

Bhidé, Amar, "Renewing Knightian Uncertainty: A Pragmatic Prospectus and Demonstration," Working Paper 21-129, Harvard Business School 2021. Preliminary draft, June 10, 2021.

_ , *Uncertainty and Enterprise: Venturing Beyond the Known*, Oxford University Press, 2024.

Bryan, Kevin, Michael Ryall, and Burkhard C. Schipper, "Value Capture in the Face of Known and Unknown Unknowns," *Strategy Science*, 2022, *7* (3), 157–189.

Camuffo, A., A. Gambardella, and A. Pignataro, "Theory-Driven Strategic Management Decisions," 2023. CEPR Discussion Paper n. 17664-3.

Chavda, Ankur, Joshua S. Gans, and Scott Stern, "Theory-Driven Entrepreneurial Search," *Strategy Science*, 2024.

Ching, Kenny, Joshua S. Gans, and Scott Stern, "Control versus Execution: Endogenous Appropriability and Entrepreneurial Strategy," *Industrial and Corporate Change*, 2019, *28* (2), 389–408.

Ehrig, Timo and Nicolai J. Foss, "Why We Need Normative Theories of Entrepreneurial Learning That Go Beyond Bayesianism," *Journal of Business Venturing Insights*, 2022, *18*, e00335.

_ , Nicolai J Foss et al., "Unknown unknowns and the treatment of firm-level adaptation in strategic management research," *Strategic Management Review*, 2022, *3* (1), 1–24.

_ , Teppo Felin, and Todd R. Zenger, "Causal Reasoning and the Scientific Entrepreneur: Beyond Bayes," in Ajay Agrawal, Arnaldo Camuffo, Alfonso Gambardella, Joshua S Gans, Erin L Scott, and Scott Stern, eds., *Bayesian Entrepreneurship*, MIT Press, 2025.

Ellsberg, Daniel, "Risk, Ambiguity, and the Savage Axioms," *The Quarterly Journal of Economics*, 1961, *75* (4), 643–669.

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

Gans, Joshua S., "On the Impossibility of Rational Choice Under Incomplete Information," *Journal of Economic Behavior & Organization*, 1996, *29* (2), 287–309.

_ , "A Theory of Visionary Disruption," 2022. Working Paper, NBER.

_ , "Experimental Choice and Disruptive Technologies," *Management Science*, 2023, *69* (11), 7044–7058.

_ , E. L. Scott, and S. Stern, *Entrepreneurship: Choice and Strategy*, New York: Norton, 2024.

_ , S. Stern, and J. Wu, "Foundations of Entrepreneurial Strategy," *Strategic Management Journal*, 2019, *40* (5), 736–756.

Gerasimou, Georgios, "On the indifference relation in Bewley preferences," *Economics Letters*, 2018, *164*, 24–26.

Gilboa, Itzhak and David Schmeidler, "Maxmin Expected Utility with a Non-Unique Prior," *Journal of Mathematical Economics*, 1989, *18*, 141–153. Originally circulated as a working paper in 1986.

_ **and Massimo Marinacci**, "Ambiguity and the Bayesian Paradigm," in Daron Acemoglu, Manuel Arellano, and Eddie Dekel, eds., *Advances in Economics and Econometrics*, Cambridge University Press, 2013, pp. 179–242.

Kay, John and Mervyn King, *Radical Uncertainty: Decision-Making Beyond the Numbers*, New York: W. W. Norton & Company, 2020.

Keynes, John Maynard, *A Treatise on Probability*, London: Macmillan, 1921.

_ , "The General Theory of Employment," *Quarterly Journal of Economics*, 1937, *51* (2), 209–223.

Knight, Frank H., *Risk, Uncertainty, and Profit*, Boston: Houghton Mifflin, 1921.

LeRoy, Stephen F. and Larry D. Singell Jr., "Knight on Risk and Uncertainty," *Journal of Political Economy*, 1987, *95* (2), 394–406.

Lopomo, Giuseppe, Luca Rigotti, and Chris Shannon, "Knightian uncertainty and moral hazard," *Journal of Economic Theory*, 2011, *146* (3), 1148–1172.

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

Neumann, John Von and Oskar Morgenstern, *Theory of Games and Economic Behavior*, Princeton, NJ: Princeton University Press, 1944.

Nishimura, Kiyohiko G. and Hiroyuki Ozaki, "Search and Knightian uncertainty," *Journal of Economic Theory*, 2004, *119* (2), 299–333.

Ok, Efe A., Pietro Ortoleva, and Gil Riella, "Incomplete Preferences Under Uncertainty: Indecisiveness in Beliefs versus Tastes," *Econometrica*, 2012, *80* (4), 1791–1808.

Olea, José Luis Montiel, Pietro Ortoleva, Mallesh M. Pai, and Andrea Prat, "Competing Models," *The Quarterly Journal of Economics*, 2022, *137* (4), 2419–2457.

Packard, Mark D., Brent B. Clark, and Peter G. Klein, "Uncertainty Types and Transitions in the Entrepreneurial Process," *Organization Science*, 2017, *28* (5), 840–856.

_ **, Per L. Bylund, and Brent B. Clark**, "Keynes and Knight on Uncertainty: Peas in a Pod or Chalk and Cheese?," *Cambridge Journal of Economics*, 2021, *45* (5), 1099–1125.

Ries, E., *The Lean Startup: How Today's Entrepreneurs Use Continuous Innovation to Create Radically Successful Businesses*, Currency, 2011.

Rigotti, Luca, "Decisive entrepreneurs and cautious investors," in "Duke University Working Paper," Citeseer, 2004.

_ **and Chris Shannon**, "Uncertainty and Risk in Financial Markets," *Econometrica*, 2005, *73* (1), 203–243.

Ryall, Michael D., "Causal Ambiguity, Complexity, and Capability-Based Advantage," *Management Science*, 2009, *55* (3), 389–403.

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

_ **and S. Kotha**, "Managing Knightian Uncertainty in the New Economy," in B. Hunt and S. Kotha, eds., *Entrepreneurship and the Internet Era: Cases and Contexts*, Amsterdam: Elsevier Science, 2001, pp. 31–50.

Savage, Leonard J., *The Foundations of Statistics*, New York: John Wiley & Sons, 1954.

Shackle, George L. S., *Decision, Order and Time in Human Affairs*, Cambridge University Press, 1961.

Sutton, John, "Flexibility, profitability and survival in an (objective) model of Knightian uncertainty," Technical Report, London School of Economics 2006.

Vaswani, Ashish, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N. Gomez, Lukasz Kaiser, and Illia Polosukhin, "Attention Is All You Need," in "Advances in Neural Information Processing Systems (NeurIPS)," Vol. 30 2017, pp. 5998–6008.

Wiltbank, Ronald N., Nicholas Dew, Stuart Read, and Saras D. Sarasvathy, "What to Do Next? The Case for Non-Predictive Strategy," *Strategic Management Journal*, 2006, *27* (10), 981–998.