Outside Options and Entrepreneurial Experimentation

Nathan Seegert[*] Robert Wuebker[†] Mu-Jeung Yang[‡]

Bayesian entrepreneurship explains entrepreneurial action, such as committing to or pivoting away from a given venture, as choice under uncertainty. Prior work in this literature treats outside options as common across entrepreneurs. This paper relaxes that assumption, introducing entrepreneur-specific outside options. Experiments then reveal not whether an idea is good, but whether it is good enough relative to what the entrepreneur could do instead. Outside options reshape the value, contingency, and design of entrepreneurial experiments. For example, more optimistic entrepreneurs may require stronger evidence before committing, not less; and outside options also determine which experiments entrepreneurs build. More broadly, phenomena that prior work interprets as excessive persistence, escalation of commitment, premature scaling, or pivoting may reflect entrepreneur-specific alternatives rather than failure to learn.

Keywords: Bayesian Entrepreneurship, Outside Option, Opportunity Cost, Identification, Entrepreneurship

JEL Codes: L20; L26; L29; D21; D23

[*]Department of Finance, D'Amore-McKim School of Business, Northeastern University

[†]Department of Entrepreneurship and Strategy, David Eccles School of Business, University of Utah

[‡]Department of Strategy, Entrepreneurship and Operations, Leeds School of Business, University of Colorado; Yang gratefully acknowledges financial support provided by the Leeds School of Business.

If I cannot have both, I will forsake fish and select bear's paw. - Mencius

1 Introduction

Although experimentation has long been central to entrepreneurship research and practice, recent work has increasingly formalized entrepreneurial experimentation as a problem of Bayesian learning under uncertainty (Agrawal, Camuffo, Gambardella, Gans, Scott, and Stern, 2026; Agrawal, Gans, and Stern, 2021; Gans, Stern, and Wu, 2019). In Bayesian entrepreneurship, entrepreneurs hold subjective priors about uncertain business ideas, generate evidence through experiments, update beliefs, and decide whether to commit to the focal idea or pivot away from it (Agrawal et al., 2026; Gambardella, Gius, and Stern, 2026; Gans, 2026). A growing literature has used this framework to analyze how prior beliefs, experiment costs, signal informativeness, and action thresholds shape whether entrepreneurs experiment, what evidence they seek, and whether they commit, continue, or pivot (Agrawal et al., 2026, 2021; Kim, 2026; Krieger, 2026). It has also connected Bayesian learning to "deep Bayesian" or "scientific" approaches to entrepreneurship, showing how hypotheses, causal theories, and disciplined updating can guide entrepreneurial choice under uncertainty (Camuffo, Cordova, Gambardella, and Spina, 2020; Felin, Gambardella, Novelli, and Zenger, 2024; Felin and Zenger, 2009; Novelli, 2026; Scott, 2026; Spina, 2026; Zellweger and Zenger, 2022). Taken together, this work has reframed entrepreneurial experimentation as a process of disciplined learning under uncertainty, rather than simply as the pursuit of promising ideas. At the same time, some users of Bayesian entrepreneurship, such as Arora (2026) recognize an important omission in the framework: "Clearly, there is more work to be done to sort out experimentation and information gathering, on the one hand, and costs and benefits, on the other. Heterogeneity in costs and benefits can lead to many of the same patterns as differences in priors."

We follow Arora's (2024) suggestion and relax a single assumption within the Bayesian entrepreneurship framework: That entrepreneurs share a common outside option, which is the value for the alternative they would pursue if they pivoted away from or abandoned the focal idea. This margin is not foreign to the Bayesian framework—-it is already present in the entrepreneur's choice problem, but is typically held fixed. We make it explicit and heterogeneous, since this allows us to better capture the incentives of very talented and capable high-growth entrepreneurs, who often have attractive outside options, such as lucrative careers or alternative ventures. Jeff Bezos left a senior role at D.E. Shaw to build Amazon, giving up an unusually attractive career. Reed Hastings had attractive alternatives in software before Netflix, and Peter Thiel had a derivatives-trading career at Credit Suisse before entrepreneurship. Additionally, serial entrepreneurs (Gompers, Kovner, Lerner, and Scharfstein, 2010) often choose among rival venture ideas, while experienced entrepreneurs may also have acquisition, employment, advisory, or investment opportunities that first-time or necessity entrepreneurs do not. The argument in this paper is that entrepreneur-specific outside options are not an arbitrary additional source of heterogeneity, but the opportunity-cost side of the same entrepreneurial choice problem that the current Bayesian entrepreneurship literature analyzes through beliefs, experimentation, and updating.

Using a formal model, we show that an entrepreneur's outside option shapes experimentation along three dimensions: its value, its contingency, and its design. It shapes value, meaning what experimentation is worth, because evidence pays off only when it can sharpen the comparison between the venture and the alternative. It shapes contingency, meaning when experimentation happens at all, because the same prior can lead one entrepreneur to commit, another to test, and a third to walk, according to the alternative each faces. And it shapes design, meaning which experiment gets built, because the alternative fixes which mistake is costly, and so whether the entrepreneur seeks credible bad news or credible good news, and how much rigor to buy. Our unifying idea is that experiments are judged not against beliefs, but against the opportunity cost of acting on them. Our model produces several new insights: We show that the presence of entrepreneur-specific outside options can overturn familiar theoretical results in the Bayesian entrepreneurship tradition. Prior work implies that entrepreneurs with stronger priors about an idea should require less evidence before committing, while entrepreneurs with weaker priors should require more demanding evidence before continuing (Agrawal et al., 2026; Gans, 2023). When outside options rise with priors, this prediction can reverse—entrepreneurs with the strongest priors may also have the strongest alternatives, and may therefore require especially strong evidence before continuing to commit. The same logic also has implications for how entrepreneurial behavior should be interpreted empirically. For example, persistence need not imply overoptimism, pivoting need not imply pessimism, and exit need not imply that the idea was bad. Each action may instead reflect the value of the entrepreneur's outside option. More generally, decisions that prior work might interpret as overconfidence (Van den Steen, 2011), bias (Åstebro, Fossen, and Gutierrez, 2026), or a lack of training (Camuffo et al., 2020) may be rational once the entrepreneur's outside option is taken into account.

This paper proceeds as follows. First, the paper reviews research on entrepreneurial experimentation and shows that, although outside options are often acknowledged in this literature, their value is typically treated as common across entrepreneurs. This review also explains why this simplification is not merely a change in notation, because it removes the opportunity-cost side of the choice problem. The paper then introduces a formal model in which entrepreneurs differ in the value of the alternative they would pursue if they pivot away from or abandon the focal idea. Analysis of the model shows how outside options change the value of experimentation, the decision to experiment, and the design of experiments. The paper concludes by discussing implications for Bayesian entrepreneurship, the broader literature on entrepreneurial experimentation and decision-making, and practice.

Related Literature

This paper is part of the Bayesian entrepreneurshiup literature (Agrawal et al., 2026). In it, entrepreneurs hold subjective priors, generate evidence, update beliefs, and then decide whether to commit to a focal idea or pivot away from it (Gans et al., 2019). Prior work shows that these decisions depend on the entrepreneur's prior beliefs, the cost and informativeness of experiments, and the evidence required to justify action under uncertainty (Agrawal et al., 2026, 2021; Gans, 2023; Gans et al., 2019). Additionally, the "deep Bayesian learning" view of Agrawal et al. (2026); Camuffo, Gambardella, and Pignataro (2024b) widens the lens from the strategies a founder might pursue to the *theories of value* that justify them. Related work on the scientific approach extends this logic by emphasizing that experiments are useful when they are tied to hypotheses, causal theories, and theories of value that give evidence meaning (Felin et al., 2024; Novelli, 2026; Spina, 2026). A central implication of this literature is that experimental strategy is contingent. Prior beliefs, experiment costs, signal informativeness, and the value of acting before uncertainty is fully resolved shape not only whether entrepreneurs experiment, but also what evidence they seek and how demanding that evidence must be (Agrawal et al., 2026; Contigiani and Levinthal, 2019; Gans et al., 2019). Subsequent studies have extended this logic to the structure of entrepreneurial search, the evidence required to justify action under uncertainty, and the sequence of experiments and pivots through which ideas are evaluated (Agarwal, Bacco, Camuffo, Coali, Gambardella, Msangi, Sonka, Temu, Waized, and Wormald, 2023; Chavda, Gans, and Stern, 2024; Chen, Elfenbein, Posen, and Wang, 2022; Valentine, Novelli, and Agarwal, 2024).

However, the traditional Bayesian entrepreneurship framework outlined in Agrawal et al. (2026) holds fixed what we consider an important part of the entrepreneur's decision problem: the value of the alternative available if the entrepreneur pivots away from or abandons the focal idea. Prior work often acknowledges that entrepreneurs have outside options, including alternative projects, employment opportunities, or other uses of their time and resources. But these alternatives are usually normalized to zero (Agrawal et al., 2026) or treated as common across entrepreneurs, allowing the analysis to focus on beliefs about the focal idea and the value of information generated by experiments (Agrawal et al., 2026; Arora, 2026; Gans et al., 2019)[1] This simplification may matter because entrepreneurs differ not only in what they believe about the focal idea, but also in what they could do instead. Others have much weaker alternatives. The same signal about the same idea may therefore justify continued investment for one entrepreneur and pivoting for another, and may also change whether the entrepreneur wants evidence that can stop commitment or evidence that can justify continuing. Explicitly analyzing outside options is also important for the normative advice the Bayesian entrepreneurship framework can offer, based on contingencies the theory incorporates. The point is not that any unobserved difference among entrepreneurs are consequential. Instead, personalized advice about experimentation requires a theoretical analysis of the entrepreneur's alternative as well as a theory of the focal idea (Snyder, Wuebker, and Zenger, 2024).

Close to our concern with the value of an experiment and the action it warrants is the pragmatist view of Zellweger and Zenger (2023, 2022). There the founder produces value through a sequence of causally inferential actions. She forms a belief about the theory of value creation; she tests it; and she responds to the feedback received. We take the same three-step sequence but treat its methodic doubts (about product-market fit, feedback validity and total error) as questions the founder can resolve by explicit Bayesian learning and optimization rather than heuristics as adovated by Zellweger and Zenger (2022). For example, we read Zellweger and Zenger (2022)'s the third doubt, that she might over- or under-react to feedback, as a question of how far to revise her belief once a result is in, and

[1]Some recent work relaxes related but distinct assumptions. Barney, Bigelow, Shelef, and Wuebker (2024), for example, examine settings in which entrepreneurs possess multiple business ideas and experimenting on one idea generates information about the value of another, "entangled" idea. Posen, Shelef, and Wuebker (2024) examine settings in which actors are uncertain about the value of assets they have already committed to an idea, and may therefore experiment to learn about those assets rather than about the focal offering, idea, or theory. The focus of this paper is different—it examines outside options whose value may differ across entrepreneurs if they pivot away from or abandon the focal idea, even when those alternatives are not informationally entangled with the focal idea and even when the entrepreneur is not learning about the residual value of assets committed to the venture.

the experimental design margin of rigor as a way to address it.

2 Experimentation Under Heterogeneous Outside Options

2.1 Model Setup

Bayesian entrepreneurship formalizes entrepreneurial action as a choice problem under uncertainty. An entrepreneur begins with a subjective prior about the quality of a venture, may run an experiment that generates an informative signal, updates beliefs by Bayes' rule, and then chooses whether to commit to the venture or pivot away from it. In the canonical setup, the payoff from pivoting or abandoning is treated as common across entrepreneurs and normalized to zero. Once that payoff is fixed, and holding the experimental technology constant, the entrepreneur's prior becomes the central state variable. It determines the value of information, whether experimentation is worthwhile, and what kind of experiment the entrepreneur should run (Agrawal et al., 2026; Gans, 2022).

This paper relaxes that normalization by introducing entrepreneur-specific outside options. In the model, the outside option is not a background payoff. It is part of the entrepreneur's action threshold. As a result, experiments do not simply provide information about the venture. They reveal whether the posterior belief generated by the experiment is strong enough to justify committing to the venture rather than taking the entrepreneur's alternative. This section analyzes how entrepreneur-specific outside options change the value of experimentation and the decision to experiment. Later sections extend the model to experimental design, including how precise the experiment should be and whether it should be designed to avoid continuing with a bad idea or abandoning a good one.

An entrepreneur begins with a prior belief, θ_S, that the venture is good. Given this prior, the entrepreneur decides whether to experiment. If the entrepreneur experiments, the experiment generates a signal about the venture's quality; the entrepreneur updates beliefs using Bayes' rule; and then chooses whether to commit resources to the venture or pivot away from it. Committing to a good venture yields the upside G, while committing to a bad venture yields no payoff. Pivoting away yields the outside option ω. In this paper, the outside option is entrepreneur-specific. It includes the expected payoff and option value of alternative business ideas, employment opportunities, or other uses of the entrepreneur's time and resources. It therefore captures the opportunity cost of continuing to commit resources to the venture.

Figure 1 summarizes the model used in this section. Panel A gives the timing of the entrepreneur's problem. The entrepreneur begins with a prior, decides whether to experiment, observes a signal if an experiment is run, and then chooses whether to commit resources to the venture or pivot away from it. Panel B represents the experiment as a noisy positive or negative signal, with true positive rate q_G and true negative rate q_B. Panel C translates beliefs into action through the threshold ω/G. The entrepreneur commits to the venture when the operative belief exceeds this threshold and pivots otherwise.

Each result below reads a different implication from this setup. The value of experimentation and the decision to experiment depend on whether the experiment can move the entrepreneur's belief across the action threshold. The accuracy of the experiment depends on how often the signal gives the correct verdict, as captured by q_G and q_B. The focus of the experiment depends on whether the experiment is better at identifying good ventures or bad ventures. Panel C illustrates the decision rule for a particular entrepreneur. In general, however, the threshold is entrepreneur-specific because it depends on the entrepreneur's outside option, ω, relative to the value of success, G. The figures below therefore add the outside option as a second dimension. Adding this second dimension shows why the same prior can sit above one entrepreneur's threshold and below another's.

(Expected) Decision Mistakes

The entrepreneur's decision can be organized around two potential regrets. The first is a *costly mistake*. The entrepreneur continues investing resources in a venture that turns out to

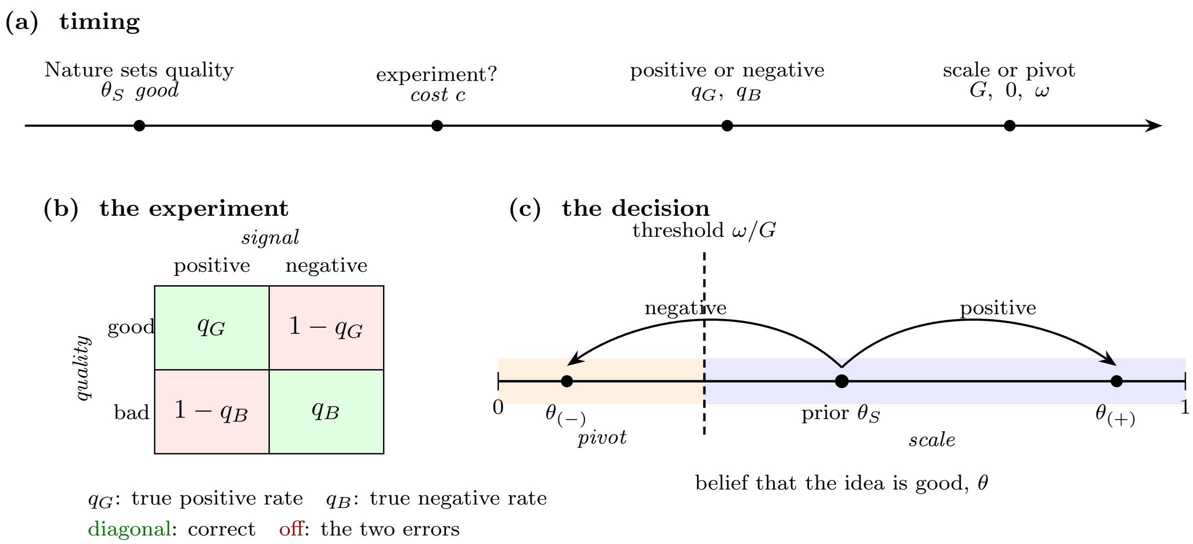

Figure 1. The model at a glance. Panel A shows timing. Nature draws the venture's quality, good with probability θ_S, after which the entrepreneur may run one experiment at cost c, observe a positive or negative result, and then commit or pivot for realized payoffs G, 0, or ω. Panel B shows the experiment as a positive or negative signal with true positive rate $q_G = \Pr(\text{positive} \mid \text{good})$ and true negative rate $q_B = \Pr(\text{negative} \mid \text{bad})$. The green diagonal is a correct verdict, and the off-diagonal cells are false positive and false negative errors. Panel C shows the decision rule. The entrepreneur commits when belief exceeds the threshold ω/G and pivots below it. A positive result raises the prior θ_S to the posterior $\theta_{(+)}$, and a negative result lowers it to $\theta_{(-)}$.

be bad. From the entrepreneur's perspective, this occurs with probability $1-\theta_S$ and costs the outside option, ω, that could have been preserved by pivoting away. The second is a *missed opportunity.* The entrepreneur pivots away from a venture that turns out to be good. From the entrepreneur's perspective, this occurs with probability θ_S and costs the upside the entrepreneur could have captured, G, net of the outside option received, ω. These two expected regrets are therefore

$$v_{\mathrm{cm}} = (1-\theta_S)\,\omega, \qquad v_{\mathrm{mo}} = \theta_S\,(G-\omega). \tag{1}$$

Each term is an expected *regret.* The costly-mistake stake is what the entrepreneur expects to lose, relative to perfect foresight, by committing to a venture that should have been abandoned. The missed-opportunity stake is what the entrepreneur expects to lose by pivoting away from a venture that should have been built. These expected regrets are conceptually distinct from the expected value of an action. Pivoting yields ω, while committing yields expected payoff $\theta_S G$. The two regrets are the two sides of a generalized version of Gans et al. (2019)'s *Paradox of Entrepreneurship*[2], a tension between learning and commitment.[3]

An entrepreneur can resolve uncertainty by continuing to invest and learning information that reveals its quality. But this form of learning requires commitment. In a one-shot venture, the lesson arrives as the payoff *itself*, too late to guide action and only after the entrepreneur has given up the outside option. Pivoting preserves the outside option, but it may also forgo the upside of a venture the entrepreneur never learned was good. Prior work in Bayesian entrepreneurship emphasizes that experiments help entrepreneurs navigate this problem by generating noisy but decision-relevant signals before commitment is irreversible (Agrawal et al., 2026, 2021; Gans et al., 2019). The outside option is central to this tradeoff because it determines what the entrepreneur gives up by continuing to commit resources to the venture.

[2]The original Paradox of Entrepreneurship centers around learning about different going to market strategies for a business idea.

[3]The more general statement is that choosing well among the paths open to a venture requires knowledge only experimentation can supply, yet experimentation "inevitably results in (at least some) commitment that forecloses strategic options" (Scott, 2026, p. 251).

The difference $v_{\text{mo}} - v_{\text{cm}}$ is the *launch return*. It is the expected surplus from committing to the venture rather than taking the outside option,[4]

$$R \;=\; v_{\text{mo}} - v_{\text{cm}} \;=\; \theta_S G - \omega. \tag{2}$$

The sign of the launch return pins down the entrepreneur's default action without experimentation. When the launch return is positive ($R > 0$), the expected value of committing to the venture exceeds the outside option. The entrepreneur therefore commits by default. This occurs when the prior belief exceeds the action threshold ω/G (see Figure 1(C)). In this case, the entrepreneur faces a larger missed opportunity than costly mistake. When the launch return is negative ($R < 0$), the outside option exceeds the expected value of committing to the venture. The entrepreneur therefore pivots by default. This occurs when the prior belief is below the action threshold ω/G. In this case, the entrepreneur faces a larger costly mistake than missed opportunity. In both cases, the default action avoids the larger expected regret while leaving the entrepreneur exposed to the smaller one.

An experiment is valuable when it can move the entrepreneur away from that default action.

Experimentation

An experiment is described by two error rates, its *true positive rate*, $q_G = \Pr(\text{positive} \mid \text{good})$, and its *true negative rate*, $q_B = \Pr(\text{negative} \mid \text{bad})$.[5]

Beginning from the prior θ_S, the entrepreneur updates beliefs by Bayes' rule. A positive result raises the entrepreneur's belief to the posterior $\theta_{(+)}$, while a negative result lowers it

[4]Under the correspondence with the canonical model ($G \leftrightarrow V$, $\omega \leftrightarrow C$), the launch return corresponds to $R(\mu) = \mu V - C$ in Agrawal et al. (2026), evaluated at the prior, and we retain the symbol R.

[5]These correspond to the true positive rate $\lambda_1 = \Pr[\text{positive} \mid \text{good}]$ and true negative rate $\lambda_0 = \Pr[\text{negative} \mid \text{bad}]$ in Gans (2022) and Agrawal et al. (2026). The notation q_G, q_B keeps the good/bad conditioning visible.

to $\theta_{(-)}$,

$$\theta_{(+)} = \frac{\theta_S\, q_G}{\theta_S\, q_G + (1-\theta_S)(1-q_B)}, \qquad \theta_{(-)} = \frac{\theta_S\,(1-q_G)}{\theta_S\,(1-q_G) + (1-\theta_S)\, q_B}, \tag{3}$$

with $\theta_{(-)} < \theta_S < \theta_{(+)}$ whenever the experiment is informative ($q_G + q_B > 1$). The entrepreneur then commits resources to the venture if the operative belief, either the posterior after an experiment or the prior if no experiment is run, exceeds the action threshold ω/G, and pivots otherwise.

Two features of the experiment matter for the analysis. The first is its *accuracy.* Accuracy is the average probability that the experiment returns the correct verdict in each state, given by $\bar{q} = \frac{1}{2}(q_G + q_B) \in (\frac{1}{2}, 1)$. The second is its *focus.* Focus captures whether the experiment is better at identifying good ventures or bad ventures, given by the difference $q_G - q_B$. Section 3 and Appendix A allow the entrepreneur to design the experiment by choosing both its accuracy and its focus. For now, the analysis focuses on symmetric experiments, $q_G = q_B = \bar{q}$.

The *value of experimenting* is the expected gain from allowing the signal to change the entrepreneur's default action, net of the experiment cost c. This value differs depending on whether the entrepreneur would commit or pivot by default. Let V_c denote the value of experimentation for an entrepreneur whose default is to commit to the venture, and let V_p denote the value for an entrepreneur whose default is to pivot. These values are

$$V_c = q_B\, v_{\text{cm}} - (1-q_G)\, v_{\text{mo}} - c, \qquad V_p = q_G\, v_{\text{mo}} - (1-q_B)\, v_{\text{cm}} - c. \tag{4}$$

For an entrepreneur who would commit by default, the experiment creates value when it correctly identifies a bad venture and allows the entrepreneur to avoid a costly mistake, $q_B\, v_{\text{cm}}$. The experiment destroys value when it incorrectly produces a negative result for a good venture and leads the entrepreneur to abandon a missed opportunity, $(1-q_G)\, v_{\text{mo}}$. For an entrepreneur who would pivot by default, the experiment creates value when it correctly identifies a good venture and allows the entrepreneur to avoid a missed opportunity, $q_G\, v_{\mathrm{mo}}$. The experiment destroys value when it incorrectly produces a positive result for a bad venture and leads the entrepreneur into a costly mistake, $(1-q_B)\, v_{\mathrm{cm}}$. The four results below follow from these two value functions. Appendix A.1 derives the expressions term by term from Bayes' rule (eqs. (14)–(15)).

Outside Options, Heterogeneity and Co-movement with Priors

The outside option varies in two analytically distinct ways. First, it can vary in *level*, holding the entrepreneur's prior fixed. Second, it can vary *with beliefs* across entrepreneurs. The second form of variation is central to the analysis because it asks how experimental strategy changes when the entrepreneurs with stronger priors are also the entrepreneurs with stronger outside options. For each of the four questions below, the paper therefore states two results. The first result in each pair, numbered ".1", holds the prior fixed and varies the outside option. The second result, numbered ".2", traces behavior across a population of entrepreneurs whose outside options move with their beliefs. We summarize this relationship by the slope of the conditional mean,

$$k \;\equiv\; \frac{\mathrm{Cov}(\omega, \theta_S)}{\mathrm{Var}(\theta_S)}, \qquad \mathbb{E}[\omega \mid \theta_S] = a + k\,\theta_S.$$

We call k the *co-movement* of outside options and beliefs. The intercept a is the baseline outside option of an entrepreneur with prior $\theta_S = 0$. This formulation is meant to summarize how outside options and beliefs move together across entrepreneurs. It is not a claim that outside options are generated by a particular distribution, nor does the analysis require affiliation, monotone likelihood ratios, or stronger assumptions about the ordering of entrepreneurs. What matters is whether, and by how much, outside options rise with beliefs.

This co-movement has a natural interpretation in entrepreneurial settings. For high-growth entrepreneurs, the same ability, ambition, experience, or market validation that supports a strong belief in the venture may also create valuable alternatives, including other ventures, acquisition opportunities, senior operating roles, or other ways to redeploy time and resources. In this case, outside options may rise steeply with beliefs. For necessity entrepreneurs, by contrast, confidence in a particular venture may say little about the value of the alternative available if the entrepreneur pivots away. The high-growth case is especially important for the results below because it is where the co-movement between beliefs and outside options can become large enough to change standard Bayesian entrepreneurship predictions.

The same threshold logic drives the comparative statics that follow. Each object of interest, including the value of experimentation, the decision to experiment, accuracy, and focus, is a function $X(\theta_S, \omega)$ of the prior and the outside option. Holding ω fixed, a change in the prior moves the entrepreneur relative to the action threshold. Holding θ_S fixed, a change in the outside option moves the threshold itself. When outside options co-move with priors, the cross-sectional slope of any object is

$$\frac{dX}{d\theta_S} = \frac{\partial X}{\partial \theta_S} + k\,\frac{\partial X}{\partial \omega}$$

as in Lemma 2. The two terms often have opposite signs. Co-movement therefore determines whether belief and threshold effects offset or reinforce each other. The ".2" propositions apply this logic to value, participation, accuracy, and focus.

2.2 The Value of Experimentation

In the standard Bayesian entrepreneurship model, the value of experimentation is single-peaked in the prior. Entrepreneurs near certainty, optimistic or pessimistic, have little to learn and do not experiment (Agrawal et al., 2026). This result follows when the outside option is the same for all actors. Given a common outside option, all entrepreneurs face the same action threshold, so variation in the value of experimentation is determined from the prior alone. Once the outside option is entrepreneur-specific, the peak moves. Experimentation is most valuable at the entrepreneur's *own* decision boundary, where the launch return is zero and the entrepreneur is indifferent between committing by default and pivoting by default. A higher outside option raises that boundary, so the entrepreneur who is indifferent before experimentation must exhibit a higher prior. An experiment is valuable only if it can change the entrepreneur's action. Let V_c denote the value of an experiment for an entrepreneur who commits by default, and let V_p denote the value of an experiment for an entrepreneur who would otherwise pivot. In both cases, the value is the expected improvement from sometimes changing that default action, net of the experiment cost c.

Proposition 1.1 (The value of experimentation peaks at the action threshold)**.** *For any fixed outside option ω, the net value of a symmetric experiment is single-peaked in the prior and reaches its maximum at the entrepreneur's action threshold,*

$$\theta_S^* = \frac{\omega}{G},$$

where the launch return is zero, $R = \theta_S G - \omega = 0$. Experimentation is therefore most valuable when the entrepreneur is indifferent between committing by default and pivoting by default. The standard result that experimentation peaks at moderate beliefs is the special case $\omega = G/2$. More generally, the entrepreneur's peak belief rises with the outside option, $\partial\theta_S^*/\partial\omega = 1/G$. A stronger outside option also raises the value of experimentation for a commit-default entrepreneur and lowers it for a pivot-default entrepreneur,

$$\frac{\partial V_c}{\partial \omega} > 0, \qquad \frac{\partial V_p}{\partial \omega} < 0.$$

The two comparative statics describe the same trade-off from opposite sides of the action threshold. A commit-default entrepreneur experiments to avoid committing to a bad idea. A negative result is valuable precisely because it moves the entrepreneur away from the default action and preserves the outside option. The stronger the outside option, the more valuable that correction becomes. A pivot-default entrepreneur experiments to avoid abandoning a good idea. A positive result is valuable precisely because it moves the entrepreneur toward committing, but that correction requires giving up the outside option. The stronger the outside option, the less valuable that correction becomes.

This is why the outside option does not simply raise or lower the value of experimentation. It changes which mistake the experiment is valuable for. A stronger outside option raises the costly-mistake stake, $v_{\text{cm}} = (1 - \theta_S)\omega$ and lowers the missed-opportunity stake, $v_{\text{mo}} = \theta_S(G - \omega)$. Since the launch return is the difference between these stakes, $R = v_{\text{mo}} - v_{\text{cm}} = \theta_S G - \omega$, a higher outside option lowers the launch return one-for-one. It moves commit-default entrepreneurs toward their action threshold and moves pivot-default entrepreneurs further away from it.

Figure 2 plots the value of experimentation over the prior and outside-option space. Each entrepreneur is a point (θ_S, ω). The feasible region is the strip $0 \leq \omega \leq G$. The outside option cannot be negative, and an outside option above G would dominate even a sure-good venture. The lens is the set of entrepreneurs for whom a symmetric experiment has positive net value. Its bold rim is the break-even contour $V = 0$. The dashed diagonal is the action threshold, $\theta_S G = \omega$. Below the diagonal, entrepreneurs commit by default, and a stronger outside option raises the value of experimentation because a negative signal protects a more valuable alternative. Above the diagonal, entrepreneurs pivot by default, and a stronger outside option lowers the value of experimentation because a positive signal must justify giving up that alternative. The diagonal itself is the ridge of the value surface. Experimentation is most valuable where the launch return is zero and the two expected regrets are equal.

At the boundary $\theta_S = \omega/G$, where V_c and V_p coincide, the net value of the symmetric experiment reaches its ridge,

$$V\Big|_{\theta_S=\omega/G} = \frac{\omega(G-\omega)}{G}\,(2\bar{q}-1) - c, \tag{5}$$

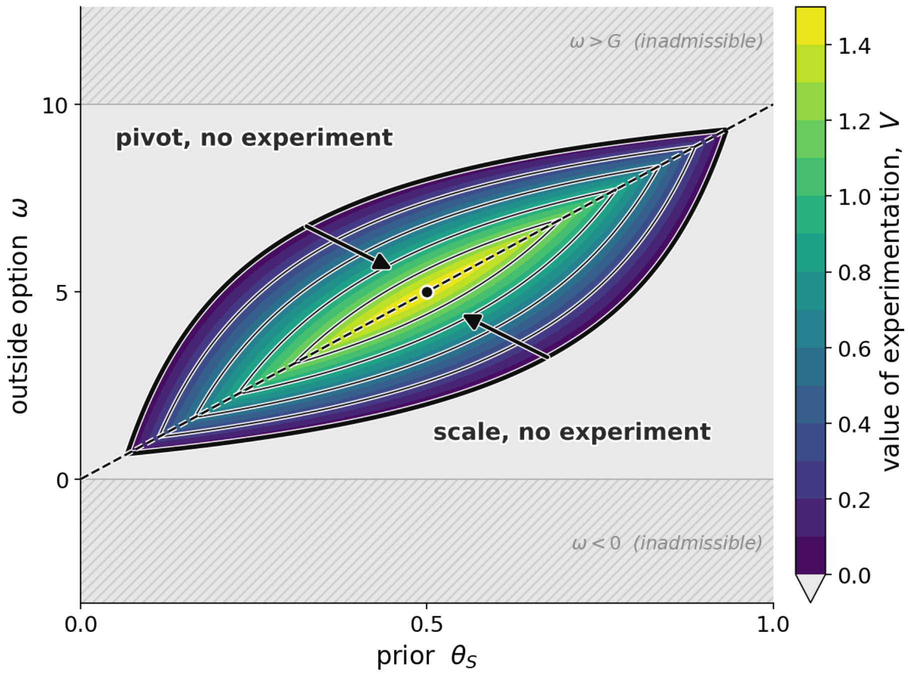

Figure 2. The value of experimentation across entrepreneurs. The shading is the net value of a symmetric experiment, V, measured relative to the entrepreneur's default action and net of the cost c, over the (θ_S, ω) plane. Brighter is higher, and the bold contour is the break-even line $V = 0$. Inside the lens, value rises to a ridge along the action threshold $\theta_S = \omega/G$ (dashed) and peaks at $(\frac{1}{2}, G/2)$, where the entrepreneur is most uncertain and the two expected regrets are equal. Outside the lens, in the deep-commit and deep-pivot corners (grey), an experiment cannot pay for itself. Parameters are $G = 10$, $\bar{q} = 0.9$, and $c = 0.5$.

the bright diagonal of Figure 2. This ridge is highest at $\omega = G/2$, where the entrepreneur is also at prior $\theta_S = 1/2$. The experiment pays where the ridge value is positive, so along the diagonal the lens runs between the two priors solving $G\,\theta_S(1-\theta_S)(2\bar{q}-1) = c$, where the rim meets the prior axis (Lemma 3, eq. (23)).

Proposition 1.2 (Co-movement moves the experimentation peak)**.** *Let outside options co-move with priors along* $\mathbb{E}[\omega \mid \theta_S] = a + k\theta_S$ *with* $0 < k < G$*. Along this cross-section the value of experimentation is concave in the prior and single-peaked. The peak lies at the belief where the population reaches the action threshold,*

$$\theta_S^* = \frac{a}{G-k},$$

whenever θ_S^* falls between the vertices of the two regret branches, and the peak belief is then strictly increasing in the co-movement k. Under stronger co-movement the peak detaches from the threshold but remains interior. In the limiting case $k = G$ and $a = 0$, the outside option rises one-for-one with perceived venture value, so every entrepreneur lies on the action threshold, $\theta_S G = \omega$. In that case, willingness to pay for a symmetric experiment is positive at every interior prior,

$$G\,\theta_S(1 - \theta_S)(2\bar{q} - 1),$$

and the value of experimentation is no longer concentrated among entrepreneurs with moderate beliefs about the value of their business idea.

Along the conditional mean $\mathbb{E}[\omega \mid \theta_S] = a + k\theta_S$, the launch return changes linearly with the entrepreneur's prior,

$$R(\theta_S) = (G - k)\,\theta_S - a, \qquad \frac{dR}{d\theta_S} = G - k, \qquad \theta_S^* = \frac{a}{G - k}. \tag{6}$$

The slope $dR/d\theta_S = G - k$ is the key cross-sectional object. It measures how much the entrepreneur's perceived venture value rises relative to the outside option as the prior rises. When $k = 0$, outside options do not co-move with priors, so the action threshold is fixed and the standard inverted-V appears, peaking at a/G. When $0 < k < G$, higher priors are partly offset by stronger outside options. More optimistic entrepreneurs have higher perceived venture value, $\theta_S G$, but they also face a higher threshold, ω. They therefore remain closer to their own decision boundaries than the common-outside-option model in prior work would imply, and for moderate co-movement the demand peak tracks the population threshold θ_S^*, moving outward as k rises. For strong co-movement the peak detaches from the threshold: demand along the tilted slice is the lower envelope of two concave regret branches, and once θ_S^* passes a branch vertex the peak stays at that vertex (Appendix A.3 gives the formulas). What strong co-movement does is not push the peak out of the belief range but spread positive demand across it.

This is why the moderate-believer result is an artifact of holding the outside option fixed. Each entrepreneur's willingness to experiment peaks at the entrepreneur's own action threshold, $\theta_S G = \omega$, where the launch return is zero and the two expected regrets are equal. A population with a common outside option is a horizontal slice through Figure 2. It crosses the ridge at one belief, generating the familiar inverted-V (Agrawal et al., 2026). A population whose outside option rises with belief is a tilted slice. As the slice approaches the ridge, more entrepreneurs sit near their own action thresholds. In the limiting case $k = G$ and $a = 0$, the outside option tracks perceived venture value one-for-one, every entrepreneur lies on the ridge, and none is too convinced, in either direction, to experiment. By treating the outside option as common, the standard model also makes the action threshold common. This paper lets that threshold vary across entrepreneurs.

2.3 When to Experiment

The standard model already delivers a non-monotonicity in the prior. Entrepreneurs with very low priors pivot regardless of the signal, and entrepreneurs with very high priors commit regardless, so experimentation is valuable only for an intermediate band of beliefs (Agrawal et al., 2026). Once outside options vary, the same logic applies to the action threshold rather than to the prior alone. Holding the prior fixed, a stronger outside option can first make experimentation more valuable by moving a commit-default entrepreneur toward the threshold, and then make experimentation less valuable by moving a pivot-default entrepreneur farther from it. The result is a non-monotonicity in the outside option, running through the same two-regret trade-off (Figure 3).

Proposition 2.1 (The outside option induces and suppresses experimentation)**.** *The outside option* induces *experimentation among commit-default entrepreneurs and* suppresses *it among pivot-default entrepreneurs, so its effect on the aggregate experimentation rate is non-monotone. Moreover, no entrepreneur experiments unless the experiment is accurate enough* *to move the posterior across the action threshold. For a symmetric experiment, this requires*

$$\frac{\bar{q}}{1-\bar{q}} > \frac{\theta_S}{1-\theta_S}\,\frac{G-\omega}{\omega}, \tag{7}$$

a minimum-accuracy floor set by the entrepreneur's prior odds and the stakes ratio $(G-\omega)/\omega$*, not by the experiment's cost.*

The accuracy floor separates usefulness from cost. An entrepreneur runs an experiment only if at least one possible result could change the action. An experiment that is too imprecise to move the posterior across the threshold is worthless, even if it is cheap. The floor (7) (derived in Appendix A.4) is therefore the threshold of usefulness. It depends on the prior odds and on the stakes ratio $(G-\omega)/\omega$, which is to say, on the outside option, but not on c. The cost determines which entrepreneurs pay for experiments that could change their action. The floor determines whose action could be changed at all. This is the margin that a literature setting $c = 0$ to study experiment design assumes away.[6]

Proposition 2.2 (Positive co-movement maximizes experimentation)**.** *Across entrepreneurs whose priors span an interval of width* $\Delta\theta$*, the dispersion of launch returns is* $|G-k|\,\Delta\theta$*. This dispersion shrinks as* k *approaches* G *from below. Positive co-movement therefore concentrates entrepreneurs near their own action thresholds, where information is most likely to change the action. Aggregate experimentation is maximized as* $k \uparrow G$*, when variation in perceived venture value is matched by variation in outside options.*

One might expect an entrepreneur with both a strong prior and an attractive outside option to skip experimentation entirely. The entrepreneur seems to have enough confidence to commit and enough outside option value to pivot, so paying to learn may seem unnecessary. The model shows the opposite. This entrepreneur can be especially likely to find the

[6]Setting $c = 0$ does more than simplify. It removes the participation question entirely, so that every entrepreneur experiments and the design of the experiment can be read as a direct signal of the prior. Reintroducing a positive cost, as we do, repopulates the non-experimenting entrepreneurs that the $c = 0$ device suppresses, and makes observed experiment design a selected, rather than universal, signal.

experiment worth running because experimentation is most valuable near the action threshold. Positive co-movement keeps high-prior entrepreneurs near that threshold by pairing high beliefs with high outside options. The population sorts toward an *indifference diagonal*, defined by the zero-launch-return locus $R \approx 0$. Along this diagonal, belief and the outside option pull in opposite directions, so the verdict of an experiment can decide the action. The entrepreneurs who skip experimentation are those away from the threshold. High-prior entrepreneurs with weak outside options commit without experimenting, and low-prior entrepreneurs with strong outside options pivot without experimenting.

The Stanford MBA. Consider an entrepreneur with prior $\theta_S = 0.6$ in a venture whose upside is $G = 10$ (million), weighing a standing job offer from a technology company worth $\omega = 8$. The perceived venture value is $\theta_S G = 6$, which falls short of the outside option. The entrepreneur is therefore pivot-default. Absent new evidence, the offer is better than committing to the venture. In the standard reading, this entrepreneur is optimistic because the prior exceeds one-half. But optimism is being evaluated against an unusually strong outside option. The entrepreneur may still experiment because a sufficiently positive signal could move the posterior across the action threshold. If the outside option falls below 6, nothing about the prior changes, but the default action reverses. The entrepreneur becomes commit-default. Belief alone does not place the entrepreneur. The comparison between belief and the outside option does.

Figure 3 follows a single entrepreneur through the lens as the outside option changes. Entrepreneurs below the lens commit without experimenting, and entrepreneurs above it pivot without experimenting. The arrow tracks a confident entrepreneur ($\theta_S = 0.88$) as the outside option rises. At low values of ω, the entrepreneur is deep commit-default and commits regardless of the signal. As ω rises, the entrepreneur enters the lens and experiments. At still higher values of ω, the entrepreneur pivots without experimenting because the venture no longer clears the action threshold. One prior supports three different postures toward

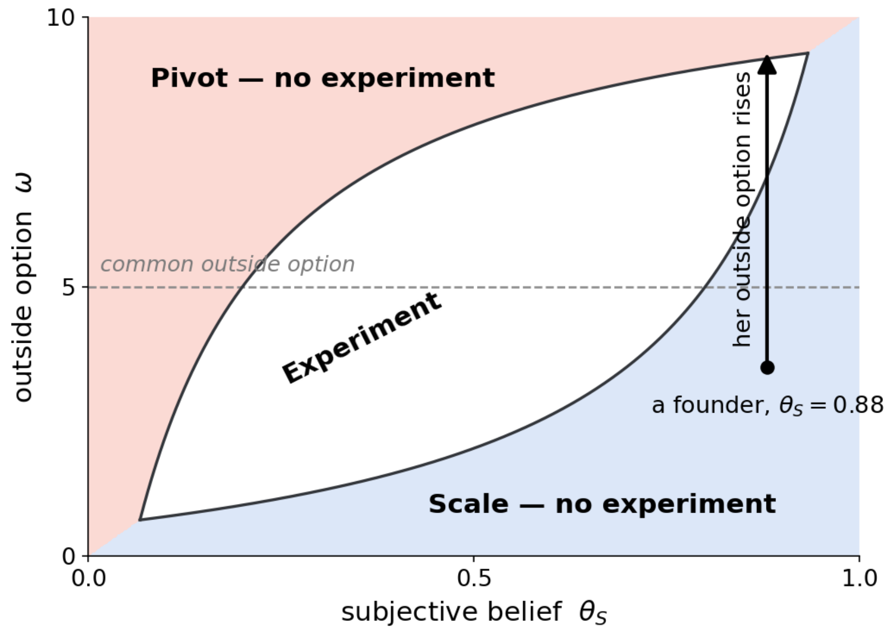

Figure 3. When to experiment as the outside option rises. The white lens is the region in which an experiment is worth running. Below the lens, the entrepreneur commits without experimenting. Above it, the entrepreneur pivots without experimenting. The dashed line marks a common outside option, $\omega = G/2$. The arrow follows a confident entrepreneur ($\theta_S = 0.88$) as the outside option rises. The entrepreneur moves from committing without experimenting, into the region where experimentation pays, and then into pivoting without experimenting. The posture toward experimentation is set by the outside option, not by belief alone. Parameters are $G = 10$, $\bar{q} = 0.9$, and $c = 0.5$.

experimentation. What moves the entrepreneur among them is the outside option. Proposition 2.2 reads the same geometry across entrepreneurs rather than along one entrepreneur. Positive co-movement places confident entrepreneurs with strong outside options near the indifference diagonal, where information can change the action.

2.4 How to Experiment: Accuracy

With the experiment held symmetric, the question is how much accuracy $\bar{q}$ the entrepreneur will pay for. The standard model of Agrawal et al. (2026) says little about this margin because experimentation is costless. If experimentation is free, every entrepreneur prefers a more accurate experiment to a less accurate one. Once experimentation is costly and entrepreneurs differ in outside options, accuracy becomes a choice about stakes. A more accurate experiment is valuable in proportion to what rides on the decision it informs (Figure 4). With a symmetric experiment, $q_G = q_B = \bar{q}$, the two value functions collapse to a single expression that is linear in accuracy,

$$V(\bar{q}) = (v_{\text{mo}} + v_{\text{cm}})\,\bar{q} - \max\{v_{\text{mo}}, v_{\text{cm}}\} - c. \tag{8}$$

A perfect experiment, with $\bar{q} = 1$, is worth the smaller of the two expected regrets net of cost (Appendix A.5, eq. (26)). The marginal value in Proposition 3.1 is the slope of this expression.

Proposition 3.1 (The marginal value of accuracy is the total stake)**.** *For an entrepreneur running a symmetric experiment of accuracy $\bar{q} = \Pr(\textit{positive} \mid \textit{good}) = \Pr(\textit{negative} \mid \textit{bad})$, the value of experimentation is linear in $\bar{q}$. A marginal increase in accuracy is worth the total expected stake riding on the decision,*

$$\frac{\partial V}{\partial \bar{q}} \;=\; \theta_S(G-\omega) + (1-\theta_S)\omega \;=\; v_{\text{mo}} + v_{\text{cm}}. \tag{9}$$

This marginal value is the same for commit-default and pivot-default entrepreneurs. A more optimistic entrepreneur places more weight on the missed-opportunity stake, $v_{\text{mo}} = \theta_S(G - \omega)$. A more pessimistic entrepreneur places more weight on the costly-mistake stake, $v_{\text{cm}} = (1-\theta_S)\omega$. Holding the outside option fixed, $\partial^2 V/\partial\bar{q}\,\partial\theta_S = G - 2\omega$. Accuracy therefore becomes more valuable with optimism when $\omega < G/2$ and less valuable with optimism when $\omega > G/2$. In the Bayesian entrepreneurship benchmark with a common outside option below $G/2$, accuracy is more valuable to more optimistic entrepreneurs. With heterogeneous outside options, entrepreneurs whose outside option exceeds half the prize value accuracy more when they are more pessimistic, because the costly mistake is the dominant stake.

The result separates accuracy from the default action. An entrepreneur values accuracy in proportion to the total stake riding on the decision, not only the mistake the default action leaves exposed. A commit-default entrepreneur will commit absent new evidence. Yet this entrepreneur still values a sharper experiment for both possible errors. A sharper experiment can correctly identify the bad venture the entrepreneur would otherwise commit to, and it can also avoid a false negative that would push the entrepreneur away from a good venture. Accuracy is therefore priced by the sum of the two expected regrets, $v_{\mathrm{mo}} + v_{\mathrm{cm}}$. The launch return, $R = v_{\mathrm{mo}} - v_{\mathrm{cm}}$, determines which regret is larger and therefore which action is the default. Accuracy is different. It is valuable because it reduces both errors.

This distinction changes who should value accuracy most. The optimist is wagering on the upside, $G - \omega$. The pessimist is wagering on the outside option, ω. A more accurate experiment is worth more to whichever entrepreneur has the larger total stake. The zero-outside-option benchmark sits at the edge of this logic. When $\omega = 0$, the costly-mistake stake carries no weight and the marginal value of accuracy rises with optimism.[7]

Figure 4 reads Proposition 3.1 from the stacked bars. In each panel, the height of the stack is the marginal value of accuracy. The lower block is the costly-mistake stake, v_{cm}. The upper block is the missed-opportunity stake, v_{mo}. The top edge is their sum, which is what one more increment of accuracy is worth. As belief rises, the costly-mistake stake shrinks and the missed-opportunity stake grows. Whether the total rises or falls depends on whether the upside stake grows faster than the outside-option stake shrinks. With a modest outside option, $\omega < G/2$, the total rises with belief and accuracy is worth more to more optimistic entrepreneurs. With a strong outside option, $\omega > G/2$, the total falls with belief and accuracy is worth more to more pessimistic entrepreneurs. The action threshold marks the point where the two stakes are equal, and the launch return changes sign. The difference between the stakes flips at that point. The sum passes through it smoothly.

Proposition 3.2 (Co-movement makes the demand for accuracy hump-shaped)**.** *When the*

[7] At $\omega = 0$, the value of experimenting is negative for every entrepreneur. Pivoting is worthless, so a negative result creates no gain from preserving the outside option, while a false negative forfeits the full prize. The ranking is therefore a statement about the marginal value of accuracy, $\partial V/\partial \bar{q} = \theta_S G$, not about who experiments. Any small common outside option keeps entrepreneurs in the experimentation region while preserving the ranking.

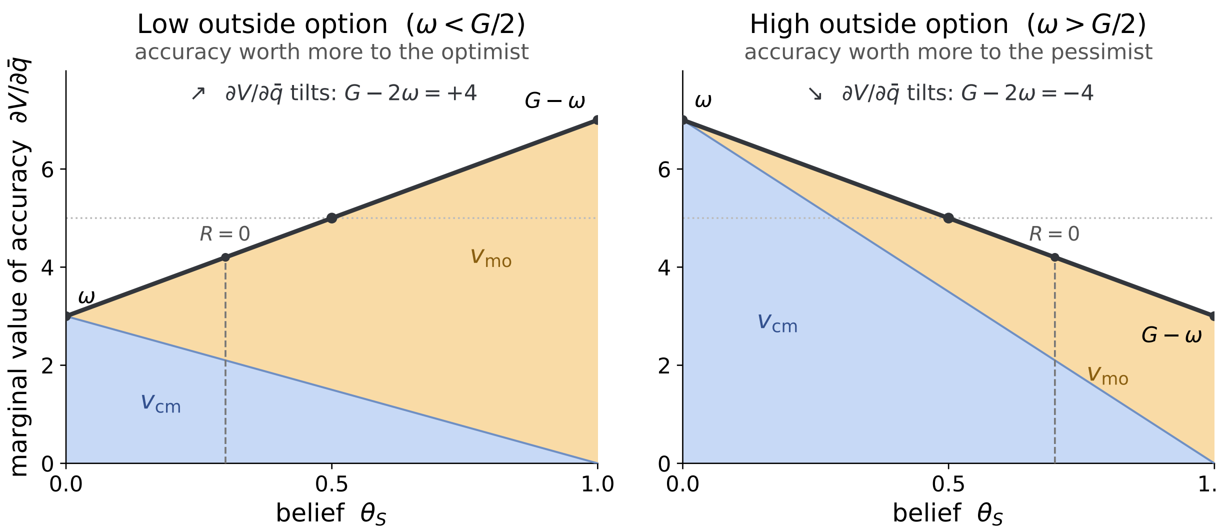

Figure 4. The marginal value of accuracy is the total stake. Each panel stacks the two expected regrets that ride on the experiment as belief θ_S varies, holding the outside option fixed. The costly mistake is $v_{\mathrm{cm}} = (1-\theta_S)\omega$ and the missed opportunity is $v_{\mathrm{mo}} = \theta_S(G-\omega)$. Their sum is the marginal value of a symmetric experiment, $\partial V/\partial\bar{q} = v_{\mathrm{mo}}+v_{\mathrm{cm}} = \omega+\theta_S(G-2\omega)$. In the left panel, a low outside option ($\omega < G/2$) makes the upside the larger wager, so the total rises with belief, and a sharper experiment is worth more to more optimistic entrepreneurs. In the right panel, a high outside option ($\omega > G/2$) reverses this pattern, so accuracy is worth more to more pessimistic entrepreneurs. The two stakes are equal at the action threshold $\theta_S = \omega/G$ (dashed), where the launch return $R = v_{\mathrm{mo}} - v_{\mathrm{cm}}$ changes sign. Parameters are $G = 10$, with $\omega = 3$ in the left panel and $\omega = 7$ in the right panel.

outside option co-moves strongly enough with the prior, the marginal value of experimental accuracy is hump-shaped in optimism. It can also decrease across the optimistic half of the entrepreneur population, so that the most optimistic entrepreneurs have the lowest marginal value of accuracy. This pattern cannot be generated by a common outside option. Along $\mathbb{E}[\omega \mid \theta_S] = a + k\theta_S$*, the marginal value* $v_{\text{mo}} + v_{\text{cm}}$ *is concave in* θ_S *with curvature* $-4k$*. It peaks at the interior prior* $\theta_S^Q = (G - 2a + k)/(4k)$*. Once* $k > G - 2a$*, the peak lies below one half and the most optimistic entrepreneurs value accuracy strictly less than the most pessimistic entrepreneurs* (eqs. (29)–(31)).

The mechanism is the same action-threshold logic seen above, now applied to the total stake. Once the outside option rises with belief, each expected regret becomes concave in optimism,

$$v_{\text{mo}} = \theta_S(G - a - k\theta_S), \qquad v_{\text{cm}} = (1 - \theta_S)(a + k\theta_S).$$

For the missed-opportunity stake, a higher prior increases the probability that the venture is good. But the same higher prior is associated with a stronger outside option, which shrinks the upside from committing rather than pivoting. These two forces work against each other inside v_{mo}. Past a critical level of optimism, the shrinking upside dominates, and the missed-opportunity stake falls. The costly-mistake stake also remains small for highly optimistic entrepreneurs because the bad state receives little weight. The most confident entrepreneurs can therefore have little of either regret to insure. Their outside options make a misstep less costly, and their confidence leaves them only weakly exposed to the costly mistake. With little at stake, they demand less accuracy.

2.5 How to Experiment: Focus

The sharpest contrast is with the standard treatment of experimental *focus*. In Agrawal et al. (2026), an entrepreneur predisposed to commit runs a low-bar experiment, designed so that only credible bad news can stop commitment. An entrepreneur predisposed to pivot runs a high-bar experiment, designed so that only credible good news can justify commitment. Because predisposition is read from the prior against a common action threshold, focus maps cleanly onto belief. The optimist runs the low-bar experiment, the pessimist runs the high-bar experiment, and an observer can use the entrepreneur's choice of experiment to infer whether the entrepreneur is optimistic or pessimistic. The results below show why that inference is under-identified once outside options are entrepreneur-specific.

Proposition 4.1 (Focus is selected by the outside option)**.** *Experimental focus is selected by the entrepreneur's position relative to the action threshold, not by the prior alone. Two entrepreneurs with the same prior and the same upside but outside options on opposite sides of $\theta_S G$ build opposite experiments. A change in ω that crosses $\theta_S G$ reverses experimental focus with no change in belief. Because a positive cost makes experimentation worthwhile only for entrepreneurs near their action threshold, observed focus identifies whether an entrepreneur is near a commit or pivot margin, not whether the prior is high or low in isolation.*

Maximizing the value of experimentation over an accuracy frontier $q_G + q_B = Q$ yields a first-order condition equal to the launch return,

$$\left.\frac{\partial V_c}{\partial q_G}\right|_{q_B=Q-q_G} = \theta_S G - \omega = R, \qquad \left.\frac{\partial V_p}{\partial q_B}\right|_{q_G=Q-q_B} = \omega - \theta_S G = -R. \tag{10}$$

Focus is therefore selected by the sign of R. A commit-default entrepreneur, with $R > 0$, drives the true positive rate to its maximum and builds a low-bar experiment. A pivot-default entrepreneur, with $R < 0$, drives the true negative rate to its maximum and builds a high-bar experiment. The dividing line is the action threshold $\theta_S G = \omega$. It is a comparison of perceived venture value with the outside option, not a threshold in belief alone. The frontier and corner solutions are derived in Appendix A.6.

Two implications follow. First, the same prior can support opposite experiments. Under a weak outside option, the entrepreneur is commit-default and builds a low-bar experiment. Under a strong outside option, the entrepreneur is pivot-default and builds a high-bar experiment. Belief has not changed. The action threshold has. Second, the entrepreneurs who actually run experiments are selected. Entrepreneurs far above the threshold commit without experimenting, and entrepreneurs far below it pivot without experimenting. A low-bar experiment is therefore not the signature of the most optimistic entrepreneur. It is the signature of an entrepreneur near the commit side of the action threshold. Focus locates an entrepreneur relative to the threshold, but it does not by itself separate a strong prior from a weak outside option.

Proposition 4.2 (Co-movement can reverse experimental focus)**.** *If the outside option rises with belief faster than perceived venture value does, so that* $\mathrm{Cov}(\omega, \theta_S)/\mathrm{Var}(\theta_S) > G$*, then the map from optimism to focus reverses. The most optimistic entrepreneurs build high-bar experiments rather than low-bar experiments. Beyond the pivot belief* θ_{pivot} (eq. (35))*, they decline to experiment and take the outside option. Read in order of increasing optimism across the full population, entrepreneurs commit without experimenting, run a low-bar experiment, run a high-bar experiment, and then pivot without experimenting.*

The standard Bayesian entrepreneurship model with a common outside option implies strict sorting. The most pessimistic entrepreneurs pivot without experimenting. Moderately pessimistic entrepreneurs run high-bar experiments. Moderately optimistic entrepreneurs run low-bar experiments. The most optimistic entrepreneurs commit without experimenting. The left panel of Figure 5 shows this pattern.

Proposition 4.2 shows that this sorting depends on the common-outside-option assumption. When more optimistic entrepreneurs also have stronger outside options, the population moves along a steeper co-movement line in the (θ_S, ω) plane. The same four behavioral regions are still present, but the population passes through them in a different order. The most optimistic entrepreneurs now sit in the high-bar region or beyond it. They build the most skeptical experiments, or they do not experiment and pivot. The same observed focus can therefore read belief forward, backward, or not at all, depending on how outside options move with priors.

The Serial Entrepreneur. Consider an entrepreneur on a third venture with a public track record. Each increment of demonstrated conviction raises the perceived value of the new venture. But it may raise the entrepreneur's alternatives by even more. The prior rises, but so do the value of an acqui-hire, the valuation attached to the entrepreneur's name, the team that would follow the entrepreneur elsewhere, and the senior operating role available outside the venture. This entrepreneur can be highly optimistic and still be pivot-default because the outside option has outrun the venture. If the entrepreneur experiments, the experiment is high-bar. It requires strong positive evidence before commitment, not because the entrepreneur doubts the idea, but because the action threshold is high. Past a point, the entrepreneur does not experiment and takes the outside option. The most optimistic entrepreneur in the population can therefore be the one running the most skeptical experiment, or none at all. This is the reverse of what a belief-only reading of focus would predict.[8]

Focus is the first design choice in the model; Section 3 completes the design problem, recording what the choice is worth and letting the entrepreneur choose the accuracy behind it.

3 Endogenous Experimental Design

Section 2 treated the experiment as given and showed how entrepreneur-specific outside options change the value of experimentation, the decision to experiment, and the focus of the experiment. We now let entrepreneurs choose the design of the experiment itself. This adds a new observable margin: rigor. Focus tells us which mistake the entrepreneur is trying to avoid; rigor tells us how much is at stake. The key result is that endogenous design does not overturn the earlier threshold logic. Instead, it sharpens it. Once the entrepreneur chooses

[8]The attenuation result does not require the flip condition. Any positive co-movement, $0 < k < G$, already weakens the map from focus to belief because the cross-sectional gradient of the launch return is $dR/d\theta_S = G - k$ (eq. (6)). The full reversal requires $k > G$. An empirically cautious statement is that positive co-movement flattens the focus-belief relationship. The reversal is the sharp case, and it matters most when belief dispersion is ability-driven rather than idea-specific.

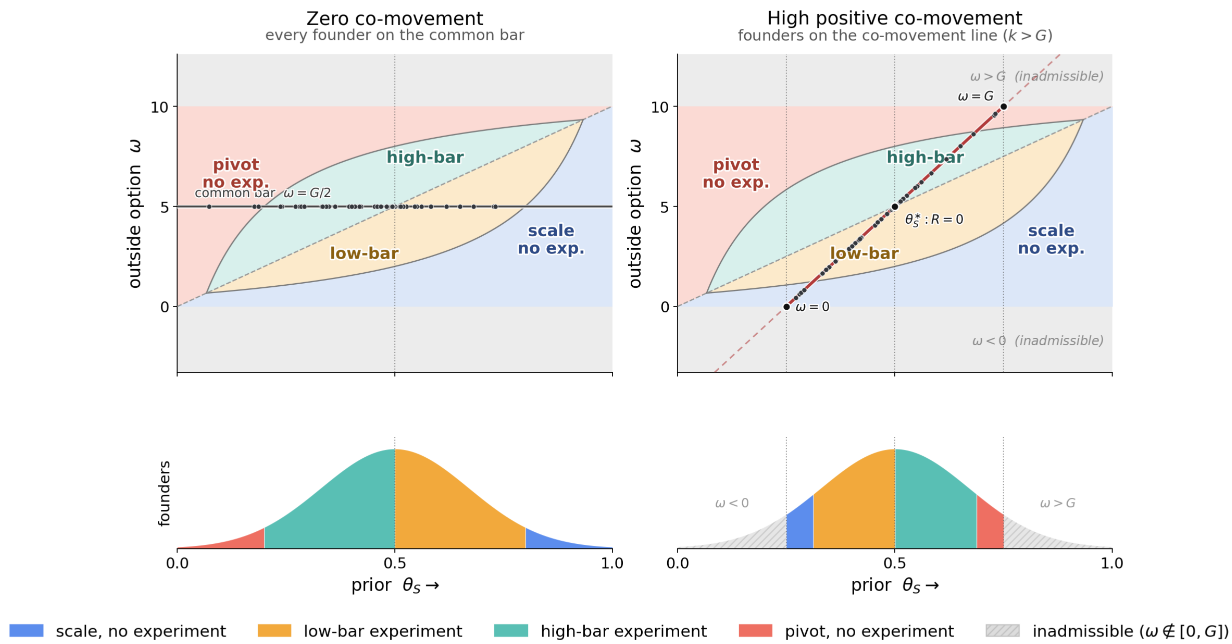

Figure 5. Focus and the source of belief dispersion. Each panel partitions the (θ_S, ω) plane into four behaviors. The entrepreneur can commit without experimenting, run a low-bar experiment, run a high-bar experiment, or pivot without experimenting. The dashed line is the action threshold $\omega = \theta_S G$, and the two break-even curves are $V_c = 0$ and $V_p = 0$. The lens between the curves is the region where experimentation pays off. Dots are entrepreneurs, and the histogram beneath each panel shows their density in θ_S, shaded by behavior. In the left panel, the outside option is common ($k = 0$), so every entrepreneur sits on the line $\omega = G/2$. Read in order of increasing optimism, behavior runs pivot, high-bar, low-bar, commit. In the right panel, the outside option is more belief-elastic than the venture ($k > G$). Entrepreneurs ride a steep co-movement line, and the order reverses. Read in order of increasing optimism, behavior runs commit, low-bar, high-bar, pivot. The most optimistic entrepreneurs build the most skeptical experiments, or decline to experiment and pivot. Parameters are $G = 10$, $\bar{q} = 0.9$, and $c = 0.5$.

focus optimally, accuracy is valuable only to the extent that it insures the regret left exposed by the default action. For commit-default entrepreneurs, rigor rises with the outside option because the experiment protects a more valuable fallback. For pivot-default entrepreneurs, rigor falls with the outside option because the incremental upside from the venture is smaller. A sloppy-looking experiment from a well-positioned entrepreneur leaning toward exit may not reflect low diligence; it may reveal a strong outside option. Observed experimental design therefore contains information about hidden outside options, not only about beliefs. Endogenous design turns experiments from mere learning devices into empirical traces of outside options: focus reveals which margin the entrepreneur is on, while rigor reveals how much is at stake.

Table 1 highlights the model implications by comparing the model with and without endogenous experimentation with the benchmark model from Agrawal et al. (2026). The common-outside-option benchmark links experimentation and focus primarily to beliefs. Heterogeneous outside options break that link by making the action threshold entrepreneur-specific. Endogenous design adds a further implication: the rigor of the experiment contains information about the hidden outside option that shapes entrepreneurial action. The proofs of all results are collected in Appendix A.7.

Table 1. What heterogeneous outside options change, and what endogenous experimental design adds.

	Common outside option	Heterogeneous outside options (Props. 1–4)	. . . with endogenous experimental design (Prop. 5)
Who experiments?	Moderate believers. Strong believers act—commit or pivot—without experimenting.	Entrepreneurs near their *own* action threshold, at any belief. Strong co-movement swaps the ends: the most pessimistic commit without experimenting, the most optimistic pivot without experimenting.	Same qualitative prediction: the same entrepreneurs, near their own thresholds.
Experimental focus	Focus reads belief. Optimists run low-bar experiments, pessimists high-bar; the map flips at the one shared threshold.	Focus reads the launch return, not belief. The same prior supports opposite experiments; an outside-option shock reverses focus with no change in belief; strong co-movement reverses the map from optimism to focus.	Same qualitative prediction: focus is set before rigor and does not change.
What is accuracy worth?	The standard model is silent: experimentation is free, so more accuracy is always preferred.	Worth the sum of the two expected regrets. Co-movement bends the sum: the most confident have the least of either regret to insure, and value accuracy least.	Experimental rigor increases in the outside option for commit-default founders and decreases in the outside option for pivot-default founders.
What does the observed experiment reveal?	Experimental focus reveals belief.	Focus reveals the balance between belief strength and outside-option attractiveness.	High rigor signals a strong outside option for commit-default entrepreneurs and a weak one for pivot-default entrepreneurs.

Table 1 separates the robustness results from the new design results. Endogenizing experimental design does not change which entrepreneurs experiment or which focus they choose. It changes what rigor is worth and what observed rigor reveals. We now solve the design problem by backward induction.

Assumption 1 (Experimental design environment)**.** *If the entrepreneur runs an experiment, she chooses total accuracy*

$$Q \equiv q_G + q_B \in (1, 2],$$

and chooses focus, $q_G - q_B$, along the feasible frontier. The cost of total accuracy is

$$C(Q) = c + \frac{\kappa}{2}(Q-1)^2,$$

where $c > 0$ is the setup cost of running any experiment and $\kappa > 0$ governs the marginal cost of accuracy. The entrepreneur chooses accuracy anticipating the optimal choice of focus.

We solve the design problem by backward induction. For a fixed level of total accuracy Q, the entrepreneur first chooses the focus of the experiment. Proposition 4.1 showed that focus is selected by the sign of the launch return,

$$R = \theta_S G - \omega.$$

The next lemma restates that result in the endogenous-design environment and records the value of the optimally focused experiment.

Lemma 1 (Optimal focus)**.** *Under Assumption 1, fix total accuracy $Q \in (1, 2]$. The value of experimentation is linear in focus. Its slope is the launch return R. Therefore, a commit-default entrepreneur, for whom $R > 0$, chooses the low-bar experiment,*

$$\{q_G, q_B\} = \{1, Q-1\},$$

while a pivot-default entrepreneur, for whom $R < 0$, chooses the high-bar experiment,

$$\{q_G, q_B\} = \{Q-1, 1\}.$$

At $R = 0$, all feasible allocations of focus with the same total accuracy are payoff-equivalent. The optimally focused experiment has net value

$$V^*(Q) = (Q-1)\min\{v_{mo}, v_{cm}\} - C(Q),$$

or equivalently,

$$V^*(Q) = \begin{cases} (Q-1)(1-\theta_S)\omega - C(Q), & \text{if } R > 0, \\ \\ (Q-1)\theta_S(G-\omega) - C(Q), & \text{if } R < 0. \end{cases}$$

Lemma 1 shows that focus removes the larger of the two regrets from the design problem. A commit-default entrepreneur runs an experiment to avoid a costly mistake: a negative result can push the entrepreneur off the venture and onto the outside option. Once focus is chosen for that purpose, the upside G no longer enters the value of experimentation for that entrepreneur. What remains at stake is the outside option that would be lost by mistakenly continuing. Conversely, a pivot-default entrepreneur runs an experiment to avoid a missed opportunity, so the relevant stake is the upside of the venture net of the outside option.[9]

Corollary 1 (Focus and accuracy are substitutes). *Under optimal focus, the marginal gross value of total accuracy is*

$$\frac{\partial}{\partial Q}\left[(Q-1)\min\{v_{mo}, v_{cm}\}\right] = \min\{v_{mo}, v_{cm}\}.$$

This is weakly below the value of raw accuracy in a symmetric experiment,

$$\min\{v_{mo}, v_{cm}\} \leq \frac{1}{2}(v_{mo} + v_{cm}),$$

with equality only at the action threshold. Focus and accuracy are therefore substitutes: by directing the experiment toward the error left exposed by the default action, focus lowers the marginal value of additional undirected accuracy.

Proposition 4.2 (Optimal rigor). *Under Assumption 1 and optimal focus, suppose the entrepreneur runs an experiment and the accuracy choice is interior. Then optimal total*

[9] In the outer stage, the entrepreneur chooses accuracy and pays for it. Because the value of a binary experiment is linear in its error rates, a constant marginal cost of accuracy produces an all-or-nothing choice; an interior accuracy requires a strictly convex cost, and the simplest is quadratic, with the same setup cost c as throughout.

accuracy is

$$Q^* = 1 + \frac{(1-\theta_S)\omega}{\kappa} \quad \textit{for a commit-default entrepreneur,}$$

and

$$Q^* = 1 + \frac{\theta_S(G-\omega)}{\kappa} \quad \textit{for a pivot-default entrepreneur.}$$

Thus rigor rises with the outside option for commit-default entrepreneurs,

$$\frac{\partial Q^*}{\partial \omega} = \frac{1-\theta_S}{\kappa} > 0,$$

and falls with the outside option for pivot-default entrepreneurs,

$$\frac{\partial Q^*}{\partial \omega} = -\frac{\theta_S}{\kappa} < 0.$$

For commit-default entrepreneurs, chosen informativeness Q^*-1 *is independent of the upside* G *and has unit elasticity with respect to the outside option.*

Proposition 4.2 separates the direction of the experiment from its intensity. Focus is determined by the sign of the launch return. Rigor is determined by the size of the regret that focus leaves exposed. A commit-default entrepreneur with a stronger outside option builds a more rigorous experiment because a costly mistake now sacrifices a more valuable alternative. A pivot-default entrepreneur with a stronger outside option builds a less rigorous experiment because the incremental upside from committing to the venture has shrunk.

Corollary 2 (Participation and minimum observed rigor)**.** *Under Assumption 1 and optimal focus, an interior experiment is run only if the exposed regret is large enough to cover the setup cost. For commit-default entrepreneurs this requires*

$$\omega > \frac{\sqrt{2\kappa c}}{1-\theta_S},$$

while for pivot-default entrepreneurs this requires

$$G - \omega > \frac{\sqrt{2\kappa c}}{\theta_S}.$$

Equivalently, an entrepreneur runs an interior experiment only if

$$\min\{v_{mo}, v_{cm}\} > \sqrt{2\kappa c}.$$

Every observed experiment therefore clears a minimum rigor floor. When $c \leq \kappa/2$, any interior experiment that is run satisfies

$$Q^* \geq 1 + \sqrt{\frac{2c}{\kappa}}.$$

When $c > \kappa/2$, any profitable experiment must be constrained at the upper bound, $Q^* = 2$.

Corollary 2 highlights selection in observed experimental accuracy. Entrepreneurs whose optimal tests would be too weak do not run experiments at all. Observed experiments are therefore selected toward sufficiently high rigor, and observed rigor understates the full cross-sectional variation in what entrepreneurs have at stake.

Proposition 4.2 (Co-movement and the rigor of experiments)**.** *Let outside options rise with optimism along*

$$\omega = a + k\theta_S.$$

Among entrepreneurs who experiment and choose interior accuracy, optimal rigor is

$$Q^* - 1 = \frac{(1-\theta_S)(a + k\theta_S)}{\kappa} \quad \text{on the commit-default branch,}$$

and

$$Q^* - 1 = \frac{\theta_S(G - a - k\theta_S)}{\kappa} \quad \text{on the pivot-default branch.}$$

The cross-sectional rigor profile is therefore single-peaked over the experimenting region, and the location of the peak depends on the co-movement between beliefs and outside options.

In particular:

1. *If outside options rise sufficiently quickly with beliefs, rigor increases with optimism among commit-default entrepreneurs, reversing the fixed-outside-option pattern.*
2. *Under strong positive co-movement, the most rigorous experiment in the population is built not by the entrepreneur exactly at the decision boundary, but by an entrepreneur who would abandon the venture by default.*
3. *If $k > G$, outside options rise faster than perceived venture value. The ordering of behavior follows Proposition 4.2: entrepreneurs move from committing without experimenting, to low-bar experiments, to high-bar experiments, and then to pivoting without experimenting. Along this path, rigor rises through the low-bar region, peaks inside the high-bar region, and then falls to the participation floor before the most optimistic entrepreneurs exit.*

The mechanism is the same threshold logic as before, now applied to the intensity of the experiment. Holding the outside option fixed, greater optimism reduces the costly mistake that a commit-default entrepreneur needs to insure. But when outside options rise with beliefs, the fallback becomes more valuable at the same time. Where this second force dominates, more optimistic commit-default entrepreneurs build more rigorous experiments. Among pivot-default entrepreneurs, the opposite force eventually dominates: a stronger outside option reduces the incremental gain from discovering that the venture is good, so rigor falls.

Corollary 3 (The kink identifies the co-movement gap)**.** *At the focus switch, the cross-sectional rigor profile is continuous but kinked. The difference between the slopes of the two rigor branches is*

$$\frac{G-k}{\kappa},$$

so the kink has magnitude

$$\frac{|G-k|}{\kappa}.$$

The kink vanishes exactly when outside options and perceived venture value move one-for-one, $k = G$. Observing both experimental focus and experimental rigor therefore identifies the co-movement gap $G-k$: focus reveals which side of the action threshold the entrepreneur occupies, while rigor reveals the size of the regret being insured.

Across the two stages of design, focus is a direction and rigor is a magnitude. Focus is set by the sign of the launch return, $R = \theta_S G - \omega$. Rigor is set by the size of the regret that the default action leaves exposed: the costly mistake for commit-default entrepreneurs and the missed opportunity for pivot-default entrepreneurs. Endogenous experimental design therefore turns experiments from mere learning devices into empirical traces of outside options. A single observed experiment cannot separately identify belief and outside-option value, but cross-sectional variation in focus and rigor reveals how tightly beliefs and outside options are bundled.

4 Discussion

The analysis makes outside options an explicit state variable in entrepreneurial experimentation. In Bayesian models, an experiment changes action by changing beliefs about the focal idea. Our results show that this is only half of the comparison: the relevant action threshold also depends on what the entrepreneur can do instead. As a result, identical priors and identical signals can rationally generate different actions, different experimental focus, and different experimental rigor across entrepreneurs. This matters theoretically because experimental strategy is actor-specific as well as idea-specific. It matters empirically because persistence, pivoting, exit, and scaling are not clean signals of beliefs or learning when outside options are heterogeneous and unobserved.

The model delivers three implications for entrepreneurial experimentation: value, contingency, and design. Outside options change the value of experimentation because evidence is useful only when it can improve the comparison between the focal idea and the entrepreneur's alternative. They change the contingency of experimentation because the same prior can imply commitment, experimentation, or pivoting depending on the outside option. They also change experimental design because the relevant alternative determines whether the entrepreneur seeks credible bad news before committing, credible good news before pivoting, and how much rigor is worth purchasing. The central point is that experiments are not evaluated against beliefs alone. They are evaluated against the opportunity cost of acting on those beliefs.

Outside options change both the entrepreneur's optimal experimental strategy and the inference researchers can draw from observed experimentation. Theoretically, they move experimental strategy from a property of the idea to a joint property of the idea and the actor evaluating it. Empirically, they make persistence, pivoting, exit, scaling, and experimental design potentially consistent with beliefs, learning, or opportunity costs. This does not displace existing accounts based on overconfidence, disciplined learning, or scientific training. It adds a missing state variable that can generate the same behaviors through rational choice. The remainder of this section develops these implications first for Bayesian entrepreneurship and experimental strategy, and then for empirical research on entrepreneurial action.

4.1 Implications for Theory

Bayesian Entrepreneurship

Bayesian entrepreneurship links evidence to action through belief updating. Entrepreneurs observe signals, update their beliefs, and then decide whether to commit to the focal idea or pivot away. That mapping is complete only when outside options are common across entrepreneurs. Once outside options differ, the posterior belief is no longer a sufficient statistic for action. The same posterior can justify commitment for an entrepreneur with a weak alternative and pivoting for an entrepreneur with a strong alternative. Action depends not only on what the entrepreneur believes about the idea, but also on whether the idea is worth pursuing relative to what the entrepreneur could do instead.

This changes what experimental evidence is valuable for. Evidence matters not only because it changes beliefs about the focal idea, but because it can change the entrepreneur's optimal action relative to the outside option. For a commit-default entrepreneur, negative evidence is valuable when it prevents continued investment in an idea that is not worth its opportunity cost. For a pivot-default entrepreneur, positive evidence is valuable when it justifies continuing with an idea that would otherwise be abandoned. Thus the same signal can have different strategic value for different entrepreneurs because it resolves different comparisons.

The effect is strongest when priors and outside options co-move. For high-growth entrepreneurs, the same ability, experience, market access, or reputation that supports a strong prior may also create valuable alternatives. In those settings, heterogeneous outside options can weaken, and in some cases reverse, the standard prediction that stronger priors require less demanding evidence (Agrawal et al., 2026; Gans, 2023). The entrepreneur who appears most confident may also have the strongest alternative and may therefore require especially strong evidence before continuing to invest.

The contribution to Bayesian entrepreneurship is therefore not to replace beliefs, signals, experiment costs, or updating. It is to add the opportunity-cost side of the choice problem to the evidence-to-action mapping. Holding outside options fixed remains useful for isolating how beliefs and information shape experimentation. Allowing outside options to vary is essential for settings in which entrepreneurs also differ in what they give up by continuing, especially high-growth and opportunity-driven entrepreneurship.

Experimental Strategy

The outside-option logic also contributes to the broader literature on experimental strategy. Recent work has moved beyond the view of experiments as simple instruments for producing information at low cost (Chavda et al., 2024; Contigiani and Levinthal, 2019). Experiments can be tied to theories of value, guide search, shape sequences of commitment and pivoting, persuade stakeholders, mobilize resources, and even change the economic potential of the idea being tested (Adner and Levinthal, 2024; Felin, Gambardella, Stern, and Zenger, 2020; Shelef, Karp, and Wuebker, 2024a; Shelef, Wuebker, and Barney, 2024b; Valentine et al., 2024; Zellweger and Zenger, 2022).

Our contribution is to add the entrepreneur's alternative as another contingency in experimental strategy. The value of an experiment depends not only on what the experiment reveals about the idea, changes about the idea, or persuades others to believe. It also depends on what the entrepreneur would otherwise do. The same experiment can therefore have different strategic value for different entrepreneurs. For one entrepreneur, credible negative evidence may be valuable because it prevents continued investment that would sacrifice a strong outside option. For another, credible positive evidence may be valuable because it justifies staying with an idea that would otherwise be abandoned.

This makes experimental strategy actor-specific as well as idea-specific. An experiment's design cannot be evaluated only by its informativeness, credibility, or financial cost (Agrawal et al., 2026). It must also be evaluated against the opportunity cost of the actor running it and the entrepreneur's position relative to the action threshold. On the commit side of the threshold, experimental design protects against costly mistakes; on the pivot side, it protects against missed opportunities. Focus indicates which margin the entrepreneur is on, while rigor indicates how much is at stake on that margin.

Implications for Empirical Work

The model also changes how empirical work should interpret entrepreneurial action. A large empirical literature uses persistence, pivoting, exit, and scaling to infer what entrepreneurs believe, what they have learned, or whether they use evidence well. Persistence after discouraging evidence is often read as confidence, limited learning, or a failure to discipline prior beliefs (Camuffo et al., 2020; Camuffo, Gambardella, Messinese, Novelli, Paolucci, and Spina, 2024a; Van den Steen, 2011). Pivoting or exit is often read as evidence that entrepreneurs have learned enough to redirect or stop investing (Camuffo et al., 2020, 2024a; Kirtley and O'Mahony, 2020). Scaling is often read as commitment after learning about the idea, strategy, or market (Agrawal et al., 2021; Gans, 2023; Gans et al., 2019; Lee and Kim, 2024). Our model shows that these interpretations are under-identified when outside options are unobserved. The same action can reflect beliefs, learning, bias, or the opportunity cost of continuing.

Consider persistence. An entrepreneur who continues after discouraging evidence may be ignoring bad news, but she may also have a weak outside option: the focal idea remains better than what she could do instead. Persistence can therefore reflect overconfidence, escalation, or low opportunity costs. Exit and pivoting raise the same identification problem. An entrepreneur with a strong outside option may abandon an idea not because the idea is bad, but because the idea is not good enough relative to the alternative. She may even pivot after positive but modest evidence if the posterior still fails to clear her action threshold. Conversely, an entrepreneur with a weak outside option may continue after discouraging evidence because the alternative is worse. Observed action is therefore not a sufficient statistic for beliefs or learning.

This logic also matters for interventions designed to improve entrepreneurial learning. Accelerators are often interpreted as helping entrepreneurs learn faster and terminate weak ideas more efficiently (Hallen, Cohen, and Bingham, 2020; Yu, 2020). Scientific training is similarly interpreted as improving hypothesis testing and evidence-based action (Camuffo et al., 2020, 2024a). These interpretations may be correct, but such programs can also change outside options. Accelerators provide networks, credentials, investor access, mentoring, and labor-market visibility. Training programs can change what entrepreneurs know, whom they meet, and what paths become available. A higher exit rate may therefore reflect better learning, stronger outside options, or both. The same logic applies to founding teams:

cofounder entry, exit, role allocation, and venture dissolution may reflect not only what the team has learned, but also the changing alternatives available to individual team members (Eesley, Hsu, and Roberts, 2014; Hellmann and Wasserman, 2017).

The identification problem is deeper than fixed unobserved heterogeneity. Outside options are entrepreneur-specific, often unobserved, and frequently time-varying. They include foregone wages, alternative venture ideas, acquisition possibilities, investor access, employment opportunities, and the subjective option value of paths not taken. Founder fixed effects may absorb stable differences across entrepreneurs, but they do not solve the problem when outside options change during the venturing process. A job offer can arrive, a new project can become salient, an investor can open another path, or a professional network can make exit more attractive. In these settings, the same entrepreneur may face different action thresholds at different moments, and those threshold changes may drive experimentation, persistence, pivoting, or exit.

The constructive implication is that empirical work should measure alternatives as well as experiments. Data on priors, signal quality, experiment cost, experiment focus, and post-experiment action remain valuable, but the model suggests that researchers also need measures of what the entrepreneur could do instead: employment prospects, alternative projects, acquisition opportunities, human capital, investor access, and changes in networks or credentials. When direct measures are unavailable, experimental design may provide indirect evidence. Focus reveals whether the entrepreneur seeks credible bad news before committing or credible good news before pivoting; rigor reveals how much is at stake on that margin. Experiments are therefore not only outcomes to be explained. They are also potential measurements of the hidden outside options that shape entrepreneurial action.

5 Conclusion

Bayesian entrepreneurship has advanced the study of entrepreneurial action by showing how entrepreneurs use experiments to learn under uncertainty. In this framework, entrepreneurs hold priors about uncertain business ideas, generate evidence, update beliefs, and then decide whether to commit, continue, pivot, or exit. This paper builds on that logic by relaxing a useful simplification: that the value of the entrepreneur's outside option is common across entrepreneurs. Once outside options are allowed to differ, experimental evidence is no longer evaluated only against beliefs about the focal idea. It is also evaluated against what the entrepreneur gives up by continuing to pursue it.

The central implication is that experiments do not simply reveal whether an idea is good. They reveal whether the idea is good enough for the entrepreneur who possesses it, given that entrepreneur's alternative. This distinction changes the value of experimentation, the decision to experiment, and the design of experiments themselves. A strong outside option can make credible bad news especially valuable because it protects the entrepreneur from continuing with an idea that is not worth its opportunity cost. A weak outside option can make credible good news especially valuable because it may justify continuing with an idea that would otherwise be abandoned. When priors and outside options co-move, the familiar prediction that stronger beliefs require less demanding evidence can weaken or reverse.

Recognizing outside options increases the explanatory power of Bayesian entrepreneurship without diminishing the importance of beliefs, learning, or scientific experimentation. Pivoting need not imply pessimism. Persistence need not imply overconfidence. Exit need not imply that the idea was bad. Each action may instead reflect the comparison between the focal idea and the entrepreneur's next-best alternative. For theory, this means that experimental strategy is actor-specific as well as idea-specific. For empirical research, it means that observed action cannot be interpreted without attention to the alternatives entrepreneurs face.

References

Adner, Ron, Daniel A Levinthal. 2024. Strategy experiments in nonexperimental settings: Challenges of theory, inference, and persuasion in business strategy. *Strategy Science*, **9**(4) 311–321.

Agarwal, Rajshree, Francesca Bacco, Arnaldo Camuffo, Andrea Coali, Alfonso Gambardella, Haji Msangi, Steven T Sonka, Anna Temu, Betty Waized, Audra Wormald. 2023. Does a Theory-of-Value Add Value? Evidence from a Randomized Control Trial with Tanzanian Entrepreneurs. *SSRN*.

Agrawal, Ajay, Arnaldo Camuffo, Alfonso Gambardella, Joshua Gans, Erin Scott, Scott Stern. 2026. The foundations of Bayesian entrepreneurship. Ajay Agrawal, Arnaldo Camuffo, Alfonso Gambardella, Joshua Gans, Erin L. Scott, Scott Stern, eds., *Bayesian Entrepreneurship*, chap. 1. 3–62.

Agrawal, Ajay, Joshua S. Gans, Scott Stern. 2021. Enabling Entrepreneurial Choice. *Management Science*, **67**(9) 5301–5967.

Arora, Ashish. 2026. Bayesian entrepreneurship: A user's perspective. Ajay Agrawal, Arnaldo Camuffo, Alfonso Gambardella, Joshua Gans, Erin L. Scott, Scott Stern, eds., *Bayesian Entrepreneurship*, chap. 4. 109–114.

Åstebro, Thomas, Frank M. Fossen, Cédric Gutierrez. 2026. Entrepreneurs: Clueless, Biased, Poor Heuristics, or Bayesian Machines? Ajay Agrawal, Arnaldo Camuffo, Alfonso Gambardella, Joshua Gans, Erin L. Scott, Scott Stern, eds., *Bayesian Entrepreneurship*, chap. 5. 117–140.

Barney, Jay B., Lyda S. Bigelow, Orie Shelef, Robert Wuebker. 2024. Experimental Strategy With Entangled Ideas.

Camuffo, Arnaldo, Alessandro Cordova, Alfonso Gambardella, Chiara Spina. 2020. A scientific approach to entrepreneurial decision making: Evidence from a randomized control trial. *Management Science*, **66**(2) 564–586.

Camuffo, Arnaldo, Alfonso Gambardella, Danilo Messinese, Elena Novelli, Emilio Paolucci, Chiara Spina. 2024a. A scientific approach to entrepreneurial decision-making: Large-scale replication and extension. *Strategic Management Journal*, **45**(6) 1209–1237.

Camuffo, Arnaldo, Alfonso Gambardella, Andrea Pignataro. 2024b. Theory-Driven Strategic Management Decisions. *Strategy Science*, **9**(4) 382–396.

Chavda, Ankur, Joshua S. Gans, Scott Stern. 2024. Theory-Driven Entrepreneurial Search. *SSRN*.

Chen, John S, Daniel Walter Elfenbein, Hart E. Posen, Ming Zhu Wang. 2022. Programs of Experimentation and Pivoting for (Overconfident) Entrepreneurs. *Academy of Management Review*.

Contigiani, Andrea, Daniel A. Levinthal. 2019. Situating the construct of lean start-up: adjacent conversations and possible future directions. *Industrial and Corporate Change*, **28**(3) 551–564.

Eesley, Charles E., David H. Hsu, Edward B. Roberts. 2014. The contingent effects of top management teams on venture performance: Aligning founding team composition with innovation strategy and commercialization environment. *Strategic Management Journal*, **35**(12) 1798–1817.

Felin, Teppo, Alfonso Gambardella, Elena Novelli, Todd Zenger. 2024. A scientific method

for startups. *Journal of Management* 01492063231226136.
Felin, Teppo, Alfonso Gambardella, Scott Stern, Todd Zenger. 2020. Lean startup and the business model: Experimentation revisited. *Long Range Planning*, **53**(4) 101889.
Felin, Teppo, Todd R. Zenger. 2009. Entrepreneurs as theorists: on the origins of collective beliefs and novel strategies. *Strategic Entrepreneurship Journal*, **3**(2) 127–146.
Gambardella, Alfonso, Luca Gius, Scott Stern. 2026. Homo entrepreneuricus. Ajay Agrawal, Arnaldo Camuffo, Alfonso Gambardella, Joshua Gans, Erin L. Scott, Scott Stern, eds., *Bayesian Entrepreneurship*, chap. 2. 63–86.
Gans, Joshua S. 2022. A Theory of Visionary Disruption.
Gans, Joshua S. 2023. A Theory of Visionary Disruption. *Management Science*.
Gans, Joshua S. 2026. Knightian uncertainty and bayesian entrepreneurship. Ajay Agrawal, Arnaldo Camuffo, Alfonso Gambardella, Joshua Gans, Erin L. Scott, Scott Stern, eds., *Bayesian Entrepreneurship*, chap. 3. 87–110.
Gans, Joshua S., Scott Stern, Jane Wu. 2019. Foundations of entrepreneurial strategy. *Strategic Management Journal*, **40**(5) 736–756.
Gompers, Paul, Anna Kovner, Josh Lerner, David Scharfstein. 2010. Performance persistence in entrepreneurship. *Journal of financial economics*, **96**(1) 18–32.
Hallen, Benjamin L., Susan L. Cohen, Christopher B. Bingham. 2020. Do accelerators work? if so, how? *Organization Science*, **31**(2) 378–414.
Hellmann, Thomas, Noam Wasserman. 2017. The first deal: The division of founder equity in new ventures. *Management Science*, **63**(8) 2647–2666.
Kim, Hyunjin. 2026. Learning from data in entrepreneurial experimentation. Ajay Agrawal, Arnaldo Camuffo, Alfonso Gambardella, Joshua Gans, Erin L. Scott, Scott Stern, eds., *Bayesian Entrepreneurship*, chap. 9. 197–214.
Kirtley, Jacqueline, Siobhan O'Mahony. 2020. What is a pivot? Explaining when and how entrepreneurial firms decide to make strategic change and pivot. *Strategic Management Journal*.
Krieger, Joshua L. 2026. Charting (and updating) the path: A bayesian perspective on entrepreneurial learning. Ajay Agrawal, Arnaldo Camuffo, Alfonso Gambardella, Joshua Gans, Erin L. Scott, Scott Stern, eds., *Bayesian Entrepreneurship*, chap. 10. 217–230.
Lee, Kim. 2024. Placeholder: supply title. *PLACEHOLDER: supply venue*.
Novelli, Elena. 2026. A scientific approach to entrepreneurial decision-making as bayesian learning. Ajay Agrawal, Arnaldo Camuffo, Alfonso Gambardella, Joshua Gans, Erin L. Scott, Scott Stern, eds., *Bayesian Entrepreneurship*, chap. 7. 161–174.
Posen, Hart E., Orie Shelef, Robert Wuebker. 2024. Sunk Costs And Experimental Strategy. *SSRN*.
Scott, Erin L. 2026. Entrepreneurial strategy pedagogy and practice: Learning, Commitment, and Choice. Ajay Agrawal, Arnaldo Camuffo, Alfonso Gambardella, Joshua Gans, Erin L. Scott, Scott Stern, eds., *Bayesian Entrepreneurship*, chap. 12. 249–268.
Shelef, Orie, Rebecca Karp, Robert Wuebker. 2024a. Business Experiments As Persuasion.
Shelef, Orie, Robert Wuebker, Jay B Barney. 2024b. Heisenberg effects in experiments on business ideas. *Academy of Management Review*.
Snyder, Jason, Robert Wuebker, Todd Zenger. 2024. Average treatment effects and unique strategies. *Available at SSRN 5050066*.
Spina, Chiara. 2026. What helps entrepreneurs learn: Reflections, new directions, and open

questions. Ajay Agrawal, Arnaldo Camuffo, Alfonso Gambardella, Joshua Gans, Erin L. Scott, Scott Stern, eds., *Bayesian Entrepreneurship*, chap. 13. 269–292.

Valentine, Jacob, Elena Novelli, Rajshree Agarwal. 2024. The Theory-Based View and Strategic Pivots: The Effects of Formalization in Cognition and Action on the Type and Nature of Pivots. *Available at SSRN 4730043*.

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

Yu, Sandy. 2020. How do accelerators impact the performance of high-technology ventures? *Management Science*, **66**(2) 530–552.

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

Zellweger, Thomas M., Todd R. Zenger. 2022. Entrepreneurs as scientists, Bayesian inference, and belief revision. *Journal of Business Venturing Insights*, **18** e00350.

A Derivations and Proofs

This appendix is self-contained. Section A.1 fixes payoffs and notation and derives the two value functions from first principles. Section A.2 states the one lemma, the chain rule applied to a population in which outside options move with beliefs, that is used in every co-movement proof. The remaining subsections prove the eight baseline propositions in the order they appear in the text, and a final subsection proves the design lemma, the two corollaries, and the pair of design propositions of Section 3. The arguments are deliberately elementary: each derivative is written out, and each sign is traced to the assumption that delivers it.

A.1 Setup: Payoffs, Actions, and the Value of Experimentation

An entrepreneur faces a single venture whose quality is either *good* or *bad*, and holds a prior $\theta_S = \Pr(\text{good})$. *Committing* to a good venture earns the upside $G > 0$; committing to a bad one earns 0; *pivoting*, redeploying to the next-best use of time, earns the outside option ω. Expected payoffs are therefore $\theta_S G$ from committing and ω from pivoting, so the entrepreneur's *default action*, the action taken with no further information, is

$$\text{commit if } \theta_S G > \omega, \qquad \text{pivot if } \theta_S G < \omega. \tag{11}$$

We call an entrepreneur *commit-default* in the first case and *pivot-default* in the second.

Before acting, the entrepreneur may run one binary experiment, at cost $c > 0$, that returns a *positive* or a *negative* result. The experiment is described by its true positive rate and true negative rate,

$$q_G = \Pr(\text{positive} \mid \text{good}), \qquad q_B = \Pr(\text{negative} \mid \text{bad}),$$

so that the false-negative rate is $1 - q_G$ (a good venture that comes back negative) and the false-positive rate is $1 - q_B$ (a bad venture that comes back positive). An experiment is *informative* when

$$q_G + q_B > 1. \tag{12}$$

Accuracy and lean. An experiment is the pair (q_G, q_B). Its overall accuracy is governed by the sum $q_G + q_B$ and its *lean* by the difference $q_G - q_B$, and we treat the two separately. For the value of experimentation, the decision to experiment, and the demand for accuracy (Sections A.3, A.4, and A.5) the experiment is *symmetric*, $q_G = q_B = \bar{q}$, so that $\bar{q} \in (\frac{1}{2}, 1)$ indexes accuracy and the lean is neutral. The choice of a lean, moving along a frontier $q_G + q_B = \text{const}$ toward the true positive rate or the true negative rate, is the question of focus, which the entrepreneur makes endogenously in Section A.6. The symmetric experiment is the midpoint of that focus frontier (32): accuracy asks how far out the frontier sits, focus which way to slide along it. We record the value functions below for general (q_G, q_B), because the focus analysis selects an asymmetric point; the value, timing, and accuracy results are the $q_G = q_B = \bar{q}$ specialization, in which the informativeness $q_G + q_B - 1$ reads $2\bar{q} - 1$ and the accuracy floor $q_B/(1 - q_G)$ reads $\bar{q}/(1 - \bar{q})$.

Maintained assumptions. Throughout, $0 < \theta_S < 1$, $0 < q_G < 1$, $0 < q_B < 1$, the informativeness condition (12) holds, $c > 0$, $G > 0$, and the outside option lies below the full upside,

$$0 < \omega < G. \tag{13}$$

The last assumption is natural: if the outside option exceeded even a sure-good venture the entrepreneur would never enter, and it holds automatically for any commit-default entrepreneur, for whom $\omega < \theta_S G < G$.

The value of experimenting. Suppose the entrepreneur runs the experiment. By Bayes' rule the probability of a positive result is $\Pr(\text{positive}) = \theta_S q_G + (1-\theta_S)(1-q_B)$, and the posterior that the venture is good, conditional on a positive result, is $\theta_{(+)} = \theta_S q_G / \Pr(\text{positive})$. Likewise $\Pr(\text{negative}) = \theta_S(1-q_G) + (1-\theta_S)q_B$ and $\theta_{(-)} = \theta_S(1-q_G)/\Pr(\text{negative})$. When each verdict is decisive, the optimal post-experiment rule is to *commit on a positive result and pivot on a negative one*, and we evaluate experimentation against this rule.[10] The expected payoff from this strategy is

$$\underbrace{\Pr(\text{positive})\,\theta_{(+)}\,G}_{=\theta_S q_G G} + \Pr(\text{negative})\,\omega - c = \theta_S q_G G + \big[\theta_S(1-q_G) + (1-\theta_S)q_B\big]\omega - c,$$

where the first equality uses $\Pr(\text{positive})\theta_{(+)} = \Pr(\text{positive and good}) = \theta_S q_G$.

The *value of experimenting* is this expected payoff minus the payoff from the default action.

For a **commit-default** entrepreneur the default earns $\theta_S G$, so

$$\begin{aligned} V_c &= \theta_S q_G G + \big[\theta_S(1-q_G) + (1-\theta_S)q_B\big]\omega - c - \theta_S G \\ &= (1-\theta_S)\,q_B\,\omega - \theta_S\,(1-q_G)\,(G-\omega) - c. \end{aligned} \tag{14}$$

The two terms read as follows. With probability $(1-\theta_S)q_B$ the experiment correctly returns negative on a bad venture the entrepreneur would otherwise have committed to; the entrepreneur pivots and keeps the outside option ω. With probability $\theta_S(1-q_G)$ the experiment wrongly returns negative on a good venture, and the entrepreneur forgoes the surplus $G-\omega$ that committing would have earned. Experimentation, for a commit-default entrepreneur, is insurance against committing to a bad idea, weighed against the risk of abandoning a good one.

For a **pivot-default** entrepreneur the default earns ω, so

$$\begin{aligned} V_p &= \theta_S q_G G + \big[\theta_S(1-q_G) + (1-\theta_S)q_B\big]\omega - c - \omega \\ &= \theta_S\,q_G\,(G-\omega) - (1-\theta_S)\,(1-q_B)\,\omega - c. \end{aligned} \tag{15}$$

Symmetrically: with probability $\theta_S q_G$ the experiment correctly returns positive on a good venture the entrepreneur would otherwise have pivoted away from, earning the surplus $G-\omega$;

[10] The rule is optimal when each verdict moves the posterior across the action threshold (11). Section A.4 shows that this decisiveness condition is exactly the condition for the experiment to have positive gross value, so the restriction is without loss for every entrepreneur who experiments.

with probability $(1-\theta_S)(1-q_B)$ it wrongly returns positive on a bad one, costing the outside option ω.

It is convenient to name the two expected stakes that recur throughout. The *expected costly mistake*, committing to a bad venture and forfeiting the outside option, and the *expected missed opportunity*, pivoting away from a good venture and forfeiting the upside, are

$$v_{\text{cm}} \;\equiv\; (1-\theta_S)\,\omega, \qquad v_{\text{mo}} \;\equiv\; \theta_S\,(G-\omega). \tag{16}$$

Each stake is an expected regret rather than the expected value of an action. Let $W^\star = \theta_S G + (1-\theta_S)\omega$ denote the perfect-foresight payoff, from committing in the good state and pivoting in the bad. The two stakes are the regrets of the two pure actions, $v_{\text{cm}} = W^\star - \theta_S G$ and $v_{\text{mo}} = W^\star - \omega$, where $\theta_S G$ and ω are the expected payoffs of committing and of pivoting. For $\omega > 0$ neither stake equals an action's expected value, but their difference does return an action comparison, because $W^\star$ cancels: $v_{\text{mo}} - v_{\text{cm}} = \theta_S G - \omega$, the launch return (20) below. At $\omega = 0$ the regret and value readings coincide ($v_{\text{mo}} = \theta_S G$, $v_{\text{cm}} = 0$). In this notation the two value functions read

$$V_c = q_B\, v_{\text{cm}} - (1-q_G)\, v_{\text{mo}} - c, \qquad V_p = q_G\, v_{\text{mo}} - (1-q_B)\, v_{\text{cm}} - c,$$

each pairing the protected stake, earned at the rate of the experiment's correct decisive verdict, against the opposite stake lost at the error rate.

Because V_c and V_p are each measured *net of the default action*, the entrepreneur runs the experiment exactly when the relevant value is positive:

$$\begin{aligned} \text{commit-default entrepreneur experiments} &\iff V_c > 0, \\ \text{pivot-default entrepreneur experiments} &\iff V_p > 0. \end{aligned} \tag{17}$$

It is convenient to name the gross value of the experiment, its value before paying c,

$$\text{WTP}_c \equiv V_c + c = (1-\theta_S) q_B\, \omega - \theta_S (1-q_G)(G-\omega), \qquad \text{WTP}_p \equiv V_p + c = \theta_S q_G (G-\omega) - (1-\theta_S)(1-q_B)\omega, \tag{18}$$

which is the entrepreneur's willingness to pay for the experiment; we write WTP for whichever branch the default selects.

Outside options across entrepreneurs. Entrepreneurs differ in both θ_S and ω. We summarize the joint distribution by the slope of the best linear predictor of ω on θ_S,

$$k \;\equiv\; \frac{\text{Cov}(\omega, \theta_S)}{\text{Var}(\theta_S)}, \qquad \text{so that} \qquad \mathbb{E}[\omega \mid \theta_S] = a + k\,\theta_S \tag{19}$$

for an intercept a. The slope k is an ordinary second moment, a co-movement divided by a variance, and not a distributional assumption; the results below need only this slope. The *fixed-option* propositions (X.1) hold ω fixed and vary it exogenously: they are statements about $\partial/\partial\omega$ at a given θ_S. The *co-movement* propositions (X.2) trace behavior across entrepreneurs *along the conditional mean* (19). Both turn on a single object, the *launch return*: the expected payoff of committing net of the outside option it forgoes, the signed distance of perceived value from the action threshold, and equivalently the difference of the two stakes (16),

$$R \;\equiv\; \theta_S G - \omega \;=\; v_{\text{mo}} - v_{\text{cm}}, \tag{20}$$

positive for commit-default entrepreneurs (the missed opportunity is the larger stake) and negative for pivot-default ones (the costly mistake is).

A.2 A Lemma on Cross-Sectional Gradients

Lemma 2 (Cross-sectional gradient)**.** *Let $X(\theta_S, \omega)$ be differentiable, and consider how X varies across entrepreneurs along the conditional mean* (19)*, i.e. as a function of θ_S with $\omega = a + k\theta_S$. Then*

$$\left.\frac{dX}{d\theta_S}\right|_{\text{c.m.}} = \frac{\partial X}{\partial \theta_S} + k\,\frac{\partial X}{\partial \omega}. \tag{21}$$

If $\partial X/\partial\omega \neq 0$, this total derivative vanishes at the critical co-movement

$$k^* = -\,\frac{\partial X/\partial \theta_S}{\partial X/\partial \omega}, \tag{22}$$

and is monotone in k. In particular, if $\partial X/\partial\theta_S > 0$ and $\partial X/\partial\omega < 0$ (opposite signs), then $k^* > 0$, the total derivative is positive for $k < k^*$, and it is negative for $k > k^*$: the cross-sectional relationship between X and the prior reverses sign *as the co-movement crosses k^*.*

Proof. Equation (21) is the chain rule: with $\omega = a + k\theta_S$ we have $d\omega/d\theta_S = k$, so $dX/d\theta_S = \partial X/\partial\theta_S + (\partial X/\partial\omega)(d\omega/d\theta_S) = \partial X/\partial\theta_S + k\,\partial X/\partial\omega$. Setting the right-hand side to zero and solving for k gives (22). The total derivative (21) is affine in k with slope $\partial X/\partial\omega$, hence monotone; when the two partials have opposite signs, $k^* = -(\partial X/\partial\theta_S)/(\partial X/\partial\omega) > 0$, and the affine function is positive below k^* and negative above it (for $\partial X/\partial\omega < 0$). $\square$

The lemma organizes every co-movement proof below. Holding ω fixed, the prior moves a behavioral object one way ($\partial X/\partial\theta_S$) and the outside option moves it the other ($\partial X/\partial\omega$); across entrepreneurs the co-movement k ties the two channels together, and the lemma determines which one prevails.

A.3 The Value of Experimentation (Propositions 1.1 and 1.2)

Proof of Proposition 1.1. Differentiate (14) in ω, treating θ_S as fixed:

$$\frac{\partial V_c}{\partial \omega} = \underbrace{(1-\theta_S)q_B}_{>0} + \underbrace{\theta_S(1-q_G)}_{\geq 0} \;>\; 0.$$

The first term is strictly positive because $0 < \theta_S < 1$ gives $1 - \theta_S > 0$ and $q_B > 0$; the second is nonnegative. (The derivative of $-\theta_S(1-q_G)(G-\omega)$ in ω is $+\theta_S(1-q_G)$ because ω enters with a minus sign inside $G - \omega$.) A stronger outside option therefore *raises* the value of experimentation for a commit-default entrepreneur. Differentiating (15) similarly,

$$\frac{\partial V_p}{\partial \omega} = -\,\theta_S q_G - \,(1-\theta_S)(1-q_B) = -\big[\underbrace{\theta_S q_G}_{>0} + \underbrace{(1-\theta_S)(1-q_B)}_{\geq 0}\big] \;<\; 0,$$

so a stronger outside option *lowers* the value of experimentation for a pivot-default entrepreneur. Neither sign uses informativeness: the result is driven by the fact that a false positive costs the entrepreneur the outside option while a false negative costs the upside, so raising ω raises the stake on one regret and lowers it on the other, and the two cases weight those regrets oppositely. The launch return $R = \theta_S G - \omega$ falls one-for-one in the outside option, $\partial R/\partial \omega = -1$. By Lemma 3, for fixed ω willingness to pay, and hence the value V, which differs from it only by the constant cost c, is single-peaked in the prior, with its maximum at $\theta_S = \omega/G$, the belief at which $R = 0$; this peak belief is strictly increasing in the outside option, $\partial(\omega/G)/\partial\omega = 1/G > 0$, and equals $\frac{1}{2}$ exactly at $\omega = G/2$. Under the symmetric experiment $q_G = q_B = \bar{q}$, the swap $(\theta_S, \omega) \mapsto (1-\theta_S, G-\omega)$ exchanges the regrets $v_{\text{cm}} \leftrightarrow v_{\text{mo}}$ and the value branches $V_c \leftrightarrow V_p$, leaving V unchanged; its unique fixed point is $(\frac{1}{2}, G/2)$. That this fixed point is the global maximum follows from (23): the value along the ridge, $\frac{\omega(G-\omega)}{G}(2\bar{q}-1) - c$, is maximized over ω at $\omega = G/2$. $\square$

The peak in Proposition 1.1 and the cross-section in Proposition 1.2 both turn on the shape of willingness to pay as a function of the prior, the binary-signal counterpart of the standard inverted-V.

Lemma 3 (Willingness to pay peaks at the decision boundary)**.** *Fix ω. As a function of the prior, willingness to pay* (18) *is continuous, strictly increasing on the pivot-default range $\theta_S < \omega/G$, and strictly decreasing on the commit-default range $\theta_S > \omega/G$. It therefore attains a unique maximum at the boundary $\theta_S = \omega/G$, where*

$$\text{WTP}_c = \text{WTP}_p = \frac{\omega(G-\omega)}{G}\,(q_G + q_B - 1) \;>\; 0. \tag{23}$$

Proof. On the commit-default range, differentiate WTP_c from (18):

$$\frac{d\text{WTP}_c}{d\theta_S} = -q_B\,\omega - (1-q_G)(G-\omega) \;<\; 0,$$

strictly negative because $q_B\omega > 0$ and, by (13), $(1-q_G)(G-\omega) > 0$. On the pivot-default range,

$$\frac{d\text{WTP}_p}{d\theta_S} = q_G\,(G-\omega) + (1-q_B)\,\omega \;>\; 0,$$

strictly positive for the same reason ($G - \omega > 0$). Thus WTP rises up to $\theta_S = \omega/G$ and falls after it. Evaluating either branch at $\theta_S = \omega/G$ and simplifying gives (23); this peak value is strictly positive precisely because the experiment is informative, $q_G + q_B > 1$. Informativeness enters only here: it guarantees that the entrepreneur at the action threshold values the experiment at all. $\square$

Proof of Proposition 1.2. Each entrepreneur's willingness to pay peaks at the entrepreneur's own action threshold, $\theta_S = \omega/G$, where the launch return (20) is zero (Lemma 3). Across entrepreneurs, substitute the conditional mean (19) into the launch return:

$$R(\theta_S) = \theta_S G - (a + k\theta_S) = (G-k)\,\theta_S - a, \qquad \frac{dR}{d\theta_S} = G - k.$$

The belief at which the population reaches the action threshold solves $R(\theta_S) = 0$,

$$\theta_S^* = \frac{a}{G-k}, \qquad \frac{d\theta_S^*}{dk} = \frac{a}{(G-k)^2} > 0. \tag{24}$$

Demand itself must be traced along the slice. Substituting $\omega = a + k\theta_S$ into (18) with $q_G = q_B = \bar{q}$ makes each branch a quadratic in the prior with common curvature $-2k(2\bar{q}-1) < 0$. The entrepreneur's willingness to pay is the branch consistent with the default, which is the smaller of the two, since $\mathrm{WTP}_p - \mathrm{WTP}_c = R$: the pivot branch applies below θ_S^* and the commit branch above it. As the minimum of two concave functions, willingness to pay is concave along the slice and single-peaked. Write θ_c and θ_p for the vertices of the two branches,

$$\theta_c = \frac{\bar{q}(k-a) - (1-\bar{q})(G-a)}{2k(2\bar{q}-1)}, \qquad \theta_p = \frac{\bar{q}(G-a) - (1-\bar{q})(k-a)}{2k(2\bar{q}-1)}, \qquad \theta_p - \theta_c = \frac{G-k}{2k(2\bar{q}-1)} > 0.$$

The peak of demand is θ_S^* clamped to the interval $[\theta_c, \theta_p]$. It sits at the population threshold θ_S^* when $\theta_c \le \theta_S^* \le \theta_p$; at θ_c, inside the commit-default region, when $\theta_S^* < \theta_c$; and at θ_p, inside the pivot-default region, when $\theta_S^* > \theta_p$. For small k the interval is wide (as $k \to 0$ it covers the belief range, recovering the fixed-option tent with peak a/G), so the peak tracks θ_S^* and rises with the co-movement by (24). For larger k the peak can detach from the threshold, but it stays interior: both vertices converge to $(G-a)/(2G)$ as $k \uparrow G$, so the peak does not leave the belief range.[11]

Consider finally the limiting alignment $k = G$ with $a = 0$, so that $\omega = G\theta_S$ exactly. Then $R(\theta_S) \equiv 0$: *every* entrepreneur sits at the entrepreneur's own action threshold, and by (23) willingness to pay is

$$\mathrm{WTP} = \frac{G\theta_S(G - G\theta_S)}{G}(q_G + q_B - 1) = G\,\theta_S(1-\theta_S)(q_G + q_B - 1) > 0$$

for every $\theta_S \in (0,1)$. No entrepreneur is confident enough, in either direction, to skip the experiment. The concentration of demand at a fixed moderate belief is the $k = 0$ slice of this surface; positive co-movement moves the peak, and in the alignment limit spreads positive demand across the entire belief distribution. □

A.4 When to Experiment (Propositions 2.1 and 2.2)

Proof of Proposition 2.1. (i) Opposite effects of the outside option. By Proposition 1.1, $\partial V_c/\partial\omega > 0$ and $\partial V_p/\partial\omega < 0$. By the experimentation rule (17), raising ω therefore moves

[11] A numerical illustration of detachment: $G = 10$, $a = \frac{1}{2}$, $k = 5$, $\bar{q} = 0.9$ gives $\theta_S^* = 0.10$, while demand along the slice peaks at $\theta_c = 0.3875$, with value 1.05 there against 0.72 at θ_S^*.

commit-default entrepreneurs *toward* experimenting (their V_c rises across zero) and pivot-default entrepreneurs *away* from it (their V_p falls across zero); it also reclassifies entrepreneurs across the boundary (11). The net effect on the population experimentation rate is the sum of two opposing flows and is therefore ambiguous in sign and generically non-monotone.

(ii) A minimum-accuracy floor. An experiment is worth running only if its verdict can change the action; otherwise the entrepreneur pays c and takes the default anyway. For a commit-default entrepreneur the only action-changing verdict is a negative result, which must be strong enough to move the posterior below the decision threshold, $\theta_{(-)}G < \omega$. Writing $\theta_{(-)} = \theta_S(1-q_G)/\big[\theta_S(1-q_G)+(1-\theta_S)q_B\big]$ and clearing the (positive) denominator,

$$\theta_S(1-q_G)\,G < \omega\big[\theta_S(1-q_G)+(1-\theta_S)q_B\big] \iff \theta_S(1-q_G)(G-\omega) < (1-\theta_S)q_B\,\omega,$$

which rearranges to the accuracy floor

$$\frac{q_B}{1-q_G} > \frac{\theta_S}{1-\theta_S}\,\frac{G-\omega}{\omega}. \tag{25}$$

Comparing with (18), condition (25) is exactly $\mathrm{WTP}_c > 0$: an experiment below the floor has nonpositive gross value and is never run, however small its cost. The pivot-default case is symmetric. There the action-changing verdict is a positive result, which must carry the posterior above the threshold, $\theta_{(+)}G > \omega$; the same algebra yields $q_G/(1-q_B) > \big[(1-\theta_S)/\theta_S\big]\big[\omega/(G-\omega)\big]$, which is exactly $\mathrm{WTP}_p > 0$. Either floor depends on the prior odds and the stakes ratio, but not on c. The cost determines *how many* entrepreneurs above the floor actually pay for an experiment; the floor itself is about whether a verdict could move the entrepreneur at all. □

Proof of Proposition 2.2. Under a symmetric experiment the two branches of (18) can be written

$$\mathrm{WTP}_c = \tfrac{1}{2}\big[(2\bar{q}-1)(v_{\mathrm{mo}}+v_{\mathrm{cm}}) - R\big], \qquad \mathrm{WTP}_p = \tfrac{1}{2}\big[(2\bar{q}-1)(v_{\mathrm{mo}}+v_{\mathrm{cm}}) + R\big],$$

and the branch consistent with the default is the smaller of the two, so willingness to pay is $\tfrac{1}{2}\big[(2\bar{q}-1)(v_{\mathrm{mo}}+v_{\mathrm{cm}}) - |R|\big]$. Experimentation therefore requires

$$|R| \;<\; (2\bar{q}-1)\,(v_{\mathrm{mo}}+v_{\mathrm{cm}}) - 2c\,:$$

a small distance to one's own action threshold, relative to the total stakes. Now measure that distance across entrepreneurs along the conditional mean (19). From (24), $R(\theta_S) = (G-k)\theta_S - a$, so for a population whose beliefs range over an interval of width $\Delta\theta \equiv \theta_{\max} - \theta_{\min}$, the spread of launch returns is

$$\big|R(\theta_{\max}) - R(\theta_{\min})\big| = |G-k|\,\Delta\theta.$$

This spread is *decreasing* in k for $k < G$ and vanishes at $k = G$. Positive co-movement therefore compresses entrepreneurs toward their thresholds: because optimism and the outside option rise together, a more optimistic entrepreneur's higher perceived value $\theta_S G$ is matched by a higher bar ω, keeping the entrepreneur near indifference. A larger share of the population then lies in the experimenting band, and the aggregate experimentation rate rises, peaking as $k \uparrow G$. (For $k > G$ the spread grows again and, as Proposition 4.2 shows, the boundary ordering reverses.)

Experimentation concentrates on the *indifference diagonal*: the optimist with a strong outside option (commit-default, near the threshold) and the pessimist with a weak one (pivot-default, near the threshold), which are exactly the combinations that positive co-movement populates. The deep-commit and deep-pivot corners, the optimist with a weak outside option and the pessimist with a strong one, lie far from any threshold, fail the floor (25), and are populated by negative co-movement. The confident entrepreneur with a strong outside option is therefore among the most likely buyers of an experiment, not among those who skip it. $\square$

A.5 How Demanding an Experiment (Propositions 3.1 and 3.2)

Here the experiment is symmetric, $q_G = q_B = \bar{q}$, the midpoint of the focus frontier of Section A.6, and $\bar{q} \in (\frac{1}{2}, 1)$ is its accuracy. Specializing the value functions (14)–(15) to $q_G = q_B = \bar{q}$ and writing them with the two stakes (16),

$$V_c(\bar{q}) = \bar{q}\, v_{\mathrm{cm}} - (1-\bar{q})\, v_{\mathrm{mo}} - c, \qquad V_p(\bar{q}) = \bar{q}\, v_{\mathrm{mo}} - (1-\bar{q})\, v_{\mathrm{cm}} - c, \tag{26}$$

each linear in $\bar{q}$. Both collapse to $V(\bar{q}) = (v_{\mathrm{mo}} + v_{\mathrm{cm}})\, \bar{q} - \max\{v_{\mathrm{mo}}, v_{\mathrm{cm}}\} - c$, so a perfect experiment ($\bar{q} = 1$) is worth the smaller of the two stakes, the residual regret the default leaves exposed, net of cost.

Proof of Proposition 3.1. Differentiating either branch of (26) in $\bar{q}$ gives the same slope, whatever the sign of the launch return (20):

$$\frac{\partial V}{\partial \bar{q}} = v_{\mathrm{mo}} + v_{\mathrm{cm}} = \theta_S(G-\omega) + (1-\theta_S)\omega. \tag{27}$$

Differentiating once more in the prior,

$$\frac{\partial^2 V}{\partial \bar{q}\, \partial \theta_S} = (G-\omega) - \omega = G - 2\omega, \tag{28}$$

positive iff $\omega < G/2$. With a common outside option ω_0 the sign $G - 2\omega_0$ is shared by every entrepreneur and is positive when $\omega_0 < G/2$ (in particular at $\omega_0 = 0$); with heterogeneous outside options it is entrepreneur-specific, and negative for any entrepreneur with $\omega > G/2$, for whom the costly mistake is the larger stake. $\square$

Proof of Proposition 3.2. Trace the marginal value of accuracy (27) across entrepreneurs along the conditional mean (19), $\omega = a + k\theta_S$:

$$\gamma(\theta_S) \;\equiv\; \big(v_{\mathrm{mo}}+v_{\mathrm{cm}}\big)\big|_{\omega=a+k\theta_S} = \theta_S(G-a-k\theta_S)+(1-\theta_S)(a+k\theta_S) = a+(G-2a+k)\,\theta_S-2k\,\theta_S^2. \tag{29}$$

Hence $\gamma''(\theta_S) = -4k < 0$ for $k > 0$: the cross-sectional marginal value of accuracy is strictly concave in optimism, with interior maximizer

$$\theta_S^Q = \frac{G - 2a + k}{4k}. \tag{30}$$

The condition $k > G - 2a$ is equivalent to $\theta_S^Q < \frac{1}{2}$, so γ is decreasing on $[\frac{1}{2}, 1]$; and

$$\gamma(1) - \gamma(0) = (G - a - k) - a = G - 2a - k \ < \ 0 \tag{31}$$

under the same condition, so the most optimistic entrepreneur's marginal value lies strictly below the most pessimistic entrepreneur's. With a common outside option ω_0, $\gamma(\theta_S) = \theta_S(G - \omega_0) + (1 - \theta_S)\omega_0$ is affine in θ_S, so $\gamma'' = 0$, γ is monotone, and no interior peak exists: the hump cannot arise.[12] □

The mechanism is visible in the two stakes taken separately. Along the conditional mean,

$$v_{\text{mo}}(\theta_S) = (G - a)\theta_S - k\theta_S^2, \qquad v_{\text{cm}}(\theta_S) = a + (k - a)\theta_S - k\theta_S^2,$$

are each concave (curvature $-2k$), and (29) is their sum. The missed opportunity is the stake an optimist loads onto, yet the same optimism inflates the outside option and shrinks the upside $G - \omega$ inside it; past $\theta_S = (G - a)/(2k)$ the shrinking upside dominates the rising weight and v_{mo} falls, while v_{cm} stays muted because a high belief downweights the bad state. The most confident entrepreneurs therefore have the least of either loss to insure, and demand the least accuracy. The reversal threshold $k > G - 2a$ collapses to $k > G$ as $a \to 0$, the same co-movement at which focus reverses (Proposition 4.2): accuracy turns on the *sum* $v_{\text{mo}} + v_{\text{cm}}$, focus on the *difference* $v_{\text{mo}} - v_{\text{cm}} = R$, and the same co-movement bends both.

A.6 Where to Aim the Experiment: Focus (Propositions 4.1 and 4.2)

Now let the entrepreneur choose the experiment's *shape* along an accuracy frontier

$$q_G + q_B \leq Q, \qquad 1 < Q < 2, \tag{32}$$

trading off the true positive rate against the true negative rate at a fixed total accuracy. The *low-bar* experiment puts all of the budget into the true positive rate, $\{q_G, q_B\} = \{1, Q - 1\}$ (no false negatives, so a negative result is fully credible); the *high-bar* experiment does the reverse, $\{Q - 1, 1\}$ (no false positives, so a positive result is fully credible).

Proof of Proposition 4.1. (i) Focus is selected by the default. A commit-default entrepreneur runs the experiment to vet a possible negative result, and chooses the focus that maximizes its value over the frontier. Substituting the binding frontier $q_B = Q - q_G$ into V_c from (14)

[12]Admissibility: $\omega = a + k\theta_S \in (0, G]$ across $\theta_S \in [0, 1]$ requires $a \in (0, G)$ and $a + k \leq G$, so the reversal regime $k \in (G - 2a,\, G - a]$ is nonempty for any $a > 0$. As $k \uparrow G - a$ the most optimistic entrepreneur's outside option approaches the whole prize, $\omega(1) \to G$, and the missed-opportunity stake $v_{\text{mo}}(1) = G - \omega(1) \to 0$.

and differentiating,

$$V_c = (1-\theta_S)(Q-q_G)\,\omega-\theta_S(1-q_G)(G-\omega)-c, \qquad \frac{\partial V_c}{\partial q_G}\bigg|_{q_B=Q-q_G} = \theta_S G-\omega = v_{\mathrm{mo}}-v_{\mathrm{cm}} = R(\theta_S).$$

For a commit-default entrepreneur $\theta_S G - \omega > 0$, so V_c increases in q_G and the true positive rate is driven to its maximum: the *low-bar* corner. Symmetrically, substituting $q_G = Q - q_B$ into V_p from (15),

$$\frac{\partial V_p}{\partial q_B}\bigg|_{q_G=Q-q_B} = \omega - \theta_S G = v_{\mathrm{cm}} - v_{\mathrm{mo}} = -R(\theta_S),$$

which is positive for a pivot-default entrepreneur, who therefore drives the true negative rate to its maximum: the *high-bar* corner. The dividing line between the two corners is the default condition $\theta_S G = \omega$ in (11), a comparison of perceived value with the outside option, *not* a threshold in belief alone. Focus is thus selected by the *sign* of the launch return $R = v_{\mathrm{mo}} - v_{\mathrm{cm}}$, the difference of the two stakes, just as the value of accuracy (Section A.5) is their *sum*.

(ii) Two corollaries the boundary delivers. Because focus turns on $\theta_S G \gtrless \omega$, two entrepreneurs with identical prior θ_S and identical upside G but outside options on opposite sides of $\theta_S G$ optimally build *opposite* experiments. And an entrepreneur hit by an exogenous shock to ω that crosses $\theta_S G$ *reverses* focus with no change in belief whatever.

(iii) Only an intermediate band of entrepreneurs experiments. Because $c > 0$, an entrepreneur runs an experiment only if the value (17) is positive. On the commit-default side, V_c is strictly decreasing in θ_S (its derivative is $-[\,\omega q_B + (1 - q_G)(G - \omega)\,] < 0$), so $V_c > 0$ holds only below an upper edge $\bar{\theta}$; on the pivot-default side, V_p is strictly increasing in θ_S (derivative $q_G(G - \omega) + (1 - q_B)\omega > 0$), so $V_p > 0$ holds only above a lower edge $\underline{\theta}$. Setting $V_c = 0$ and $V_p = 0$ and solving,

$$\bar{\theta} = \frac{q_B\,\omega - c}{q_B\,\omega + (1 - q_G)(G - \omega)}, \tag{33}$$

$$\underline{\theta} = \frac{c + (1 - q_B)\,\omega}{q_G(G - \omega) + (1 - q_B)\,\omega}. \tag{34}$$

Experimentation occurs only on the band $\theta_S \in (\underline{\theta}, \bar{\theta})$. Provided the band is nonempty, that is, provided the peak willingness to pay (23) exceeds the cost c, it straddles the focus boundary ω/G. Entrepreneurs above $\bar{\theta}$ are so optimistic, relative to their outside option, that they commit without experimenting; entrepreneurs below $\underline{\theta}$ are so pessimistic that they pivot without experimenting. Consequently *observed* focus, which is defined only for entrepreneurs who run an experiment, reflects an entrepreneur who is optimistic or pessimistic only in moderation: a low-bar experiment is the signature of a moderately optimistic entrepreneur near the threshold, not of the most optimistic entrepreneur, who runs no experiment at all. This is the selection half of the identification problem: focus locates an entrepreneur relative to the action threshold but cannot, on its own, separate a strong prior from a weak outside option. □

Proof of Proposition 4.2. (i) The map from optimism to focus reverses. Trace focus across entrepreneurs along the conditional mean (19). From (24) the launch return is $R(\theta_S) = (G-k)\theta_S - a$ with

$$\frac{dR}{d\theta_S} = G - k.$$

When $k > G$ this derivative is negative: as the prior rises, the launch return *falls*. A more optimistic entrepreneur is therefore more likely to be *pivot-default* and to build a *high-bar* experiment, the opposite of the fixed-ω pattern of Proposition 4.1, in which rising optimism carries an entrepreneur from high-bar to low-bar. Under $k > G$ the outside option rises with belief faster than perceived value does: the gain in $\theta_S G$ per unit of belief is G, while the bundled gain in ω is $k > G$, so the more optimistic entrepreneur is, on net, pulled further below the entrepreneur's own bar. The most optimistic entrepreneurs are the high-bar testers.[13]

(ii) The most optimistic pivot without experimenting. Consider an entrepreneur who, under $k > G$, has crossed into the pivot-default region; whether the entrepreneur experiments depends on the sign of V_p evaluated along the conditional mean. By the chain rule, $dV_p/d\theta_S|_{\text{c.m.}} = \partial V_p/\partial\theta_S + k\,\partial V_p/\partial\omega$, in which $\partial V_p/\partial\theta_S = q_G(G-\omega) + (1-q_B)\omega > 0$ but $\partial V_p/\partial\omega < 0$ by Proposition 1.1. By Lemma 2 there is a co-movement threshold above which this total derivative is negative; for such k, V_p decreases as the prior rises. A zero of V_p exists in any case: as ω approaches G along the conditional mean, $V_p \to -(1-\theta_S)(1-q_B)G - c < 0$, so whenever V_p is positive anywhere on the pivot-default side it crosses zero within the feasible range, at the *pivot belief* θ_{pivot} solving

$$V_p\big(\theta_S,\ a + k\theta_S\big) = 0. \tag{35}$$

For $\theta_S > \theta_{\text{pivot}}$ the value of experimenting is negative and the entrepreneur simply pivots, taking the now large, bundled outside option. Reading the regions in order of increasing optimism under $k > G$ therefore gives the sequence *low-bar experiment* (while still commit-default) $\to$ *high-bar experiment* (once pivot-default) $\to$ *no experiment* (beyond θ_{pivot}). This mirrors the standard result that the very optimistic do not experiment, but for the opposite reason: there they would commit regardless; here they would pivot regardless. $\square$

A.7 Endogenous Experimental Design (Lemma 1, Propositions 4.2 and 4.2)

The baseline subsections hold the experiment fixed and symmetric. Section 3 lets the entrepreneur design it, in two stages, stating one lemma, two corollaries, and one pair of propositions; this subsection proves them. Total accuracy is $Q = q_G + q_B$, with $Q \in (1,2)$ on the fixed frontier and $Q \leq 2$ once accuracy is chosen; informativeness is $Q - 1$, and the symmetric experiment of Section A.5 is the midpoint $Q = 2\bar{q}$. Write $m \equiv \min\{v_{\text{mo}}, v_{\text{cm}}\}$ for the smaller of the two stakes (16), the regret the default action leaves exposed.

[13] With $\omega = a + k\theta_S \in (0, G)$ and $k > G$, the belief support must be a strict subinterval of $[0,1]$, and the projection intercept a may then be negative; only the slope of the conditional mean over the observed support matters, as in the right panel of Figure 5.

Proof of Lemma 1. Along the binding frontier $q_B = Q - q_G$, the corner argument in the proof of Proposition 4.1 gives $\partial V_c/\partial q_G = \theta_S G - \omega = R$ (20), so a commit-default entrepreneur ($R > 0$) drives the true positive rate to its maximum (the low-bar corner $q_G = 1$, $q_B = Q - 1$) and a pivot-default entrepreneur ($R < 0$) the true negative rate (the high-bar corner $q_B = 1$, $q_G = Q - 1$). Substituting each corner into the stake-form value functions (14)–(15),

$$V_c\Big|_{q_G=1,\, q_B=Q-1} = (Q-1)\, v_{\mathrm{cm}} - c, \qquad V_p\Big|_{q_B=1,\, q_G=Q-1} = (Q-1)\, v_{\mathrm{mo}} - c.$$

The error term vanishes at each corner: focusing the experiment perfectly on one side removes that side's regret. The surviving stake is the smaller of the two, since the default's sign of R orders them, so both cases collapse to

$$V^* = (Q-1)\, m - c. \tag{36}$$

Writing the stakes out by (16) gives the two displayed values; the commit-default value contains ω but not G, and equals $-c$ at $\omega = 0$. □

Proof of Corollary 1. Differentiating (36) in total accuracy,

$$\frac{\partial V^*}{\partial Q} = \min\{v_{\mathrm{mo}}, v_{\mathrm{cm}}\} = m, \tag{37}$$

which is at most $\frac{1}{2}(v_{\mathrm{mo}} + v_{\mathrm{cm}})$, with equality exactly where the two stakes are equal, the action threshold. With the experiment held symmetric the marginal value of accuracy is the sum (27), and the two are in like units: the symmetric experiment carries $Q = 2\bar{q}$, so its value per unit of Q is $\frac{1}{2}(v_{\mathrm{mo}} + v_{\mathrm{cm}}) \geq m$. Holding ω fixed, v_{mo} rises and v_{cm} falls in the prior, so m rises up to the boundary $\theta_S = \omega/G$ and falls beyond it, peaking at $m = \omega(G - \omega)/G$, whereas the sum is affine in the prior and monotone. □

Proof of Proposition 4.2. Because the value of a binary experiment is linear in its error rates, a constant marginal cost of accuracy yields a corner choice (a perfect experiment or none); an interior accuracy requires a strictly convex cost, the simplest being the quadratic

$$C(Q) = c + \frac{\kappa}{2}\,(Q-1)^2, \qquad \kappa > 0, \tag{38}$$

with c the setup cost of running any experiment, as in the baseline, and κ the marginal cost of accuracy: the designed experiment costs its setup plus its accuracy. The entrepreneur maximizes $\Pi(Q) = (Q-1)m - \frac{\kappa}{2}(Q-1)^2 - c$. Writing $x \equiv Q - 1 \in (0, 1]$, Π is strictly concave ($\Pi'' = -\kappa < 0$) with first-order condition $m - \kappa x = 0$, so

$$Q^* = 1 + \frac{m}{\kappa} \tag{39}$$

when $m \leq \kappa$, and the corner $Q^* = 2$ (the perfect experiment $q_G = q_B = 1$) when $m > \kappa$. Substituting $x^* = m/\kappa$,

$$\Pi(Q^*) = \frac{m^2}{2\kappa} - c. \tag{40}$$

Substituting the surviving stake for each default into (39) gives the two displayed accuracy formulas; for a commit-default entrepreneur the survivor is the costly mistake, $m = v_{\rm cm} = (1-\theta_S)\omega$, so

$$\frac{\partial Q^*}{\partial \omega} = \frac{1}{\kappa}\frac{\partial v_{\rm cm}}{\partial \omega} = \frac{1-\theta_S}{\kappa} > 0, \tag{41}$$

and for a pivot-default entrepreneur the survivor is the missed opportunity, $m = v_{\rm mo} = \theta_S(G-\omega)$, giving $\partial Q^*/\partial\omega = -\theta_S/\kappa < 0$. For a commit-default entrepreneur $Q^* - 1 = (1-\theta_S)\omega/\kappa$ contains no G, and $\ln(Q^*-1) = \ln(1-\theta_S) + \ln\omega - \ln\kappa$, so $d\ln(Q^*-1)/d\ln\omega = 1$. The interior value (40) is positive exactly when $m > \sqrt{2\kappa c}$; at the cap the value $m - \kappa/2 - c$ is positive only if $m > \kappa/2 + c \geq \sqrt{2\kappa c}$, the last step by the arithmetic–geometric mean inequality, so the floor is necessary in both regimes and exact at an interior optimum. Substituting $m = (1-\theta_S)\omega$ and $m = \theta_S(G-\omega)$ and rearranging gives the two displayed floors; at $\omega = 0$ the commit-default floor fails for every θ_S and every G. Finally, at the participation margin $m = \sqrt{2\kappa c}$ the interior choice is $Q^* = 1 + \sqrt{2c/\kappa}$, and $Q^* - 1 = m/\kappa$ rises with the stake, so no participant builds a weaker test; when $c > \kappa/2$ the margin lies in the capped region and every participant builds the perfect experiment. □

Proof of Proposition 4.2. Parts (i)–(iii) correspond to the three regimes below; the displayed thresholds and peak locations are the content deferred from the main-text statement. The two-channel form is Lemma 2 applied to the values of Lemma 1, divided by κ: on the commit side $\partial Q^*/\partial\theta_S = -\omega/\kappa$ and $\partial Q^*/\partial\omega = (1-\theta_S)/\kappa$, so $dQ^*/d\theta_S|_{\rm c.m.} = [k(1-\theta_S) - \omega]/\kappa$; on the pivot side the partials are $(G-\omega)/\kappa$ and $-\theta_S/\kappa$, giving $[(G-\omega) - k\theta_S]/\kappa$. Substituting $\omega = a + k\theta_S$, the two stakes become the concave quadratics

$$v_{\rm mo}(\theta_S) = (G-a)\theta_S - k\theta_S^2, \qquad v_{\rm cm}(\theta_S) = a + (k-a)\theta_S - k\theta_S^2,$$

each with curvature $-2k$ (they appear in the proof of Proposition 3.2), and the channel expressions become $[(k-a) - 2k\theta_S]/\kappa$ and $[(G-a) - 2k\theta_S]/\kappa$, vanishing at $(k-a)/(2k)$ and $(G-a)/(2k)$: the vertices are the beliefs at which the outside-option channel exactly offsets the direct one. The difference of the stakes is the launch return, $v_{\rm mo} - v_{\rm cm} = R(\theta_S) = (G-k)\theta_S - a$, so the two parabolas cross exactly at $\theta_S^* = a/(G-k)$, the minimum is the surviving stake of the default, and $Q^* = 1 + \min\{v_{\rm mo}, v_{\rm cm}\}/\kappa$ is concave and single-peaked as the affine image of a minimum of concave functions.

(i) At $k = 0$ the commit-side derivative is $-\omega/\kappa < 0$ and the pivot-side derivative $(G-\omega)/\kappa > 0$, with the peak at the threshold where the stakes are equal. Along the cross-section the commit-side derivative is positive iff $k(1-\theta_S) > \omega$, i.e. iff $\theta_S < (k-a)/(2k)$; commit-default entrepreneurs satisfy $\theta_S > \theta_S^*$, so the two conditions overlap iff $\theta_S^* < (k-a)/(2k)$, which cross-multiplies to $k(G-k) > a(G+k)$. The condition $k > \omega/(1-\theta_S)$ is the critical co-movement (22) for $X = v_{\rm cm}$, since $-(\partial v_{\rm cm}/\partial\theta_S)/(\partial v_{\rm cm}/\partial\omega) = \omega/(1-\theta_S)$.

(ii) The peak sits at the crossing if and only if the left branch is still rising and the right branch already falling there, $v'_{\rm mo}(\theta_S^*) \geq 0 \geq v'_{\rm cm}(\theta_S^*)$. Computing,

$$v'_{\rm mo}(\theta_S^*) = \frac{G(G-k) - a(G+k)}{G-k}, \qquad v'_{\rm cm}(\theta_S^*) = \frac{k(G-k) - a(G+k)}{G-k},$$

so the two conditions are the displayed sandwich, with mutually exclusive failures for $k < G$. If the first fails, the peak is the commit-side balance belief $(k-a)/(2k) > \theta_S^*$. The second fails iff $a(G+k) > G(G-k)$, i.e. iff $k > \bar{k} = G(G-a)/(G+a)$; the peak is then the pivot-side balance belief $(G-a)/(2k) < \theta_S^*$. The threshold is admissible: $\bar{k} < G-a$ for every $a \in (0,G)$, so the regime is nonempty under the feasibility bound $a+k \leq G$. In the alignment limit $\omega = G\theta_S$, both stakes equal $G\theta_S(1-\theta_S)$.

(iii) For $k > G$ the launch return falls in the prior, so the population reads commit-default first and pivot-default second (Proposition 4.2), and the same minimum applies with the branches swapped: the low-bar branch v_{cm} survives below θ_S^* and the high-bar branch v_{mo} above it. Rigor is continuous at the switch, since $R(\theta_S^*) = 0$ makes the stakes equal there, and its slope drops by $(k-G)/\kappa$, since the branch slopes differ by $R'(\theta_S) = G-k$ at every belief. Writing $a = -|a|$ (feasibility of $\omega \in (0,G)$ under $k > G$ requires a negative intercept, with the belief support inside $(|a|/k,\ (G+|a|)/k)$), the vertex comparisons

$$\frac{G-a}{2k} - \theta_S^* = \frac{G(G-k)+|a|(G+k)}{2k(G-k)}, \qquad \frac{k-a}{2k} - \theta_S^* = \frac{k(G-k)+|a|(G+k)}{2k(G-k)}$$

give the three regimes: for $|a|(G+k) < G(k-G)$ both vertices exceed θ_S^* and the peak is the high-bar vertex $(G-a)/(2k)$; for $G(k-G) < |a|(G+k) < k(k-G)$ the crossing lies between the vertices and the peak is the switch itself; and for $|a|(G+k) > k(k-G)$ both vertices fall short of θ_S^* and the peak is the low-bar vertex $(k-a)/(2k)$. Beyond the surviving vertex, high-bar rigor declines in optimism, and $v_{\mathrm{mo}} \to 0$ as $\omega \to G$ along the conditional mean, so the design value $v_{\mathrm{mo}}^2/(2\kappa) - c$ turns negative at the pivot belief (35), beyond which no experiment is built. $\square$

Proof of Corollary 3. The two rigor branches are $1 + v_{\mathrm{cm}}/\kappa$ and $1 + v_{\mathrm{mo}}/\kappa$, and $v_{\mathrm{mo}} - v_{\mathrm{cm}} = R(\theta_S) = (G-k)\theta_S - a$ at every belief, so their slopes differ by $R'(\theta_S)/\kappa = (G-k)/\kappa$ throughout. At the focus switch θ_S^* the two branches meet, since $R(\theta_S^*) = 0$, so chosen rigor is continuous there with a kink of magnitude $|G-k|/\kappa$, vanishing at $k = G$. Proposition 4.2 signs $G-k$ from the order of the focus regions; the kink gives its magnitude. The identity holds at any design stage, since the branches differ by the launch return before accuracy is chosen as well. $\square$

Across the two stages, focus is a *direction*, set by the sign of the launch return, and accuracy is a *magnitude*, set by the size of the surviving stake m. The outside option enters the design of the experiment only through that stake.