NBER WORKING PAPER SERIES

# MANAGING THE ENTREPRENEUR'S MIRACLE BUDGET

Joshua S. Gans

Working Paper 34850
http://www.nber.org/papers/w34850

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
February 2026

Thanks to Refine.ink, ChatGPT 5.2 and Claude Opus 4.6 for valuable research assistance, and to SSHRC for funding. Responsibility for all errors remains my own. The views expressed herein are those of the author and do not necessarily reflect the views of the National Bureau of Economic Research.

Managing the Entrepreneur's Miracle Budget
Joshua S. Gans
NBER Working Paper No. 34850
February 2026
JEL No. D81, D83, L26

## ABSTRACT

Entrepreneurial ventures typically require multiple independent conditions to hold simultaneously for success. This paper develops a formal framework for decision-making under such multiplicative uncertainty. When a venture's success depends on several factors—each of which must resolve favourably—even modestly uncertain factors compound to produce low overall success probabilities. I introduce the concept of a *miracle budget*: the maximum aggregate uncertainty a venture can carry while remaining viable. The miracle budget increases with a venture's payoff-to-cost ratio but only logarithmically, so even transformative opportunities buy only modest additional slack in aggregate uncertainty. The framework characterises optimal experimentation strategies, including the value of information from different experimental designs, conditions for optimal sequencing of experiments, and the allocation of experimental effort across uncertain factors. A hierarchical model distinguishes *real miracles*—genuinely uncertain factors that can only be resolved by trying—from *apparent miracles*—factors that seem uncertain but can be resolved through investigation. The framework yields actionable principles for early-stage entrepreneurs: compute your budget, prioritise low-cost tests with high “kill” probability (and, when experiment costs are similar, start with the most uncertain factors), and focus early-stage work on learning how many miracles you actually face.

Joshua S. Gans
University of Toronto
Rotman School of Management
and NBER
joshua.gans@rotman.utoronto.ca

# 1 Introduction

Entrepreneurial ventures are characterised by multiplicative uncertainty: success typically requires that several independent conditions hold simultaneously, and failure of any single condition is sufficient to sink the enterprise. This structure—where overall viability is the product of component probabilities—means that even modestly uncertain factors compound rapidly. Seven factors, each at 90 per cent confidence, yield an overall probability of just 48 per cent; at 70 per cent each, the product falls to 8 per cent.

The compounding problem has an analogue in production theory. Kremer (1993) formalised a similar multiplicative structure in his O-Ring theory of economic development, showing that production processes requiring many tasks to be performed correctly exhibit extreme sensitivity to the weakest link. The entrepreneurial setting shares this structure: a single factor at near-zero probability can render the entire venture unviable, regardless of how favourable the remaining factors may be.

This observation motivates the central question of the paper: given a fixed payoff-to-cost ratio, how much aggregate uncertainty can a venture sustain, and how should the entrepreneur allocate experimental effort to reduce that uncertainty? Martínez (2016) offered a practitioner's heuristic: count the number of "miracles" required for the venture to succeed, where a miracle is a factor whose outcome is genuinely uncertain. In Section 3, I operationalise this idea by introducing a practical cutoff: fix a confidence threshold and treat any necessary condition below that threshold as a "miracle." Under this convention, "counting miracles" becomes counting the number of factors whose uncertainty is large enough to meaningfully consume the venture's slack. Most successful startups, he argued, depend on exactly one miracle; ventures requiring two or more are rarely viable. For Airbnb, the miracle was persuading strangers to stay in one another's homes. For Google, it was building an exponentially better search algorithm. The present paper formalises this intuition and derives its implications for optimal experimentation. The purpose here is to explore specifically and formally how compound probabilities can be managed by entrepreneurs and how they impact entrepreneurial learning strategies.

I introduce the concept of a *miracle budget*: the maximum aggregate uncertainty a venture can carry while still justifying a commitment decision. The key result is that the budget increases with the venture's payoff-to-cost ratio but only logarithmically, so increasing upside can relax the constraint only slowly. I then characterise optimal experimentation strategies for reducing uncertainty to fit within this budget, including the value of information from different experimental designs, conditions for optimal sequencing of experiments, and portfolio allocation of experimental effort across uncertain factors.

The framework contributes to several strands of the entrepreneurship literature. First, it provides a formal foundation for the widely discussed but loosely specified notion that entrepreneurs should "de-risk" their ventures through experimentation (Ries, 2011; Blank, 2013; Gambardella et al., 2026). The lean startup methodology advocates for rapid experimentation and validated learning, but offers limited formal guidance on which experiments to run, in what order, and when to stop. The miracle budget framework fills this gap by characterising optimal experimentation strategies from first principles. Second, the paper connects to the theoretical literature on entrepreneurship under uncertainty (Knight, 1921; McMullen and Shepherd, 2006; Alvarez and Barney, 2007; Agrawal et al., 2026; Gans, 2026), providing a tractable model that distinguishes between uncertainty that can be resolved through investigation and uncertainty that can only be resolved by trying. Third, the framework relates to the growing literature on entrepreneurial strategy (Gans et al., 2019; Agrawal et al., 2021), offering a precise characterisation of when and how entrepreneurs should commit to a course of action versus continue experimenting. Finally, the paper connects to work on the value of experimentation in innovation (Thomke, 2003; Kerr et al., 2014; Camuffo et al., 2020; Bailey et al., 2026), providing conditions under which different types of experiments are valuable and characterising optimal experimental design.

The remainder of the paper proceeds as follows. Section 2 establishes the basic model and derives a GO/NO GO condition under multiplicative uncertainty. Section 3 introduces the log-risk formulation and establishes that the miracle budget scales logarithmically with the payoff-to-cost ratio. Section 4 analyses the value of information from experiments, distinguishing perfect and imperfect information. Section 5 characterises optimal experimentation strategies, including sequencing and portfolio allocation. Section 6 develops the hierarchical model distinguishing real from apparent miracles, illustrated with case studies. Section 7 synthesises the results into actionable principles for entrepreneurs and discusses implications for theory and practice. The conclusion summarises the results and outlines directions for future research.

# 2 The Basic Model

## 2.1 Setup

Consider an entrepreneur facing a venture that requires $N$ independent factors to succeed. Each factor represents a necessary condition—if any single factor fails, the entire venture fails. Let $\theta_i \in \{0, 1\}$ denote the true (unknown) state of factor $i$, where $\theta_i = 1$ means factor $i$ will work and $\theta_i = 0$ means it will not. The venture succeeds if and only if all factors resolve favourably:

$$\text{Success} \iff \prod_{i=1}^{N} \theta_i = 1$$

This multiplicative structure captures a fundamental feature of entrepreneurial ventures: they are systems of complementary components. A brilliant technology paired with a non-existent market produces nothing. A huge market paired with an unworkable technology produces nothing. The structure is analogous to a historical problem in commerce: merchants sailing the Rhine River faced dozens of castles, each levying its own toll. A ship starting with cargo worth 100 units would lose 10 per cent at the first castle, leaving 90. The second castle took 10 per cent of the remainder, leaving 81. By the tenth castle, even though each individual toll was "small," the cumulative effect had reduced the cargo's value to barely 35 per cent of its original worth. Entrepreneurs face the same compounding, but in the domain of probability rather than cargo value.

The entrepreneur holds beliefs about each factor. Let $\pi_i \in (0, 1)$ denote the entrepreneur's subjective probability that factor $i$ will work, i.e., $\pi_i = P(\theta_i = 1)$.

**Assumption 1** (Independence)**.** *The factors $\theta_1, \ldots, \theta_N$ are independent in the entrepreneur's belief distribution.*

Independence is a strong assumption that warrants justification. In practice, factors may be correlated: if the technology works, the market may be more likely to exist (both driven by underlying demand), and a positive signal about one factor may partially resolve another. The focus on independence is motivated by two considerations. First, independence provides a useful conservative benchmark relative to many settings with positively correlated factors, where learning about one factor can partially resolve others; negative correlation is possible but appears less common in entrepreneurial contexts. Second, independence is the natural setting in which to define a *budget*, since each factor's contribution to total uncertainty is separable. Under Assumption 1, the probability of venture success is $P(\text{Success}) = \prod_{i=1}^{N} \pi_i$.

**Beyond Independence: Correlation and Information Spillovers.** While correlation and causal dependence are something left for future work, it is useful to comment on what independence is restricting.Independence is a benchmark, not a claim about reality. If the factors are correlated, the venture success probability is no longer pinned down by the marginals $(\pi_i)$ alone. A simple way to see this is via the Fréchet bounds: letting $A_i = \{\theta_i = 1\}$ with $P(A_i) = \pi_i$,

$$\max\left\{0, \sum_{i=1}^{N} \pi_i - (N-1)\right\} \leq P\left(\bigcap_{i=1}^{N} A_i\right) \leq \min_i \pi_i.$$

The product $\prod_i \pi_i$ is the independence benchmark that lies between these extremes. With positive dependence, the true joint success probability can exceed the product (making the independence-based miracle budget conservative); with negative dependence it can fall below it (making the budget optimistic).[1]

## 2.2 The GO/NO GO Decision

The entrepreneur faces a binary choice: commit to the venture (GO), paying cost $C > 0$ and receiving payoff $V (> C)$ if successful and 0 if unsuccessful; or abandon the venture (NO GO), receiving payoff 0. The binary framing captures the essential structure of the entrepreneurial commitment decision. In practice, "GO" may represent launching a product, accepting venture capital, or scaling operations—any irreversible commitment of resources. The distinction between "GO" and "continue experimenting" is central to the entrepreneurial strategy literature: entrepreneurs must balance the value of additional information against the cost of delay (Dixit and Pindyck, 1994; Gans et al., 2019; Folta, 1998). This model focuses on the information acquisition that precedes the commitment decision.

The expected value of committing to the venture is $\mathbb{E}[\text{GO}] = V \prod_{i=1}^{N} \pi_i - C$. It is straightforward to see that the entrepreneur chooses GO if and only if $\mathbb{E}[\text{GO}] \geq \mathbb{E}[\text{NO GO}] = 0$:

$$V \prod_{i=1}^{N} \pi_i - C \geq 0 \iff \prod_{i=1}^{N} \pi_i \geq \frac{C}{V} \tag{1}$$

The ratio $C/V$ on the right-hand side is the venture's *hurdle probability*: the minimum overall success probability required to justify commitment. A venture with a 100x payoff ratio ($V/C = 100$) requires only a 1 per cent success probability; a venture with a 10x ratio requires 10 per cent. This connects to the venture capital logic of seeking "home run" returns: high payoff ratios lower the hurdle probability, making it rational to pursue ventures that would be unjustifiable under more modest return expectations.

**Example 1.** *Suppose* $V/C = 100$ *(the venture pays off 100 times the investment if successful). Then the entrepreneur requires* $\prod_i \pi_i \geq 0.01$*. With five factors each at* $\pi_i = 0.4$*, we have* $0.4^5 \approx 0.01$*, which is marginal. Six such factors would yield* $0.4^6 \approx 0.004$*, falling well below the threshold. Even a venture with enormous upside can tolerate only a handful of substantially uncertain factors.*

[1] Correlation also breaks the "one-for-one" trade-off logic of Proposition 4 (below). When learning about one factor partially resolves another, de-risking one component can effectively change the "price" of uncertainty elsewhere. In such settings, the budget still provides a useful accounting device for marginal beliefs, but managerial prescriptions should be interpreted with care: a factor that looks non-pivotal in isolation may be worth testing because it creates information spillovers.

# 3 The Log-Risk Formulation and Miracle Budget

## 3.1 Converting to Additive Form

The multiplicative structure of (1) is difficult to work with directly, particularly when thinking about how experimental effort on one factor trades off against effort on another. Taking logarithms converts the multiplicative constraint into an additive one, revealing a natural "budget" structure.

For factor $i$ with belief $\pi_i$, I define the *log-risk* as $\lambda_i = -\log(\pi_i) \geq 0$, where log denotes the natural logarithm. This quantity—sometimes called *surprisal* in information theory—measures how "surprising" it would be for factor $i$ to succeed. Log-risk has intuitive properties: $\lambda_i = 0$ when $\pi_i = 1$ (certainty of success, no surprise) and $\lambda_i \to \infty$ as $\pi_i \to 0$ (certainty of failure). The inverse relationship is $\pi_i = e^{-\lambda_i}$. A factor at 50 per cent probability consumes $\lambda_i = \log(2) \approx 0.693$ units of log-risk; a factor at 90 per cent consumes only $\lambda_i \approx 0.105$ units.

Note that the GO condition $\prod_i \pi_i \geq C/V$ is equivalent to:[2]

$$\sum_{i=1}^{N} \lambda_i \leq \log\left(\frac{V}{C}\right) \tag{2}$$

This additive formulation is the key step. It transforms the entrepreneur's problem from a multiplicative constraint (which is difficult to decompose and reason about) into a budget constraint (which is familiar and tractable). Each factor $i$ "spends" $\lambda_i$ units from a finite budget, and the total spending must not exceed the budget.

Using this, I now define the main formal object of the paper:

**Definition 1** (Miracle Budget)**.** *The* miracle budget *is:*

$$B = \log\left(\frac{V}{C}\right)$$

The GO condition is now simply: total log-risk must fit within budget $B$. This framing immediately suggests the entrepreneurial task: reduce the sum $\sum_i \lambda_i$ (through experimentation, restructuring, or abandonment of the venture) until it fits within the available budget $B$.

[2]Taking logs of both sides of the GO condition:

$$\prod_{i=1}^{N} \pi_i \geq \frac{C}{V} \iff \sum_{i=1}^{N} \log(\pi_i) \geq \log\left(\frac{C}{V}\right) \iff \sum_{i=1}^{N} (-\lambda_i) \geq -\log\left(\frac{V}{C}\right)$$

Multiplying by $-1$ and reversing the inequality gives (2).

**Definition 2** (Budget Slack)**.** *Define the venture's* budget slack *as*

$$S(\lambda) \equiv B - \sum_{i=1}^{N} \lambda_i.$$

*GO is optimal exactly when $S(\lambda) \geq 0$; when $S(\lambda) < 0$ the venture is short by $|S(\lambda)|$ units of log-risk.*

Slack makes "proximity to viability" explicit: experiments and restructurings matter to the extent that they increase $S(\lambda)$.

**Proposition 1** (Pivotality and Proximity to the Budget Frontier)**.** *Suppose the entrepreneur is currently in the NO GO region ($S(\lambda) < 0$). Resolving factor $i$ perfectly and favourably (setting $\pi_i = 1$, hence $\lambda_i = 0$) makes GO feasible if and only if $\lambda_i \geq -S(\lambda)$ (equivalently, $\sum_{j \neq i} \lambda_j \leq B$). Thus, only factors whose log-risk is large enough to close the slack can be pivotal in the one-shot sense of Proposition 5.*

*Proof.* Perfect favourable resolution of factor $i$ increases slack by exactly $\lambda_i$, since it removes $\lambda_i$ from the sum. The post-resolution slack is $S' = S(\lambda) + \lambda_i$, which is nonnegative if and only if $\lambda_i \geq -S(\lambda)$. □

The budget condition depends only on $\sum_i \lambda_i$, but experimentation often has fixed per-factor set-up costs (time, attention, and an experiment-design overhead). Holding constant the overall success probability $\Pi$, concentrating uncertainty into fewer factors can reduce the number of distinct experiments required to reach a GO decision (and hence expected experimentation cost), while spreading uncertainty across many small risks can create "death by a thousand cuts" in experimentation effort. Thus, the *distribution* of log-risk across factors is payoff-irrelevant at commitment but can be decision-relevant once learning technologies are introduced.

## 3.2 Properties of the Miracle Budget

The miracle budget has a crucial structural property: it scales only logarithmically with the venture's payoff-to-cost ratio. This means that even dramatic increases in a venture's upside translate into modest expansions of the budget available to absorb uncertainty.

**Proposition 2** (Logarithmic Scaling)**.** *The miracle budget scales logarithmically with the payoff-to-cost ratio:*

*(a) Doubling $V/C$ increases the budget by* $\log(2) \approx 0.693$*.*

*(b) A 10x increase in $V/C$ increases the budget by* $\log(10) \approx 2.303$*.*

*Proof.* Direct from the definition. If $B = \log(V/C)$, then:

(a) $B' = \log(2 \cdot V/C) = \log(2) + \log(V/C) = \log(2) + B$.

(b) $B' = \log(10 \cdot V/C) = \log(10) + B$.

□

The logarithmic scaling is particularly consequential for entrepreneurs who believe that massive opportunities should permit massive risk-taking. Going from a 10x opportunity to a 100x opportunity—a tenfold increase in potential payoff—adds only $\log(10) \approx 2.3$ to the budget, which accommodates roughly three additional coin-flip miracles. Going from 100x to 1000x adds only another $\log(10)$. The budget grows, but it grows slowly. Founders often reason that pursuing a larger opportunity justifies taking on proportionally more risk; the miracle budget shows that the relationship is sharply concave. A 100x opportunity does not permit 10 times more uncertainty than a 10x opportunity. It permits only modestly more.

A natural question is how many distinct uncertain factors—"miracles"—a venture can carry within its budget. If we impose a lower bound on each factor's uncertainty, the budget directly limits the number of unresolved factors that can coexist with a positive expected value.

**Proposition 3** (Miracle Count Bound)**.** *Suppose all unresolved factors have log-risk at least* $\bar{\lambda} > 0$ *(i.e.,* $\pi_i \leq e^{-\bar{\lambda}}$ *for all unresolved* $i$*). Then the maximum number of unresolved factors consistent with GO is:*

$$k^* = \left\lfloor \frac{B}{\bar{\lambda}} \right\rfloor = \left\lfloor \frac{\log(V/C)}{\bar{\lambda}} \right\rfloor$$

*Proof.* If $k$ factors each contribute at least $\bar{\lambda}$ to the sum, then $\sum_i \lambda_i \geq k\bar{\lambda}$. For the GO condition to hold, we need $k\bar{\lambda} \leq B$, i.e., $k \leq B/\bar{\lambda}$. Since $k$ must be an integer, $k \leq \lfloor B/\bar{\lambda} \rfloor$. □

**Example 2.** *Let* $V/C = 100$*, so* $B = \log(100) \approx 4.605$*. If we define a "miracle" as a factor with* $\pi_i \leq 0.5$ *(so* $\bar{\lambda} = \log(2) \approx 0.693$*), then the entrepreneur can carry at most* $\lfloor 4.605/0.693 \rfloor = 6$ *miracles. This gives precise content to Martínez's heuristic: even with a 100x payoff ratio—the kind of outcome venture capitalists dream about—the venture can sustain at most six coin-flip miracles. A seventh tips the expected value below zero. His one-miracle rule corresponds roughly to the budget available for more modest payoff ratios, around* $V/C = 2$ *to* $4$ *under this coin-flip definition.*

### 3.3 The Substitution Structure

The additive structure of the budget constraint reveals an important flexibility: the entrepreneur can trade off certainty across factors.

**Proposition 4** (Budget Allocation)**.** *The GO condition defines a simplex-like feasible region in* $(\lambda_1, \ldots, \lambda_N)$ *space. The entrepreneur can trade off certainty across factors: reducing* $\lambda_i$ *(increasing confidence in factor* $i$*) allows increasing* $\lambda_j$ *(accepting more uncertainty in factor* $j$*) while remaining feasible. The marginal rate of substitution is* $d\lambda_j/d\lambda_i = -1$ *along the constraint boundary.*

*Proof.* The constraint $\sum_i \lambda_i \leq B$ with $\lambda_i \geq 0$ defines a simplex in $\mathbb{R}^N_+$. The marginal rate of substitution is $d\lambda_j/d\lambda_i = -1$ along the constraint boundary: one unit of reduced log-risk on factor $i$ permits exactly one unit of increased log-risk on factor $j$. □

This unit-for-unit substitution in log-risk corresponds to a *multiplicative* substitution in probability space. Increasing $\pi_i$ by a factor $\alpha$ allows decreasing $\pi_j$ by a factor $\alpha$ while maintaining the same product. The practical implication is that nailing down one miracle—say, establishing with high confidence that the technology works—frees budget for more uncertainty elsewhere, perhaps in market acceptance or competitive dynamics. But the total budget is fixed by the payoff-to-cost ratio.

Table 1 illustrates the relationship between belief probability and budget consumption. The key insight is the nonlinearity: a factor at 50 per cent probability consumes roughly seven times more budget than a factor at 90 per cent. This means that a small number of highly uncertain factors can exhaust the budget, while many nearly-certain factors consume almost nothing.

# 4 The Value of Experimentation

I now turn to the entrepreneur's problem *before* the GO/NO GO decision: should they run experiments to learn about the factors? The lean startup movement has popularised the idea that entrepreneurs should "get out of the building" and test assumptions through rapid experimentation (Ries, 2011; Blank, 2013). The miracle budget framework provides formal foundations for this intuition, characterising precisely when experiments are valuable and how much information they provide.

The value of experimentation in entrepreneurship has been studied empirically by Camuffo et al. (2020), who showed in a randomised controlled trial that entrepreneurs trained in a "scientific approach" to decision-making—formulating hypotheses, designing rigorous tests,

| $\pi_i$ | $\lambda_i = -\log(\pi_i)$ | Budget consumption |
|---|---|---|
| 0.99 | 0.010 | Very low |
| 0.95 | 0.051 | Low |
| 0.90 | 0.105 | Low |
| 0.80 | 0.223 | Moderate |
| 0.70 | 0.357 | Moderate |
| 0.50 | 0.693 | High |
| 0.30 | 1.204 | Very high |
| 0.10 | 2.303 | Extreme |

Table 1: Relationship between belief probability and budget consumption. A factor at 50 per cent confidence consumes nearly seven times more budget than a factor at 90 per cent confidence. This nonlinearity drives the key practical implication: reducing a factor from highly uncertain to moderately uncertain frees far more budget than polishing an already-confident factor.

and updating beliefs based on evidence—performed significantly better than a control group. The present framework provides the theoretical scaffolding for such a scientific approach.

## 4.1 Perfect Information Experiments

A *perfect information experiment* on factor $i$ has cost $c_i > 0$ and perfectly reveals $\theta_i$ (the primary state variable defined earlier). After such an experiment, with probability $\pi_i$ the entrepreneur learns $\theta_i = 1$ and updates to $\pi_i' = 1$ (the factor is resolved positively, consuming zero budget); with probability $1 - \pi_i$ the entrepreneur learns $\theta_i = 0$, and the venture is doomed (making NO GO the optimal action). The asymmetry is instructive: a positive result eliminates the factor's budget consumption entirely, while a negative result terminates the venture, saving the entrepreneur from committing $C$ to a doomed enterprise. Both outcomes are valuable—the first by freeing budget, the second by preventing a costly mistake.

**Proposition 5** (Value of Perfect Information)**.** *Let $\Pi = \prod_{j=1}^{N} \pi_j$ be the current success probability and $\Pi_{-i} = \prod_{j \neq i} \pi_j$ the success probability excluding factor $i$. Consider the one-shot policy that runs a perfect information experiment on factor $i$ (cost $c_i$) and then immediately chooses GO or NO GO optimally. The value of this experiment relative to acting immediately is:*

$$VOI_i = -c_i + \pi_i \max(V\Pi_{-i} - C, 0) - \max(V\Pi - C, 0). \tag{3}$$

*Equivalently,*

$$VOI_i = \begin{cases} C(1-\pi_i) - c_i, & \textit{if } V\Pi \geq C, \\ \pi_i(V\Pi_{-i} - C) - c_i, & \textit{if } V\Pi < C \textit{ and } V\Pi_{-i} \geq C, \\ -c_i, & \textit{if } V\Pi < C \textit{ and } V\Pi_{-i} < C. \end{cases} \tag{4}$$

*Proof.* Let $S(\pi) = \max(V\Pi - C, 0)$ denote the optimal payoff from acting immediately (GO if $V\Pi \geq C$, otherwise NO GO).

If we run a perfect information experiment on factor $i$:

- With probability $\pi_i$ we learn $\theta_i = 1$, update to $\pi_i' = 1$, and the remaining success probability becomes $\Pi_{-i}$. The optimal continuation payoff is then $\max(V\Pi_{-i} - C, 0)$.
- With probability $1 - \pi_i$, we learn $\theta_i = 0$, in which case the venture cannot succeed, and the optimal continuation payoff is 0.

Therefore, the expected payoff from experimenting (and then acting optimally) is

$$\mathbb{E}[\text{Experiment on } i] = -c_i + \pi_i \max(V\Pi_{-i} - C, 0).$$

Subtracting the no-experiment value $S(\pi)$ gives

$$\text{VOI}_i = -c_i + \pi_i \max(V\Pi_{-i} - C, 0) - \max(V\Pi - C, 0).$$

To obtain the piecewise form, note that $\Pi_{-i} = \Pi/\pi_i \geq \Pi$ since $\pi_i \leq 1$.

There are three cases to consider:

1. $V\Pi \geq C$: Then $V\Pi_{-i} \geq V\Pi \geq C$, so $\max(V\Pi_{-i} - C, 0) = V\Pi_{-i} - C$ and $\max(V\Pi - C, 0) = V\Pi - C$. Hence,

$$\begin{aligned} \text{VOI}_i &= -c_i + \pi_i(V\Pi_{-i} - C) - (V\Pi - C) \\ &= -c_i + \pi_i V\Pi_{-i} - \pi_i C - V\pi_i\Pi_{-i} + C \\ &= C(1 - \pi_i) - c_i. \end{aligned}$$

2. $V\Pi < C$ and $V\Pi_{-i} \geq C$: Then $\max(V\Pi - C, 0) = 0$. If $V\Pi_{-i} \geq C$, then after $\theta_i = 1$ we would choose GO, giving

$$\text{VOI}_i = -c_i + \pi_i(V\Pi_{-i} - C).$$

3. $V\Pi < C$ and $V\Pi_{-i} < C$: Then even after learning $\theta_i = 1$ we would still choose NO GO, so $\max(V\Pi_{-i} - C, 0) = 0$ and $\text{VOI}_i = -c_i$.

□

The three cases have natural interpretations. In the first ($V\Pi \geq C$), the venture is already viable—the entrepreneur would choose GO even without the experiment. The experiment's value comes entirely from the *option to abandon*: with probability $1 - \pi_i$, the experiment reveals that the factor will fail, saving the entrepreneur from committing $C$ to a doomed venture. This case corresponds to "pre-mortem" testing—checking assumptions before making an irreversible commitment. In the second case ($V\Pi < C$ but $V\Pi_{-i} \geq C$), the venture is not currently viable, but it would become viable if factor $i$ were resolved positively. The experiment's value comes from the *option to proceed*: with probability $\pi_i$, the experiment reveals good news that tips the venture into the GO region. This is the classic startup scenario where the entrepreneur needs to demonstrate feasibility before committing.

In the third case ($V\Pi < C$ and $V\Pi_{-i} < C$), the venture is not viable even if factor $i$ is resolved. No single experiment on factor $i$ can change the decision. In this myopic, one-shot analysis, the experiment has no informational value. The condition $V\Pi_{-i} < C$ means that resolving factor $i$ alone cannot make GO optimal. This is a *myopic* sense in which factor $i$ is "non-pivotal." In sequential settings, several individually non-pivotal factors can be worth testing jointly, because passing multiple tests can move the venture into the GO region.

Proposition 5 characterises the value of experimentation in each regime. The following corollary distils the conditions under which the entrepreneur should actually run the experiment—that is, when the informational benefit exceeds the experimental cost.

**Corollary 1** (When Experimentation Is Valuable)**.** *In the one-shot problem of Proposition 5, a perfect information experiment on factor $i$ has positive value only if:*

*(a) $V\Pi \geq C$ and $c_i < C(1 - \pi_i)$ (the experiment is cheaper than the expected savings from avoiding commitment when the factor fails); or*

*(b) $V\Pi < C$, $V\Pi_{-i} \geq C$, and $c_i < \pi_i(V\Pi_{-i} - C)$ (the experiment can make GO feasible and is cheaper than the resulting option value).*

*Proof.* Directly from the piecewise characterisation in Proposition 5. In the GO case, $\text{VOI}_i = C(1 - \pi_i) - c_i$, which is positive iff $c_i < C(1 - \pi_i)$. In the NO GO case with $V\Pi_{-i} \geq C$, $\text{VOI}_i = \pi_i(V\Pi_{-i} - C) - c_i$, which is positive iff $c_i < \pi_i(V\Pi_{-i} - C)$. If $V\Pi < C$ and $V\Pi_{-i} < C$, then $\text{VOI}_i = -c_i < 0$. □

## 4.2 Imperfect Information Experiments

In practice, experiments rarely provide perfect information. A customer interview may suggest interest without guaranteeing purchase. A prototype may demonstrate technical feasibility without proving scalability. A pilot programme may indicate market demand without confirming unit economics at scale. These imperfect experiments provide noisy signals that update beliefs without fully resolving them.

I model an imperfect experiment on factor $i$ as a *binary signal experiment* with cost $c_i$ that generates signal $s \in \{+, -\}$ with likelihood $P(s = +|\theta_i = 1) = \alpha$ (true positive rate) and $P(s = -|\theta_i = 0) = \beta$ (true negative rate), where $\alpha, \beta \in (0, 1)$. The experiment is *informative* if $\alpha + \beta > 1$, ensuring that the signal is better than random guessing. As $\alpha$ and $\beta$ approach 1, the signal approaches perfect information.

**Proposition 6** (Bayesian Update from Signal)**.** *After observing signal $s$ from an experiment on factor $i$, the posterior belief is:*

$$\pi_i^+ \equiv P(\theta_i = 1|s = +) = \frac{\alpha\pi_i}{\alpha\pi_i + (1-\beta)(1-\pi_i)} \tag{5}$$

$$\pi_i^- \equiv P(\theta_i = 1|s = -) = \frac{(1-\alpha)\pi_i}{(1-\alpha)\pi_i + \beta(1-\pi_i)} \tag{6}$$

*Proof.* By Bayes' rule:

$$P(\theta_i = 1|s = +) = \frac{P(s = +|\theta_i = 1)P(\theta_i = 1)}{P(s = +)} = \frac{\alpha\pi_i}{\alpha\pi_i + (1-\beta)(1-\pi_i)}$$

The formula for $\pi_i^-$ follows similarly. □

The Bayesian update has several properties that clarify how imperfect signals interact with the entrepreneur's decision problem. The next result establishes that informative signals shift beliefs in the expected direction, preserve the prior mean, and nest perfect information as a limiting case.

**Proposition 7** (Properties of Informative Signals)**.** *For an informative experiment ($\alpha + \beta > 1$):*

*(a) $\pi_i^+ > \pi_i > \pi_i^-$: a positive signal increases belief, a negative signal decreases it.*

*(b) The posterior is a mean-preserving spread: $P(s = +)\pi_i^+ + P(s = -)\pi_i^- = \pi_i$.*

*(c) As $\alpha, \beta \to 1$: $\pi_i^+ \to 1$ and $\pi_i^- \to 0$ (approaching perfect information).*

*Proof.* For (a), we show $\pi_i^+ > \pi_i$. We need:

$$\frac{\alpha\pi_i}{\alpha\pi_i + (1-\beta)(1-\pi_i)} > \pi_i$$

Since $\pi_i \in (0,1)$, this is equivalent to:

$$\alpha(1-\pi_i) > (1-\beta)(1-\pi_i)$$

Since $1-\pi_i > 0$, this holds iff $\alpha > 1-\beta$, i.e., $\alpha+\beta > 1$, which is our informativeness condition. The proof that $\pi_i^- < \pi_i$ is analogous.

For (b), let $P(s=+) = \alpha\pi_i + (1-\beta)(1-\pi_i)$ and $P(s=-) = (1-\alpha)\pi_i + \beta(1-\pi_i)$. Then: $P(s=+)\pi_i^+ + P(s=-)\pi_i^- = \alpha\pi_i + (1-\alpha)\pi_i = \pi_i$. We verify that the posterior is a mean-preserving spread. The probability of a positive signal is:

$$\begin{aligned} P(s=+) &= P(s=+|\theta_i=1)\pi_i + P(s=+|\theta_i=0)(1-\pi_i) \\ &= \alpha\pi_i + (1-\beta)(1-\pi_i) \end{aligned}$$

Similarly: $P(s=-) = (1-\alpha)\pi_i + \beta(1-\pi_i)$. The expected posterior is:

$$\begin{aligned} P(s=+)\pi_i^+ + P(s=-)\pi_i^- &= P(s=+)\cdot\frac{\alpha\pi_i}{P(s=+)} + P(s=-)\cdot\frac{(1-\alpha)\pi_i}{P(s=-)} \\ &= \alpha\pi_i + (1-\alpha)\pi_i = \pi_i \end{aligned}$$

This confirms that information does not change the mean belief (law of iterated expectations) but spreads the posterior distribution.

For (c), as $\alpha \to 1$ and $\beta \to 1$: $\pi_i^+ = \frac{1\cdot\pi_i}{1\cdot\pi_i + 0\cdot(1-\pi_i)} = 1$ and $\pi_i^- = \frac{0\cdot\pi_i}{0\cdot\pi_i + 1\cdot(1-\pi_i)} = 0$. □

Part (b) is an instance of the law of iterated expectations: on average, the experiment does not change the entrepreneur's beliefs. However, information is valuable precisely because it creates a *spread* in beliefs: after a positive signal, the entrepreneur is more confident and more likely to choose GO; after a negative signal, less confident and more likely to abandon. It is this conditional adaptation of decisions to information that generates value.

Given these properties, a natural question is how the value of an imperfect experiment compares to the value of its perfect-information counterpart. The following result establishes that perfect information provides an upper bound on the value of any experiment about the same factor, with the bound approached as the signal becomes arbitrarily accurate.

**Proposition 8** (Value of Imperfect Information). *The value of an imperfect experiment is* *bounded above by the value of the corresponding perfect experiment:*

$$VOI_i^{imperfect} \leq VOI_i^{perfect}$$

*with equality as $\alpha, \beta \to 1$.*

*Proof.* Fix factor $i$ and hold $\pi_{-i}$ fixed, writing $\Pi_{-i} = \prod_{j \neq i} \pi_j$. For any posterior success probability $x \in [0, 1]$ for factor $i$, define the optimal stop payoff $u(x) := \max(V\Pi_{-i}x - C, 0)$. This function is convex on $[0, 1]$ (it is the maximum of two affine functions) and satisfies $u(0) = 0$.

Let an imperfect experiment about $\theta_i$ generate a signal $s$ and posterior $x_s = \pi_i(s) \in [0, 1]$. By the law of total expectation, $\mathbb{E}[x_s] = \pi_i$. For each $x \in [0, 1]$, we can write $x = x \cdot 1 + (1 - x) \cdot 0$. Convexity and $u(0) = 0$ imply

$$u(x) = u(x \cdot 1 + (1 - x) \cdot 0) \leq x\, u(1) + (1 - x)\, u(0) = x\, u(1).$$ [3]

Hence, $u(x_s) \leq x_s u(1)$ almost surely, and, therefore,

$$\mathbb{E}[u(x_s)] \leq u(1)\, \mathbb{E}[x_s] = \pi_i u(1).$$

The right-hand side is exactly the expected stop payoff under perfect information, since perfect information yields posterior 1 with probability $\pi_i$ and 0 with probability $1 - \pi_i$.

Both experiments have the same cost $c_i$ and the same no-experiment baseline $\max(V\Pi - C, 0) = u(\pi_i)$, so $\text{VOI}_i^{\text{imperfect}} \leq \text{VOI}_i^{\text{perfect}}$. As $\alpha, \beta \to 1$, the binary signal becomes perfect, and equality obtains. □

This result provides a useful benchmark: perfect information provides an upper bound on the value of any noisy experiment about the same factor, and the bound is approached as the signal becomes arbitrarily accurate. But it also implies that imperfect experiments can still be worthwhile, provided they are sufficiently cheap relative to their informational content. An entrepreneur who cannot afford to build a full prototype (perfect information) may still

[3] The mapping $\pi \mapsto \max(V \prod_j \pi_j - C, 0)$ is generally *not* convex in the full vector $\pi$ when $N \geq 2$, so Jensen's inequality cannot be applied directly in multiple dimensions. However, for any fixed $i$ and fixed $\pi_{-i}$, the optimal stop payoff as a function of the single coordinate $\pi_i$, $u(\pi_i) := \max(V\Pi_{-i}\pi_i - C, 0)$, is convex on $[0, 1]$ (it is the maximum of two affine functions) and satisfies $u(0) = 0$. If an experiment about factor $i$ induces a posterior random variable $\pi_i(s)$ with $\mathbb{E}[\pi_i(s)] = \pi_i$, then Jensen implies $\mathbb{E}[u(\pi_i(s))] \geq u(\pi_i)$, meaning that information weakly increases the optimal stop payoff before accounting for experiment cost. Moreover, convexity plus $u(0) = 0$ gives the pointwise bound $u(x) \leq x\, u(1)$ for all $x \in [0, 1]$, which implies $\mathbb{E}[u(\pi_i(s))] \leq \pi_i u(1)$. The right-hand side is the expected stop payoff under perfect information, and this is the key inequality used in Proposition 8.

benefit from running customer interviews (imperfect information), as long as the interviews are informative enough to justify their cost.

# 5 Optimal Experimentation Strategy

## 5.1 Sequential Experimentation

The entrepreneur can run experiments sequentially, updating beliefs and deciding at each stage whether to continue experimenting, commit (GO), or abandon (NO GO). This sequential structure is a distinctive feature of entrepreneurship: unlike a one-shot investment decision, the entrepreneur can adaptively choose what to learn and when to stop learning (Kerr et al., 2014; Manso, 2011).

At each stage, the entrepreneur chooses which factor to experiment on (if any) and whether to stop and choose GO or NO GO, aiming to maximise expected payoff net of experimentation costs.

**Proposition 9** (Optimal Stopping Characterisation)**.** *In the sequential experimentation problem, beliefs* $\pi = (\pi_1, \ldots, \pi_N)$ *are a sufficient state variable. Let* $\Pi = \prod_{j=1}^{N} \pi_j$ *and let* $W(\pi)$ *denote the optimal expected payoff starting from beliefs* $\pi$*. Then* $W$ *satisfies the Bellman equation*

$$W(\pi) = \max \left\{ \underbrace{\max(V\Pi - C, 0)}_{stop},\ \max_i \underbrace{\left(-c_i + \mathbb{E}\left[W(\pi^{(i)}(s))\right]\right)}_{experiment\ on\ i} \right\},$$

*where* $\pi^{(i)}(s)$ *denotes the posterior after running experiment* $i$ *and observing outcome* $s$*.*

*An optimal policy is:*

*(a) Stop if* $\max(V\Pi - C, 0)$ *is at least as large as the best continuation value* $\max_i\{-c_i + \mathbb{E}[W(\pi^{(i)}(s))]\}$*. Upon stopping, choose GO if* $V\Pi \geq C$ *and choose NO GO otherwise.*

*(b) Otherwise, run an experiment* $i$ *that attains the maximum continuation value.*

*Proof.* At any stage, the current belief vector $\pi$ summarises all payoff-relevant information. If the entrepreneur stops immediately, the optimal payoff is $S(\pi) = \max(V\Pi - C, 0)$, since GO yields expected payoff $V\Pi - C$ and NO GO yields 0.

If instead the entrepreneur runs experiment $i$, they pay $c_i$ and observe an outcome $s$, inducing posterior beliefs $\pi^{(i)}(s)$. From that point onward, optimal behaviour yields continuation value $W(\pi^{(i)}(s))$. Taking expectations over $s$ gives the continuation value $Q_i(\pi) = -c_i + \mathbb{E}\big[W(\pi^{(i)}(s))\big]$.

By the principle of optimality, the value at state $\pi$ is the maximum of stopping and the best continuation action:

$$W(\pi) = \max\left\{S(\pi), \max_i Q_i(\pi)\right\},$$

which is exactly the stated Bellman equation. The policy description follows by choosing an action that attains the maximand. □

The Bellman equation makes explicit the dynamic trade-off at the heart of the entrepreneurial problem. At each stage, the entrepreneur weighs the value of acting now (committing or abandoning) against the value of learning more (running another experiment). This is a formalisation of the commitment versus flexibility trade-off emphasised in the entrepreneurial strategy literature (Gans et al., 2019): commitment captures value but forecloses options, while experimentation preserves options but incurs costs and delay.

## 5.2 Experiment Sequencing

When multiple experiments have positive value, which should be conducted first? The sequencing question is practically important: running experiments in the wrong order wastes resources by investing in experiments whose results may be rendered moot by earlier failures.

To isolate the sequencing logic, assume experiments take negligible time, and there is no discounting or cost of delay between them.

**Proposition 10** (Experiment Sequencing with Perfect Information)**.** *Suppose the entrepreneur currently should GO but has multiple factors where experimentation is valuable (i.e., $c_i < C(1-\pi_i)$ for several factors). The optimal sequence satisfies:*

*(a) Experiments should be ordered by decreasing "kill rate" $(1-\pi_i)$, provided cost differences are small.*

*(b) More precisely, experiment on factor $i$ before factor $j$ if:*

$$\frac{c_i}{1-\pi_i} < \frac{c_j}{1-\pi_j}$$

*Proof.* Consider two factors $i$ and $j$ where experimentation is valuable. The entrepreneur will run both experiments (in some order) unless one kills the project. **Order $i$ then $j$:**

Expected cost $= c_i + \pi_i \cdot c_j$. We always pay $c_i$, and pay $c_j$ only if factor $i$ passes (probability $\pi_i$).

**Order $j$ then $i$:** Expected cost $= c_j + \pi_j \cdot c_i$. The remaining expected payoff, conditional on both factors passing, is the same either way. So we prefer $i$ first if and only if:

$$c_i + \pi_i c_j < c_j + \pi_j c_i \Leftrightarrow c_i(1 - \pi_j) < c_j(1 - \pi_i) \Leftrightarrow \frac{c_i}{1-\pi_i} < \frac{c_j}{1-\pi_j}$$

The ratio $c_i/(1 - \pi_i)$ is the cost per unit kill probability. We should run cheaper (per unit kill probability) experiments first. □

An immediate and practically important special case arises when all experiments cost the same. The sequencing rule then reduces to a simple ordering by uncertainty: the entrepreneur should test the factor most likely to kill the project first.

**Corollary 2** (Kill Quickly, Fail Cheaply)**.** *If all experiments have equal cost, run the most uncertain experiments first (lowest $\pi_i$, highest kill probability).*

*Proof.* With $c_i = c$ for all $i$, the ordering condition becomes $(1 - \pi_i) > (1 - \pi_j)$, i.e., $\pi_i < \pi_j$. □

Note that Corollary 2 applies in the "pre-mortem" GO regime of Proposition 10, where the entrepreneur would otherwise commit and experimentation is valuable mainly because it can stop a bad investment. In the search regime ($V\Pi < C$), the relevant ordering can instead prioritise experiments with the fastest expected path to crossing the viability frontier, especially when runway or time-to-market constraints matter.

Within the "pre-mortem" GO domain, the "kill quickly, fail cheaply" principle formalises a widely held intuition in entrepreneurial practice.[4] It is better to spend a small amount discovering a fatal flaw early than to spend a large amount building toward one. The sequencing result makes this intuition precise: the optimal ordering minimises expected experimentation costs by frontloading experiments that are most likely to terminate the venture.

This principle is illustrated by contrasting cases from entrepreneurial practice. Drew Houston, the founder of Dropbox, faced two distinct risks: the technology was genuinely hard (reliable file synchronisation across devices was an unsolved problem), and demand was uncertain (perhaps people did not care enough to switch from emailing files to themselves). Rather than spending years solving the technical problem only to discover that nobody cared, Houston made a three-minute demonstration video showing how Dropbox would work. The

[4] Also empirically confirmed in the recent study by Bailey et al. (2026).

waiting list jumped from 5,000 to 75,000 people overnight. He had tested the uncertain demand factor for the cost of a video, before committing to the expensive technical factor (Ries, 2011).

By contrast, Quibi—a premium short-form video service—raised \$1.75 billion and built its entire infrastructure before testing whether consumers actually wanted short-form premium content behind a paywall. The venture lasted six months. The cheap killer information ("Will people pay for short-form video when TikTok is free?") was never acquired before the expensive commitment was made.

## 5.3 Portfolio Allocation of Experimental Effort

When learning reduces uncertainty gradually (rather than providing perfect information), the entrepreneur faces a portfolio allocation problem: which factors should receive attention, and how much, in order to reach a viable GO decision.

To align the continuous-effort problem with the binary commitment decision, consider the minimum-effort problem of reaching the viability frontier. Starting from initial log-risks $\lambda_i^0$, choose effort levels $e_i \geq 0$ that reduce log-risk to $\lambda_i(e_i)$ and solve:

$$\min_{e_i \geq 0} \sum_{i=1}^{N} e_i \quad \text{s.t.} \quad \sum_{i=1}^{N} \lambda_i(e_i) \leq B. \tag{7}$$

This formulation asks: what is the cheapest way (in total experimental effort) to buy enough log-risk reduction to fit within the miracle budget?

For tractability, assume *exponential de-risking*:

$$\lambda_i(e_i) = \lambda_i^0 e^{-\kappa_i e_i},$$

where $\kappa_i > 0$ captures how quickly effort reduces uncertainty on factor $i$ (a higher $\kappa_i$ means factor $i$ is easier to de-risk per unit of effort). The exponential form captures a common feature of learning: early investigation produces large reductions in uncertainty, but further effort yields diminishing improvements.

**Proposition 11** (Minimum-Effort De-risking to Reach Viability)**.** *Under exponential de-risking, an optimal allocation for problem* (7) *satisfies*

$$\kappa_i \lambda_i^0 e^{-\kappa_i e_i^*} = \frac{1}{\nu} \quad \textit{for all } i \textit{ with } e_i^* > 0,$$

*for some multiplier $\nu > 0$ chosen so that $\sum_i \lambda_i(e_i^*) = B$ whenever the constraint binds.*

*Equivalently,*

$$e_i^* = \frac{1}{\kappa_i}\left[\log\big(\nu\,\kappa_i\lambda_i^0\big)\right]^+, \qquad \lambda_i^* \equiv \lambda_i(e_i^*) = \frac{1}{\nu\,\kappa_i} \quad \textit{for all active } i.$$

*Proof.* With $\lambda_i(e_i) = \lambda_i^0 e^{-\kappa_i e_i}$, problem (7) is convex. The Lagrangian is

$$\mathcal{L} = \sum_i e_i + \nu\left(\sum_i \lambda_i^0 e^{-\kappa_i e_i} - B\right) - \sum_i \eta_i e_i,$$

with $\nu \geq 0$ and $\eta_i \geq 0$. The first-order condition is

$$\frac{\partial \mathcal{L}}{\partial e_i} = 1 - \nu\,\kappa_i\lambda_i^0 e^{-\kappa_i e_i} - \eta_i = 0.$$

If $e_i^* > 0$, then $\eta_i = 0$ and $1 = \nu\,\kappa_i\lambda_i^0 e^{-\kappa_i e_i^*}$, yielding the stated condition. Solving for $e_i^*$ gives

$$e_i^* = \frac{1}{\kappa_i}\log\big(\nu\,\kappa_i\lambda_i^0\big),$$

which is positive if and only if $\nu\,\kappa_i\lambda_i^0 > 1$; otherwise the complementary slackness conditions imply $e_i^* = 0$, giving the $[\cdot]^+$ form. For any active factor,

$$\lambda_i^* = \lambda_i^0 e^{-\kappa_i e_i^*} = \lambda_i^0 \cdot \frac{1}{\nu\,\kappa_i\lambda_i^0} = \frac{1}{\nu\,\kappa_i}.$$

□

Proposition 11 has direct prescriptions for "how to get to GO" when effort is scarce: allocate attention to the factors that deliver the largest feasible reductions in log-risk per unit effort.

**Corollary 3** (Allocation Principles for Reaching GO)**.** *Under optimal minimum-effort de-risking:*

*(a) Factors are active ($e_i^* > 0$) only if their "bang-for-buck" index $\kappa_i\lambda_i^0$ is sufficiently large relative to the multiplier $\nu$.*

*(b) Among active factors, residual log-risk satisfies $\lambda_i^* = 1/(\nu\kappa_i)$, so easier-to-de-risk factors (higher $\kappa_i$) are driven to lower residual uncertainty.*

*(c) At the optimum, the marginal reduction in total log-risk per unit effort is equalised across all active factors:*

$$-\frac{\partial}{\partial e_i}\lambda_i(e_i^*) = \kappa_i\lambda_i^* = \frac{1}{\nu} \quad \textit{for all } i \textit{ with } e_i^* > 0.$$

*Proof.* All parts follow directly from Proposition 11 and the KKT conditions. Part (a) uses $e_i^* > 0 \iff \nu\kappa_i\lambda_i^0 > 1$. Parts (b) and (c) follow from $\lambda_i^* = 1/(\nu\kappa_i)$ and the first-order condition $1 = \nu\kappa_i\lambda_i(e_i^*)$ for active factors. □

Relative to "minimise $\sum_i \lambda_i$ given a fixed effort budget," this formulation makes the stopping logic explicit: the entrepreneur expends effort only until the venture reaches $\sum_i \lambda_i \le B$, at which point additional de-risking is wasted from the narrow perspective of reaching GO.[5]

## 5.4 Diminishing Returns

The equimarginal principle of Corollary 3 relies on a deeper structural property of the de-risking technology: the marginal reduction in log-risk from an additional unit of effort is decreasing. This diminishing-returns property has direct implications for where the entrepreneur should concentrate effort at any given point in the experimentation process.

**Proposition 12** (Diminishing Returns to De-risking)**.** *In the continuous de-risking model, suppose each factor has log-risk $\lambda_i(e_i) = \lambda_i^0 g_i(e_i)$ where $g_i$ is decreasing and convex on $[0,\infty)$. Then $\lambda_i(e_i)$ is decreasing and convex in $e_i$, so the marginal reduction in log-risk exhibits diminishing returns:*

$$\frac{\partial}{\partial e_i}\left(-\frac{\partial \lambda_i}{\partial e_i}\right) \le 0.$$

*In particular, under exponential de-risking $\lambda_i(e_i) = \lambda_i^0 e^{-\kappa_i e_i}$,*

$$-\frac{\partial \lambda_i}{\partial e_i} = \kappa_i \lambda_i(e_i),$$

*so (holding $\kappa_i$ fixed) factors with higher current log-risk $\lambda_i$ yield a larger marginal reduction per unit effort.*

*Proof.* If $g_i$ is decreasing and convex, then $\lambda_i(e_i) = \lambda_i^0 g_i(e_i)$ is also decreasing and convex in $e_i$, so $\frac{\partial^2 \lambda_i}{\partial e_i^2} \ge 0$. Therefore

$$\frac{\partial}{\partial e_i}\left(-\frac{\partial \lambda_i}{\partial e_i}\right) = -\frac{\partial^2 \lambda_i}{\partial e_i^2} \le 0,$$

which is the diminishing-returns statement.

For the exponential specification $\lambda_i(e_i) = \lambda_i^0 e^{-\kappa_i e_i}$,

$$-\frac{\partial \lambda_i}{\partial e_i} = \kappa_i \lambda_i^0 e^{-\kappa_i e_i} = \kappa_i \lambda_i(e_i),$$

[5] Entrepreneurs may still de-risk beyond the frontier for strategic, financing, or execution reasons; problem (7) isolates the minimum effort required to make GO privately optimal under the model's static payoff structure.

so the marginal reduction is proportional to current $\lambda_i(e_i)$. □

The diminishing returns result has a clear practical implication: improving the strongest factor is rarely the right move. Reducing a factor from $\pi = 0.3$ to $\pi = 0.5$ (moving log-risk from 1.20 to 0.69, saving 0.51 units of budget) frees far more budget than reducing a factor from $\pi = 0.8$ to $\pi = 0.9$ (moving log-risk from 0.22 to 0.11, saving only 0.11 units). Entrepreneurs should concentrate effort on the factors with the highest marginal reductions in log-risk, which under exponential de-risking are those with high current log-risk (and, all else equal, high $\kappa_i$).

# 6 Distinguishing Real from Apparent Miracles

An important aspect of entrepreneurial learning is distinguishing genuinely uncertain factors from factors that merely *appear* uncertain due to lack of information. This distinction echoes Knight (1921)'s contrast between risk and uncertainty, but here the emphasis is operational: entrepreneurs often do not know whether a perceived risk is a hard unknown or largely an information gap that can be narrowed through investigation. A former Apache helicopter pilot turned entrepreneur (of Jane Technologies), Socrates Rosenfeld, captured this distinction vividly: "When you're flying a helicopter you have to understand risk, and understand the differences between risk and uncertainty. On missions you are going into uncertain areas, into uncertain environments, and that is different than risk." (Stern and Zyontz, 2019)

A useful implication for entrepreneurs is that the number of genuinely unresolved contingencies is often smaller than it initially appears. What may seem like an imposing set of uncertainties frequently reflects a combination of truly uncertain "miracles" and factors whose apparent uncertainty is attributable to incomplete information and can be reduced through investigation. Accordingly, the entrepreneur's central task is not to "will" favourable outcomes into existence, but to identify which contingencies are genuinely pivotal and to determine how many such miracles the venture in fact requires.

## 6.1 A Hierarchical Model

To formalise the distinction, I introduce a simple hierarchical (mixture) structure. For each factor $i$, let $\gamma_i \in \{0, 1\}$ indicate whether the factor is a "real miracle." If $\gamma_i = 1$, factor $i$ is genuinely uncertain and, even after investigation, succeeds only with baseline probability $p \in (0, 1)$. If $\gamma_i = 0$, factor $i$ is an *apparent miracle*: once properly understood it succeeds with high probability $p_0 \in (p, 1)$, though not with certainty.[6] Let $q_i = P(\gamma_i = 1)$ be the entrepreneur's prior that factor $i$ is a real miracle.

Assume that the pairs $(\gamma_i, \theta_i)$ are independent across $i$ in the entrepreneur's prior, so that overall venture success remains the product of the marginal beliefs. An *apparent miracle* is a factor that looks uncertain because the entrepreneur has not done the work yet. It can be substantially resolved through research, prototypes, customer conversations, technical investigation, or clever restructuring of the business model. The uncertainty is about the entrepreneur's knowledge, not about the world. When one digs into an apparent miracle, one typically finds that the answer already exists somewhere—in the technical literature, in customer behaviour, in competitive dynamics—and the task is simply to uncover it.

A *real miracle* is different. It is a factor that remains genuinely uncertain even after investigation. Real miracles typically involve market acceptance at scale, timing that depends on forces outside the entrepreneur's control, competitive responses that cannot be predicted, or execution challenges that only reveal themselves under pressure. The uncertainty is about what the world will do, not about what the entrepreneur could learn. No amount of research will resolve a real miracle because the answer does not yet exist—it will only be created when the entrepreneur tries.

**Proposition 13** (Belief Structure)**.** *Under the hierarchical model, the entrepreneur's marginal belief about factor $i$ is:*

$$\pi_i = P(\theta_i = 1) = q_i \cdot p + (1 - q_i) \cdot p_0 = p_0 - q_i(p_0 - p)$$

*Proof.* By the law of total probability:

$$\begin{aligned} P(\theta_i = 1) &= P(\theta_i = 1 | \gamma_i = 1)P(\gamma_i = 1) + P(\theta_i = 1 | \gamma_i = 0)P(\gamma_i = 0) \\ &= p \cdot q_i + p_0 \cdot (1 - q_i) = p_0 - q_i(p_0 - p) \end{aligned}$$

□

A *miracle discovery experiment* on factor $i$ has cost $c_i$ and reveals $\gamma_i$ (whether the factor is a real miracle). If it reveals $\gamma_i = 0$ (an apparent miracle), then the entrepreneur updates to $\pi_i' = p_0$, sharply reducing the factor's budget consumption (though not necessarily

[6] This hierarchical structure is a Bayesian mixture capturing classification uncertainty about whether a perceived risk is a hard unknown (a real miracle) or largely a research task (an apparent miracle). It is not meant as a strong formalisation of Knightian ambiguity; one could alternatively represent a real miracle with a set of priors $\pi_i \in [\underline{p}, \overline{p}]$ and study ambiguity-averse policies. I also abstract from factors that are certainly impossible (e.g., $P(\theta_i = 1) = 0$), which can be treated in the baseline framework as factors with $\pi_i = 0$ and screened out immediately.

to zero). If it reveals $\gamma_i = 1$ (a real miracle), the experiment classifies the factor as irreducibly uncertain and the posterior remains $\pi_i' = p$. In both branches the realisation of $\theta_i$ remains unknown; what changes is the relevant probability for planning. This captures investigations aimed at determining whether a perceived risk is "real" or mostly an artefact of incomplete understanding. For example, technical diligence might reveal that a supposedly risky component relies on well-understood engineering (an apparent miracle) versus remaining a genuine R&D unknown (a real miracle). Likewise, early market or regulatory research can distinguish whether adoption or approval is effectively routine conditional on basic compliance (apparent) versus genuinely uncertain (real).

**Proposition 14** (Value of Miracle Discovery)**.** *The value of discovering whether factor $i$ is a real miracle is:*

$$\begin{aligned} VMD_i = &-c_i + q_i \cdot \max(Vp\Pi_{-i} - C, 0) + (1 - q_i) \cdot \max(Vp_0\Pi_{-i} - C, 0) \\ &- \max(V\pi_i\Pi_{-i} - C, 0) \end{aligned}$$

*where* $\Pi_{-i} = \prod_{j \neq i} \pi_j$.

*Proof.* After discovering $\gamma_i$:

- If $\gamma_i = 1$ (real miracle): update $\pi_i' = p$, then face GO decision with success prob $p\Pi_{-i}$.
- If $\gamma_i = 0$ (apparent miracle): update $\pi_i' = p_0$, then face GO decision with success prob $p_0\Pi_{-i}$.

Expected payoff from discovery:

$$\mathbb{E}[\text{Discovery}] = -c_i + q_i \max(Vp\Pi_{-i} - C, 0) + (1 - q_i) \max(Vp_0\Pi_{-i} - C, 0)$$

Payoff without discovery:

$$\mathbb{E}[\text{No Discovery}] = \max(V\pi_i\Pi_{-i} - C, 0)$$

The VMD is the difference. □

Miracle discovery experiments can be valuable even when they do not directly resolve whether a factor will work. Learning that something is *not* a miracle (it just appeared uncertain) provides reassurance without resolving the underlying state. This is because discovering $\gamma_i = 0$ shifts beliefs toward $p_0$, sharply reducing that factor's budget consumption and sometimes tipping the venture into GO.

## 6.2 Counting Miracles

The expected number of real miracles is $M = \sum_{i=1}^{N} q_i$. This summary statistic provides a tractable measure of the venture's total miracle exposure. How does the GO condition depend on $M$? Because the mixture probabilities $\pi_i = 1 - q_i(1-p)$ are neither all equal nor all binary, the exact GO condition involves the full vector $(q_1, \ldots, q_N)$. However, tight bounds expressed in terms of the aggregate $M$ alone can be obtained, providing simple sufficient conditions for viability and for failure.

**Proposition 15** (Miracle Count Bounds in the Hierarchical Model)**.** *If real miracles have base rate $p \in (0,1)$ and apparent miracles have base rate $p_0 \in (p, 1)$, then the GO condition in the hierarchical model is*

$$\prod_{i=1}^{N} \big((1-q_i)p_0 + q_i p\big) \geq \frac{C}{V}.$$

*Let $M = \sum_{i=1}^{N} q_i$ denote the expected number of real miracles. Then:*

*(a) (**Sufficient for GO**) If $p^M p_0^{N-M} \geq \frac{C}{V}$, then the GO condition holds.*

*(b) (**Sufficient for NO GO**) If $p_0^N \exp\big(-(1-p/p_0)M\big) < \frac{C}{V}$, then the GO condition fails.*

*Proof.* Under the hierarchical model, factor $i$ is a real miracle ($\gamma_i = 1$) with probability $q_i$ and otherwise an apparent miracle ($\gamma_i = 0$). If $\gamma_i = 1$ then $P(\theta_i = 1) = p$; if $\gamma_i = 0$ then $P(\theta_i = 1) = p_0$. Hence

$$\pi_i = q_i p + (1-q_i)p_0 = (1-q_i)p_0 + q_i p,$$

and the GO condition $\prod_i \pi_i \geq C/V$ becomes $\prod_i \big((1-q_i)p_0 + q_i p\big) \geq C/V$.

For part (a), weighted AM–GM gives

$$(1-q_i)p_0 + q_i p \;\geq\; p_0^{1-q_i} p^{q_i}.$$

Multiplying across $i$ yields

$$\prod_{i=1}^{N} \big((1-q_i)p_0 + q_i p\big) \geq \prod_{i=1}^{N} p_0^{1-q_i} p^{q_i} = p_0^{\sum_i (1-q_i)} p^{\sum_i q_i} = p_0^{N-M} p^M,$$

so if $p_0^{N-M} p^M \geq C/V$ then the GO condition holds.

For part (b), rewrite $\pi_i = p_0\big(1 - q_i(1-p/p_0)\big)$. For $x \in [0,1)$ we have $\log(1-x) \leq -x$.

Applying this with $x = q_i(1 - p/p_0)$ gives

$$\log\bigl(1 - q_i(1 - p/p_0)\bigr) \le -(1 - p/p_0)q_i.$$

Summing over $i$ and exponentiating,

$$\prod_{i=1}^{N} \pi_i = p_0^N \prod_{i=1}^{N} \bigl(1 - q_i(1 - p/p_0)\bigr) \le p_0^N \exp\left(-(1 - p/p_0)\sum_{i=1}^{N} q_i\right) = p_0^N \exp\bigl(-(1 - p/p_0)M\bigr).$$

Therefore if $p_0^N \exp(-(1-p/p_0)M) < C/V$, then $\prod_i \pi_i < C/V$ and the GO condition fails. $\square$

The sufficient condition for GO in part (a) provides a simple viability check: if the venture would be viable under the pessimistic assumption that every uncertain factor is a real miracle (each with probability $p$), then it is certainly viable under the actual mix of real and apparent miracles. This gives entrepreneurs a conservative rule of thumb: treat apparent miracles as high-probability ($p_0$) rather than certain, count real miracles conservatively, and check whether $p^M p_0^{N-M} \ge C/V$.

## 6.3 Case Studies: Restructuring and Discovery

The distinction between real and apparent miracles is not merely theoretical—it directly shapes entrepreneurial strategy. Two cases illustrate the practical implications.

### 6.3.1 Restructuring to Eliminate Miracles: Jane Technologies

Socrates Rosenfeld, after leaving the US Army, sought to build a cannabis e-commerce company in 2015. At that time, cannabis remained federally illegal in the United States, banking was nearly impossible for cannabis businesses, and the regulatory environment was a patchwork of state-by-state rules. The apparent miracle count was daunting: federal legalisation, consumer trust in an online cannabis platform, dispensary adoption of new technology, banking for cannabis transactions, and regulatory permission for e-commerce in a controlled substance. Each looked like a coin flip at best.

Rosenfeld made a strategic choice that transformed the problem (Stern and Zyontz, 2019). Rather than building the "Amazon of cannabis"—a centralised marketplace—he built what he called the "Shopify of cannabis." Jane Technologies would not sell cannabis. It would power other dispensaries' storefronts, providing e-commerce infrastructure, inventory management, and analytics. The actual sale would happen through local dispensaries, using their licences, their banking relationships, and their customer trust.

This architectural decision converted multiple miracles into someone else's problem. Federal legalisation became less critical because Jane operated state by state through legally licensed dispensaries. Consumer trust stayed with the local retailer—customers were buying from their neighbourhood dispensary, not from an unknown technology company. Banking flowed through the dispensaries' accounts, not Jane's. What Rosenfeld kept were two genuine uncertainties: could the team build a platform that worked, and would dispensaries adopt it? These were still real uncertainties, but they were uncertainties that could be tested, iterated on, and resolved through direct engagement.

In the language of the model, Rosenfeld reduced his miracle count $M$ not by resolving the miracles through investigation, but by restructuring the venture to make them irrelevant—effectively setting $q_i = 0$ for the factors he re-assigned to the dispensaries. By 2021, Jane's platform powered roughly 20 per cent of all legal cannabis sales in the United States, with over 2,500 dispensary partners across 39 states.

### 6.3.2 Discovering the Real Constraint: Sheertex

Katherine Homuth set out to create indestructible pantyhose—a product that would end the frustration of hosiery that rarely survived a single wearing. The surface-level miracle was obvious: can you actually make a sheer fabric that does not tear? Most investors assumed the answer was no (Gans et al., 2024).

Homuth attacked the apparent miracle directly. She researched industrial polymers, discovered ultra-high-molecular-weight polyethylene (the material used in bulletproof vests), and eventually produced a fabric that was, as she described it, "insanely strong." Each subsequent engineering challenge—making the material sheer, getting it to stretch, achieving the right colour—turned out to be an engineering problem with a solution. The apparent miracle resolved through investigation: $\gamma_i = 0$ for the technology factors.

But the investigation revealed the real miracle—the one nobody had asked about initially. The technology worked, and consumers wanted the product (a Kickstarter campaign raised nearly \$200,000). The constraint was manufacturing cost. "Our tights once cost over one hundred dollars each to make," Homuth reflected. "Today, they cost thirteen dollars. By next year, they'll cost five. But we're still fighting tooth and nail to survive." The company has raised more than \$250 million and built one of the most advanced textile factories in North America, yet profitability remains elusive.

In the model's terms, the miracle everyone focused on (whether $\gamma_{\text{tech}} = 1$) turned out to be an apparent miracle: $\gamma_{\text{tech}} = 0$. The miracle nobody asked about (whether $\gamma_{\text{cost}} = 1$) turned out to be the real one. Identifying the true constraint earlier would not have changed the physics, but it might have changed the company's resource allocation and fundraising strategy.

These cases illustrate a general principle: the miracle that consumes entrepreneurial attention may not be the miracle that determines the venture's fate. Rosenfeld succeeded by making miracles irrelevant through restructuring. Homuth discovered that the visible miracle was illusory and the invisible one was real. Both cases demonstrate that the entrepreneur's primary task is not to make miracles happen, but to learn which factors are genuinely miraculous.

# 7 Discussion: Principles and Implications

## 7.1 Principles for Entrepreneurs

The formal results yield several actionable principles that synthesise the framework's implications for practice. The first is to *know your budget*. Before evaluating any factor, the entrepreneur should compute $B = \log(V/C)$. This determines how much total uncertainty the venture can carry. Large payoffs expand the budget only logarithmically: a 100x opportunity is not permission to believe in ten uncertain things. Relative to a 10x opportunity, it buys only about three additional coin-flip miracles (e.g., roughly six instead of three under the $\pi = 0.5$ benchmark).

The second is to *focus on pivotal factors*. As a first-pass rule, the entrepreneur should prioritise experiments whose outcome could change the GO/NO GO decision. As a log-risk diagnostic, a factor is pivotal when its favourable resolution could close the current slack $-S(\lambda)$ and bring the venture to the frontier (Proposition 1). In the one-shot problem of Proposition 5, an experiment on factor $i$ is worth running only if: the venture is currently at NO GO but $V\Pi_{-i} \geq C$ (learning $i$ works could enable GO), or the venture is currently at GO and $c_i < C(1 - \pi_i)$ (the option value of avoiding bad commitment exceeds cost). Several individually non-pivotal factors can still be worth testing jointly in a sequential plan, because a sequence of successes can move the venture into the GO region.

The third is to *kill quickly and fail cheaply*. When running experiments, the entrepreneur should prioritise those with high kill probability and low cost. The optimal ordering is by increasing $c_i/(1 - \pi_i)$—the cost per unit probability of killing the project (Proposition 10). When experiment costs are roughly comparable, this reduces to running the most uncertain (highest kill-probability) tests first; otherwise, begin with the lowest $c_i/(1 - \pi_i)$ experiments before committing to expensive ones.

The fourth is to *buy the slack you need, as cheaply as possible*. When effort is scarce, allocate attention toward factors with a high "bang-for-buck" index $\kappa_i\lambda_i^0$ (easy to de-risk and currently consuming lots of budget), and stop once $\sum_i \lambda_i \leq B$. Diminishing returns (Proposition 12) imply that reducing a factor from $\pi = 0.3$ to $\pi = 0.5$ frees more budget than reducing from $\pi = 0.7$ to $\pi = 0.9$.

The fifth is to *distinguish real from apparent miracles*. Many factors that look like miracles can be de-risked through information. The entrepreneur's core skill is distinguishing real miracles (genuinely uncertain outcomes that cannot be resolved without trying) from apparent miracles (uncertainties that can be substantially resolved through experiments, research, or restructuring). Early-stage work should focus on discovering which category each factor belongs to.

Finally, *the budget is endogenous but constrained*. There is no fixed definition of "too risky." The miracle budget scales with $V/C$. A venture with 100x potential payoff can sustain more unresolved uncertainty than one with 10x payoff. However, because the budget grows only as $B = \log(V/C)$, even transformative opportunities buy surprisingly little additional slack in aggregate uncertainty.

## 7.2 Implications for Theory

The miracle budget framework connects to several strands of the entrepreneurship and strategy literatures. First, the framework provides formal content for a practical distinction that founders and investors regularly make: quantifiable multiplicative risk (captured by the log-risk budget) versus *classification uncertainty* about whether a perceived risk is a hard unknown or largely under-researched (captured by the hierarchical mixture in Section 6). Investigation can shift a factor from the "real miracle" distribution to the "apparent miracle" distribution, even if it does not fully resolve the underlying state, and the budget constraint then governs how much unresolved risk can remain.

Second, the framework formalises the core mechanism of the lean startup methodology (Ries, 2011; Blank, 2013): the idea that entrepreneurs should acquire information before committing resources. The model shows precisely when experiments are valuable (Proposition 5), what order they should be run in (Proposition 10), and how effort should be allocated across them (Proposition 11). This provides a normative foundation for practices such as minimum viable products, customer discovery, and pivoting.

Third, the results illuminate the commitment-versus-flexibility trade-off central to entrepreneurial strategy (Gans et al., 2019; Dixit and Pindyck, 1994; Folta, 1998; McGrath, 1999). The Bellman equation (Proposition 9) characterises the precise conditions under which the entrepreneur should commit (GO), abandon (NO GO), or continue experimenting. The optimal policy balances the value of additional information against the cost of delay, providing a formal version of the "leap of faith" that entrepreneurship requires.

Fourth, the hierarchical model (Section 6) contributes to the debate between "discovery" and "creation" theories of entrepreneurial opportunity (Alvarez and Barney, 2007; Shane and Venkataraman, 2000). In the discovery view, opportunities exist independently of the entrepreneur and can be found through search. In the creation view, opportunities are constructed through the entrepreneur's actions. The miracle budget framework suggests a synthesis: apparent miracles correspond to opportunities for discovery (the answer exists and can be found through investigation), while real miracles correspond to opportunities for creation (the answer does not yet exist and will only be determined through action).

## 7.3 Implications for Investors

The framework also has implications for investors evaluating early-stage ventures. An investor's role can be understood as helping to expand or manage the miracle budget. Capital expands the budget by increasing $V$ (funding growth that raises the potential payoff) or reducing $C$ (subsidising experimentation that would otherwise be too expensive). Mentorship and networks help convert apparent miracles into resolved questions, reducing $M$. Staged financing—providing capital in tranches conditional on milestones—is a natural implementation of the sequential experimentation strategy characterised in Proposition 9: each funding round represents an experiment whose outcome determines whether to continue.

The logarithmic scaling result (Proposition 2) explains a puzzle in venture capital: why experienced investors are so focused on reducing the number of uncertainties rather than on increasing the potential payoff. The model shows that $B = \log(V/C)$ grows only slowly with $V/C$, so large increases in market size translate into relatively modest increases in budget slack. In budget terms, eliminating one coin-flip miracle (removing $\lambda = \log 2$ of log-risk) creates the same slack as doubling $V/C$ (which increases $B$ by $\log 2$). De-risking can be more valuable than payoff expansion when the relevant miracle is worse than a coin flip ($\pi_i < 0.5$), when resolving one factor also reduces other correlated uncertainties, or when the investor can buy log-risk reductions more cheaply than they can expand $V/C$.

# 8 Conclusion

This paper has developed a formal framework for entrepreneurial decision-making under multiplicative uncertainty. The central contribution is the *miracle budget*: the maximum aggregate uncertainty a venture can sustain, determined by the logarithm of its payoff-to-cost ratio. Because the budget grows only as $B = \log(V/C)$, even transformative opportunities buy surprisingly little additional slack in aggregate uncertainty, providing precise content for the practitioner's intuition that ventures requiring "too many things to go right" are not viable regardless of their potential payoff.

The framework characterises optimal experimentation strategies through several results. The value of information analysis (Propositions 5 and 8) shows when experiments are worth running and establishes that imperfect information is always weakly less valuable than perfect information. The sequencing result (Proposition 10) establishes the "kill quickly, fail cheaply" principle: experiments with high kill probability and low cost should be run first. The portfolio allocation result (Proposition 11) aligns effort allocation with the GO decision by solving for the minimum experimental effort required to reach $\sum_i \lambda_i \leq B$; it implies that attention should go to factors with high "bang-for-buck" $\kappa_i \lambda_i^0$ and high current log-risk, where the marginal reduction in budget consumption is largest.

Several directions for future research follow naturally from the framework. First, the independence benchmark can be relaxed. Correlated factors, causal dependence (where some conditions only matter conditional on others), and hierarchical structures alter both the "accounting" of aggregate uncertainty and the design of informative experiments, since learning about one factor may partially resolve others. Second, payoffs and costs can be endogenised. Allowing $V$ and $C$ to evolve with time, competition, and the entrepreneur's own commitments would embed the miracle budget logic into a dynamic stopping problem in which information acquisition trades off against delay and strategic pre-emption.

Third, the model can be extended to settings with ambiguity rather than risk, where beliefs $\pi_i$ are themselves ill-defined or contested (Knight, 1921). Robust or ambiguity-averse policies would clarify when "counting miracles" should be conservative, and how much evidence is required before treating an apparent miracle as resolved. Fourth, strategic interactions deserve explicit treatment: investors, employees, and rivals can create signalling incentives that shift experimentation away from what is privately optimal, and staged financing can be interpreted as an endogenous experimentation policy.

Finally, the framework suggests several empirical questions. How many factors do founders and investors effectively treat as "real miracles" in early-stage decisions? Which experimental programmes (customer discovery, prototypes, pilots) actually reduce the product of beliefs rather than merely generate activity? And when ventures fail, is it typically because one low-probability factor was never addressed, or because many modest uncertainties compounded beyond the budget? Answering these questions would help connect the model's log-risk accounting to observed entrepreneurial practice.

## References


Agrawal, A., A. Camuffo, A. Gambardella, J. S. Gans, E. L. Scott, and S. Stern (2026): "Foundations of Bayesian Entrepreneurship," in *Bayesian Entrepreneurship*, ed. by A. Agrawal, A. Camuffo, A. Gambardella, J. S. Gans, E. L. Scott, and S. Stern, Cambridge, MA: MIT Press, 1–54.

Agrawal, A., J. S. Gans, and S. Stern (2021): "Enabling Entrepreneurial Choice," *Management Science*, 67, 5510–5524.

Alvarez, S. A. and J. B. Barney (2007): "Discovery and Creation: Alternative Theories of Entrepreneurial Action," *Strategic Entrepreneurship Journal*, 1, 11–26.

Bailey, E., D. Fehder, E. Floyd, Y. Hochberg, and D. J. Lee (2026): "Learning to Quit? A Multi-Year, Multi-Site Field Experiment with Innovation-Driven Entrepreneurs," Working Paper 34755, National Bureau of Economic Research.

Blank, S. (2013): "Why the Lean Start-Up Changes Everything," *Harvard Business Review*, 91, 63–72.

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

Dixit, A. K. and R. S. Pindyck (1994): *Investment Under Uncertainty*, Princeton, NJ: Princeton University Press.

Folta, T. B. (1998): "Governance and Uncertainty: The Tradeoff Between Administrative Control and Commitment," *Strategic Management Journal*, 19, 1007–1028.

Gambardella, A., L. Gius, and S. Stern (2026): "Homo Entrepreneuricus," in *Bayesian Entrepreneurship*, ed. by A. Agrawal, A. Camuffo, A. Gambardella, J. S. Gans, E. L. Scott, and S. Stern, MIT Press, 55–76.

Gans, J. S. (2026): "Knightian Uncertainty and Bayesian Entrepreneurship," in *Bayesian Entrepreneurship*, ed. by A. Agrawal, A. Camuffo, A. Gambardella, J. S. Gans, E. L. Scott, and S. Stern, Cambridge, MA: MIT Press, forthcoming.

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

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

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

Knight, F. H. (1921): *Risk, Uncertainty, and Profit*, Boston: Hart, Schaffner and Marx.

Kremer, M. (1993): "The O-Ring Theory of Economic Development," *The Quarterly Journal of Economics*, 108, 551–575.

Manso, G. (2011): "Motivating Innovation," *Journal of Finance*, 66, 1823–1860.

Martínez, A. G. (2016): *Chaos Monkeys: Obscene Fortune and Random Failure in Silicon Valley*, New York: Harper Business.

McGrath, R. G. (1999): "Falling Forward: Real Options Reasoning and Entrepreneurial Failure," *Academy of Management Review*, 24, 13–30.

McMullen, J. S. and D. A. Shepherd (2006): "Entrepreneurial Action and the Role of Uncertainty in the Theory of the Entrepreneur," *Academy of Management Review*, 31, 132–152.

Ries, E. (2011): *The Lean Startup: How Today's Entrepreneurs Use Continuous Innovation to Create Radically Successful Businesses*, New York: Crown Business.

Shane, S. and S. Venkataraman (2000): "The Promise of Entrepreneurship as a Field of Research," *Academy of Management Review*, 25, 217–226.

Stern, S. and S. Zyontz (2019): "Jane Technologies," MIT Sloan Case 19-183, mIT Sloan School of Management.

Thomke, S. H. (2003): *Experimentation Matters: Unlocking the Potential of New Technologies for Innovation*, Boston: Harvard Business School Press.