NBER WORKING PAPER SERIES

BELIEF DISPERSION AND ENTREPRENEURIAL ENTRY

Joshua S. Gans

Working Paper 35091
http://www.nber.org/papers/w35091

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

Thanks to Refine.ink, ChatGPT 5.4 and Claude Opus 4.6 for valuable research assistance. Thanks for Luca Gius and Scott Stern for valuable discussions. Responsibility for all errors remain my own. The views expressed herein are those of the author and do not necessarily reflect the views of the National Bureau of Economic Research.

Belief Dispersion and Entrepreneurial Entry
Joshua S. Gans
NBER Working Paper No. 35091
April 2026
JEL No. D81, D83, L13, L21, O36

ABSTRACT

When should a founder act on a strong belief about an opportunity, knowing that rivals assessing the same opportunity may hold very different views? This paper studies entry decisions when entrepreneurs hold heterogeneous beliefs about an opportunity's value and each founder knows only the range of views rivals might hold. In equilibrium, a founder enters only when their conviction exceeds a threshold set by anticipated rival optimism. The relationship between belief dispersion and entry is surprisingly rich: depending on the founder's conviction and the cost of entry, there may be no level of dispersion that supports entry, all levels may support it, or only a middle range may, so that an outside observer may see the most entry at intermediate levels of belief dispersion. When founders can delay, high dispersion that deters immediate entry need not prevent it altogether: the absence of rival action gradually reveals that competitors are less bullish than feared. Finally, not all conviction-building is equal. Validation that only the founder sees strengthens entry incentives fully, whereas validation visible to the whole market partly backfires by encouraging rivals. The paper formalises the intuition that entrepreneurial value comes not from optimism alone but from optimism that the founder anticipates rivals will not share, and derives predictions linking belief dispersion to entry patterns.

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

1 Introduction

When is an entrepreneur's conviction strong enough, relative to the range of beliefs that rivals in the same environment may hold, to justify entry? The entrepreneurial moment is not an isolated judgment about whether an opportunity looks good, but a relative judgment about whether it looks good enough given how others are likely to see it. Amazon's internal team in 2003 was weighing whether to build scalable cloud infrastructure, while IBM, HP and Dell remained committed to selling hardware. Reed Hastings in 2007 was betting that streaming would displace DVDs, while broadcasters believed bandwidth and rights obstacles would last for years. Neither founder needed to ask "is this opportunity valuable?" The harder question was whether their own reading of the opportunity was sufficiently different from the readings that well-resourced rivals were likely to reach. If every serious player sees the same promise, entry invites costly duplication; if almost nobody else sees the promise, the founder may act as a near-monopolist. Value accrues not to optimism per se but to optimism that the founder anticipates rivals will not share. This paper builds a strategic model of that insight.

The paper takes the stance that heterogeneous prior beliefs are a primitive feature of entrepreneurial environments and should be modelled as such. Two founders looking at the same technology, market or regulatory setting can rationally come to different views about its commercial potential, not because they have different information but because they hold different causal theories about how the world works. The *Bayesian Entrepreneurship* research program, synthesised in Agrawal et al. (2025), treats entrepreneurship as a sequence of decisions under uncertainty in which founders are systematically contrarian relative to the resource-holders they face. Gans et al. (2026) provide formal microfoundations for this heterogeneity through a theory-based view of disagreement: differences in which causal mechanisms an entrepreneur believes are active and in beliefs about the magnitude of those mechanisms generate rational, persistent divergence in prior means even absent any information asymmetry. The present paper takes the output of that program, a distribution of prior means over a common fundamental, as its primitive, and asks what strategic entry looks like when two such entrepreneurs each know their own prior mean but only the distribution of the rival's.

Within this setup, the natural parameter to vary is the environment's *variance*, which is related to how opaque the opportunity is to anyone who approaches it. Low-variance environments correspond to legible opportunities with well-understood unit economics and familiar business models, such as franchised retail formats or copycat mobile applications. Moderate-variance environments correspond to opportunities with imperfect but meaningful evidence: early customer traction, pilot outcomes, and regulatory milestones. High-variance environments correspond to moonshots where early evidence is extremely noisy or non-transferable, including unproven scientific breakthroughs and technologies whose bottlenecks appear only at scale. In a risk-neutral model with linear payoffs and heterogeneous priors taken as primitive, variance enters strategic decisions only by shaping the range of rival prior means that an entrepreneur regards as plausible. Variance becomes a measure of the *dispersion of rational disagreement*, not simply of payoff risk.

Informal reasoning can suggest that wider belief dispersion "helps the founder" by thinning the pool of optimistic rivals, or "hurts" by thickening it, but cannot resolve which force dominates or when. The reason is that the competitive threshold a founder must clear is itself an equilibrium object: it depends on the dispersion of rival beliefs, which simultaneously determine how many rivals enter and, in turn, set the threshold. Signing the effect of belief dispersion, therefore, requires solving for equilibrium. The paper's results follow a single logical arc. In the simultaneous-move baseline, there is a unique anonymous monotone cutoff equilibrium. Each entrepreneur's decision reduces to comparing its own prior mean to a threshold that depends on the anticipated range of rival beliefs: the level of conviction at which the marginal entrant's gain from entry exactly offsets the expected duplication cost created by rival entry. When the sunk cost of entry is low relative to the average of the monopoly and duplication returns, higher belief dispersion *helps* entry because a wider range of rival priors thickens the pessimistic tail and thins the optimistic one. When sunk costs are high, only the optimistic tail of rival priors is decision-relevant, so higher dispersion *hurts* entry by thickening that tail and raising the competition-adjusted threshold.

A dynamic extension of the baseline model allows entrepreneurs to enter irreversibly or wait. Silence becomes informative: each period without entry reveals that the rival's prior mean lies below the most recent cutoff. Immediate entry is weakly harder than in the static game, and strictly harder whenever waiting has strictly positive marginal option value, but *eventual* entry is easier. Any feasible, optimistic entrepreneur with strictly positive monopoly value eventually enters at every variance level. High variance that blocks immediate entry in the simultaneous game may, therefore, only postpone entry in the dynamic game, not forbid it.

These findings are restated as a theory of *contrarian surplus*: the gap between an entrepreneur's own prior mean and the competition-adjusted cutoff. Contrarianism in this model is relative, not absolute. A positive prior mean is not enough; what matters is the distance between that prior mean and the optimism tax levied by expected rival entry. The final substantive section separates *private validation*, which raises only the entrepreneur's own prior mean, from *public validation*, which also shifts rival beliefs and therefore lifts the cutoff. A unit of private validation raises contrarian surplus one-for-one. A unit of public validation raises it only by a fraction because part of the gain is dissipated through a higher cutoff. The fraction retained shrinks as the duplication wedge grows: in markets where consensus entry destroys most of the value, public legibility is especially expensive to the founder.

The paper contributes to two literatures. First, it develops a clean strategic foundation for the contrarian-conviction intuition that animates the Bayesian Entrepreneurship program (Agrawal et al., 2025; Gans et al., 2026) and the theory-based view of strategy (Felin and Zenger, 2009, 2017), and connects the model's predictions to the recent empirical finding of Gius (2025) that evaluator disagreement in venture competitions positively predicts startup success, controlling for mean quality. Second, it extends the timing logic of Bolton and Farrell (1990) to an environment with heterogeneous prior means rather than private cost types, and shows that in a benchmark timing game the upper variance bound of the static game becomes a bound on immediate entry, not on eventual entry.

The remainder of the paper proceeds as follows. Section 2 introduces four running real-world examples, the model primitives, and the independence-of-priors assumption that underlies the analysis. Section 3 characterises the anonymous cutoff equilibrium in the simultaneous-move benchmark. Section 4 derives closed-form results under a bounded-range family of prior means and characterises the variance levels consistent with entry by an optimistic entrepreneur. Section 5 develops the dynamic extension with irreversible entry and delay. Section 6 reinterprets the results through contrarian surplus, connects them to the empirical evidence in Gius (2025), and studies private versus public validation. Section 7 concludes. Proofs are collected in the Appendix.

2 Motivating Examples and Model Primitives

2.1 Four running examples

Four cases recur throughout the paper. Each illustrates a different configuration of the objects the model will isolate, and each is returned to as results accumulate.

Cloud infrastructure. In the early 2000s, Amazon's internal view was that scalable pay-as-you-go compute, delivered as a primary good rather than as a complement to hardware, was a viable standalone business. Hardware incumbents such as IBM, HP and Dell held causal theories in which enterprise information technology spending was overwhelmingly a hardware-plus-services story, and assigned close to zero weight to the cloud-infrastructure mechanism (Gans et al., 2026). This is a high-variance environment and a very high own prior for the founder. In the model's terms, the key point is not that high variance thins the upper tail of rival priors, but that Amazon's conviction sits exceptionally far into that upper tail relative to incumbents whose dominant theories ruled the cloud mechanism out.

Streaming video. By 2007, Reed Hastings (Netflix's founder) held a strong belief that streaming would displace DVDs within a decade, conditional on bandwidth and rights obstacles being overcome. Broadcasters and cable operators believed those obstacles would bind for much longer. The informational environment was public, but the causal theories were incompatible, and the dispersion of rational prior means across potential entrants was wide.

Franchised retail and copycat applications. Low-variance environments with well-understood unit economics, legible business models and extensive prior data. Almost any entrant can estimate a plausible return; what scares a potential entrant is not whether the opportunity is real but how many other founders will reach a similar conclusion and compete rents away.

Moonshots that fail. Theranos, Segway and Webvan were contrarian bets that turned out wrong. Founders held high own priors in environments of extreme variance, and the ventures failed because the underlying causal theories did not hold up. These examples matter because they discipline the paper's message: the model does not claim that contrarian conviction is valuable because it is contrarian. The claim is that, conditional on the founder's own prior being high enough to cover the monopoly breakeven, contrarianism relative to the expected distribution of rival priors is what generates surplus.

Each example tags differently along four dimensions: where the founder's own prior sat, what the dispersion of rival priors was, whether entry was immediate or delayed, and what monopoly value looked like conditional on the theory being right.

2.2 Opportunity, players and priors

There are two risk-neutral entrepreneurs, indexed by $i \in \{1, 2\}$.[1] The commercial opportunity yields a common payoff shifter $\theta \in \mathbb{R}$, interpreted as market potential. Entrepreneur i holds a subjective prior

$$\theta \sim N(m_i, v),$$

where $m_i \in \mathbb{R}$ is entrepreneur i's prior mean and $v > 0$ is a common prior variance. The prior mean is private information. The variance v is publicly known and is the same for both entrepreneurs.

Two modelling choices deserve comment. First, the heterogeneity sits in the prior mean, not in the variance. This matches the empirical content of the theory-based view: two founders facing the same opportunity can rationally disagree about which causal mechanisms are active and how large they are (the mean of θ) while sharing a common view of how predictable the mechanisms they accept actually are (the variance of θ). Two founders who both believe in the cloud-infrastructure mechanism can still disagree about how big the cloud market will be without disagreeing about how noisy the signals from each pilot customer are. Second, the common variance v is interpreted as a feature of the environment the founders both inhabit: opaque environments rationally admit a wider range of prior means than legible ones.

The key modelling object is not a common prior over θ, but a common *meta-belief* about the distribution of rival prior means. Let F_v denote the cumulative distribution function of prior means in an environment with variance v. The strategic predictions below require only that each entrepreneur believes the rival's prior mean is drawn from F_v. They do not require the entrepreneurs to agree on a common prior over θ itself.

2.3 The belief-dispersion family

Throughout the paper the family $\{F_v\}_{v>0}$ satisfies the following condition.

Assumption 1 (Belief-dispersion family)**.** For each $v > 0$, F_v is continuous and supported on a compact interval $[\underline{m}(v), \overline{m}(v)]$ with $\underline{m}(v) < \overline{m}(v)$. The support expands weakly with v and is symmetric around zero unless otherwise stated.

Assumption 1 captures three economically natural properties. First, compact support matches an entrepreneur's practical ability to bound the set of plausible rival views: even when a founder cannot fully describe the distribution of other founders' theories, the founder can typically name a plausible range, and the support of F_v encodes that range. Second, weak expansion of the support

[1] An $N > 2$ entrepreneur model is considered below.

with v encodes the idea that more uncertain environments admit wider rational disagreement: if the environment is tightly understood, prior means cluster, and if it is opaque, founders can rationally begin further apart. Third, symmetry around zero is an expositional convenience: because the primitives X, μ, and S already carry the positive payoff structure, the sign content of the results is not driven by skewness in the belief-dispersion family, and the symmetric parameterisation makes the comparative statics transparent. Section 4 uses a uniform family on $[-\kappa\sqrt{v}, \kappa\sqrt{v}]$ to derive closed forms.

2.4 The independence-of-priors assumption and its scope

An important modelling choice deserves to be made explicit before the strategic analysis begins. Each entrepreneur's own prior mean m_i conveys no information about the rival's prior mean m_j: conditional on the environment's variance v, the two prior means are drawn independently from F_v. Knowing its own m_i does not lead an entrepreneur to update its belief about m_j. The entrepreneur's expected distribution over m_j is F_v regardless of the entrepreneur's own level of optimism.

This assumption matches a world in which prior heterogeneity comes from idiosyncratic sources: different causal theories held by different founders, different tacit knowledge, different analogical reasoning, and different tastes for risk-taking in the construction of a mental model. Under that reading, a founder that happens to be very optimistic learns nothing about other founders' optimism from its own conviction, because its conviction is generated by an idiosyncratic internal theory and not by a common external signal. This is exactly the world to which the Thiel framing speaks (Thiel and Masters, 2014). His question "what important truth do very few people agree with you on?" presumes that the founder's conviction comes from a private, non-shared source, so that the strength of that conviction is not itself evidence about rival convictions. Independence of priors is therefore substantively, not just technically, the right modelling choice for the idiosyncratic-priors regime.

If, by contrast, heterogeneity comes from a common source (a public news event, a shared technology disclosure, or an industry report that everyone reads but interprets somewhat differently), then a high own prior should rationally shift the entrepreneur's belief about the rival's prior upward. In that case, m_i and m_j are positively correlated conditional on v, and the competition-adjusted cutoff would have to be computed against the conditional distribution $F_v(\cdot \mid m_i)$ rather than the unconditional F_v. The competitive threshold would then itself depend on the founder's own prior, and the linear best-response structure that delivers the cutoff equilibrium would no longer obtain in the same clean form.[2]

[2]The public-validation exercise in Section 6 deliberately takes a limited step toward the common-source case. It preserves the independence of the idiosyncratic component of prior heterogeneity while letting a public signal shift every founder's mean by the same amount. The shared component represents a correlated external signal; the residuals remain independent. The full common-source case, in which the correlation between m_i and m_j is endogenous to an underlying signal structure, is left as a natural extension (flagged in Section 7). The paper's results should therefore be read as characterising the idiosyncratic-priors regime cleanly, and as demonstrating via the validation extension that the logic carries over at least partially once a shared component is introduced.

2.5 The simultaneous entry game

Each entrepreneur simultaneously chooses either *Enter* or *Stay Out*. Entry requires sunk cost $S > 0$. If exactly one entrepreneur enters, the sole entrant receives

$$\theta + X - S,$$

where $X > 0$ is the incremental value of being the only entrant. If both entrepreneurs enter, each receives

$$\theta + \mu - S,$$

where $\mu < X$ captures competitive dissipation under duplication. An entrepreneur who stays out receives 0. Let $\Delta \equiv X - \mu > 0$ denote the duplication wedge, and let $M \equiv \frac{X+\mu}{2}$ denote the average of monopoly and duplication per-period returns.

For entrepreneur i, the expected payoff from entering depends on the prior mean m_i and on the probability that the rival also enters. A convenient back-of-envelope anchoring of Δ to the running examples is useful here. In the cloud-infrastructure example, Δ is plausibly small: the market grew fast enough in the mid-2000s that even simultaneous entry by two infrastructure providers would have left both profitable. In a copycat mobile-application example, Δ is large: head-to-head entry by two effectively identical offerings dissipates most of the available rents.

2.6 Anonymous monotone strategies

Because the two entrepreneurs are ex ante exchangeable in the meta-distribution of prior means, the natural benchmark is an *anonymous* strategy: a single decision rule used by either entrepreneur. This is not a claim that realised beliefs are identical, only that the mapping from a privately known prior mean to an action should not depend on an economically meaningless player label. A strategy that depends on the player label would reflect arbitrary coordination rather than any economic primitive. Here we define this equilibrium for both the baseline model (Sections 3 and 4) and the later dynamic extension (Section 5).

Definition 1 (Anonymous monotone equilibrium)**.** Fix $v > 0$.

Simultaneous-move game. An *anonymous monotone Bayesian Nash equilibrium* is a Bayesian Nash equilibrium in pure strategies for which there exists a function

$$\sigma_v : [\underline{m}(v), \overline{m}(v)] \to \{\text{Enter}, \text{Stay Out}\}$$

such that: (i) both entrepreneurs use the same strategy σ_v; (ii) beliefs about the rival's prior mean are given by the common meta-distribution F_v; and (iii) σ_v is weakly increasing in the entrepreneur's own prior mean, in the sense that if $m' > m$ and $\sigma_v(m) = \text{Enter}$, then $\sigma_v(m') = \text{Enter}$. Because the action set is binary, any such strategy is equivalent to a cutoff rule: there exists $c(v) \in \mathbb{R}$ such that $\sigma_v(m) = \text{Enter} \quad \Longleftrightarrow \quad m \geq c(v)$.

Dynamic timing game. (Applied in Section 5 below) An *anonymous monotone perfect Bayesian equilibrium* is a perfect Bayesian equilibrium in which, after every public history with no entry through date $t-1$, both entrepreneurs use the same continuation strategy, that continuation strategy is weakly increasing in the entrepreneur's own prior mean, and beliefs are updated from F_v by Bayes' rule whenever possible. Since the only non-terminal public histories are silent histories, such an equilibrium can be represented by a non-increasing sequence of cutoffs $\{c_t(v,\delta)\}_{t\geq 1}$ such that, after no entry through date $t-1$, entrepreneur i enters at date t if and only if $m_i \geq c_t(v,\delta)$.

The anonymity restriction is natural because the entrepreneurs are ex ante symmetric except for their privately observed prior means, so any label-dependent prediction is an artifact of arbitrary coordination. The monotonicity restriction is also economically natural. At every history, the expected payoff from immediate entry is strictly increasing in the entrepreneur's own prior mean, and the payoff differences across alternative entry dates satisfy a single-crossing property. A more optimistic type should therefore never be less willing to enter, or enter strictly later, than a less optimistic type facing the same beliefs about the rival. Focusing on anonymous monotone equilibria, therefore, isolates the comparative statics generated by heterogeneous priors and common variance while setting aside label-asymmetric equilibria that are not decision-relevant for an entrepreneur with no ex ante strategic role.

3 Baseline Model

The baseline model is the simultaneous entry game just described.

3.1 Variance as an indirect channel

The basic economics of variance can be stated immediately.

Proposition 1 (Variance matters only through rival entry beliefs)**.** *Fix entrepreneur i's prior mean at m_i and let $p \in [0,1]$ denote the entrepreneur's subjective probability that the rival enters. The expected payoff from entering is*

$$\Pi_E(m_i, p) = m_i + X - S - \Delta p. \tag{1}$$

Common variance v does not directly affect expected entry payoff. It matters only through the entrepreneur's belief about rival entry, which is itself generated by the distribution of rival prior means.

The proofs of all results are in Appendix A. Proposition 1 isolates the paper's conceptual pivot. Under risk neutrality and linear payoffs, the uncertainty a founder has about θ washes out of the expected payoff conditional on entry: the founder integrates over θ with expectation m_i regardless of how noisy the underlying belief is. What v actually does, in this model, is not introduce payoff risk but determine how broad the entrepreneur's beliefs are about the range of rival optimism. Variance is a *competition-dispersion* parameter.

A closely related point, worth stating explicitly because it protects against a common misreading, is that an entrepreneur's own optimism carries no implication for the expected optimism of the rival. Under the independence assumption of Section 2.4, a more optimistic entrepreneur does not rationally expect a more optimistic rival. Its subjective probability of rival entry is pinned down entirely by F_v and the rival's cutoff, and does not depend on its own prior. This is what makes variance a clean comparative static parameter rather than one that is entangled with own conviction, and it is the technical expression of the Thiel intuition that conviction is, by assumption, an idiosyncratic input. Return to the cloud-infrastructure example: if Amazon had believed that hardware incumbents were also drawn from a pool of potential cloud believers (a high rival-mean dispersion above its own), its entry calculation would have been different. The independence assumption is what allows Amazon's high own conviction to remain unaffected by worrying about whether IBM shared it.

3.2 The anonymous cutoff equilibrium

Suppose the rival uses a cutoff c. Then the rival enters if and only if $m_j \geq c$, and so the entrepreneur assigns probability $1 - F_v(c)$ to rival entry. Substituting this into (1) yields the expected payoff from entry for type m_i when the rival uses cutoff c:

$$\Pi_E(m_i; c, v) = m_i + X - S - \Delta\big(1 - F_v(c)\big). \tag{2}$$

Strict linearity in m_i gives a monotone best response and therefore a cutoff equilibrium.

Proposition 2 (Unique anonymous monotone cutoff equilibrium)**.** *For every $v > 0$, there exists a unique anonymous monotone Bayesian Nash equilibrium in cutoff strategies. There is a unique threshold $c(v) \in \mathbb{R}$ satisfying*

$$c(v) = S - X + \Delta\big(1 - F_v(c(v))\big). \tag{3}$$

In the anonymous equilibrium, entrepreneur i enters if and only if

$$m_i \geq c(v).$$

Three features of Proposition 2 are worth drawing out. First, the best response is a cutoff because expected entry payoff is strictly affine in m_i given any rival strategy. More conviction always strictly increases the expected net value of entering, and so types are linearly ordered. Second, the equilibrium cutoff is unique because the best-response map $B_v(c) \equiv S - X + \Delta(1 - F_v(c))$ is continuous and weakly decreasing in c, so $G_v(c) \equiv c - B_v(c)$ is continuous and strictly increasing and crosses zero exactly once. Third, the cutoff $c(v)$ has a direct economic interpretation. It is the *optimism tax* levied by expected rival entry: the prior mean at which the marginal entrant's gross expected return $m_i + X - S$ is exactly eroded by the expected duplication cost $\Delta(1 - F_v(c(v)))$.

A numerical vignette makes this concrete. Take $X = 10$, $\mu = 2$, $S = 7$ and $\kappa = 1$ (a parameterisation that will recur in Sections 4 and 6). Then $\Delta = 8$ and $M = 6$. If the rival's cutoff is c and

F_v is the uniform distribution on $[-\kappa\sqrt{v}, \kappa\sqrt{v}]$, as will be imposed below, the fixed-point equation $c = S - X + \Delta(1 - F_v(c))$ becomes $c = -3 + 8 \cdot (1 - (c + \sqrt{v})/(2\sqrt{v}))$. At $v = 1$ the interior solution is $c = 0.2$; at $v = 4$ it is $c = 1/3 \approx 0.333$. A founder with own prior $m = 1$ clears the cutoff in both cases, but the equilibrium rent falls from 0.8 to 2/3 as variance rises, consistent with Proposition 4 because here $S > M$. The following corollary records the payoff identity in general.

Corollary 1 (Equilibrium payoff of type m_i)**.** *In the anonymous cutoff equilibrium, entrepreneur i's equilibrium payoff is*

$$\Pi(m_i, v) = \max\{m_i - c(v), 0\}. \tag{4}$$

Equation (4) is the object the rest of the paper tracks. The cutoff $c(v)$ is the optimism tax created by expected competition, and the entrepreneur's equilibrium gain from entry is its prior mean net of that tax. An entrepreneur whose prior mean exceeds the tax enters, and the distance above the tax is the entrepreneur's equilibrium rent.

4 Closed-Form Analysis

Proposition 2 holds for any family satisfying Assumption 1, but sharper comparative statics require a specific belief-dispersion family. This section uses a uniform family whose support is proportional to the standard deviation of the environment.

4.1 The uniform belief-dispersion family

Assumption 2 (Uniform prior-mean range)**.** For each $v > 0$, prior means are uniformly distributed on

$$[-\kappa\sqrt{v}, \kappa\sqrt{v}],$$

where $\kappa > 0$ is a constant. Equivalently,

$$F_v(m) = \begin{cases} 0, & m < -\kappa\sqrt{v}, \\ \dfrac{m + \kappa\sqrt{v}}{2\kappa\sqrt{v}}, & -\kappa\sqrt{v} \le m \le \kappa\sqrt{v}, \\ 1, & m > \kappa\sqrt{v}. \end{cases}$$

Assumption 2 is a tractable reduced form that encodes the idea, discussed in Section 2, that more uncertain environments admit a wider range of reasonable prior views. The half-width $\kappa\sqrt{v}$ scales with the standard deviation of θ: opaque environments support wider rational disagreement. Two practical properties of the uniform family should be noted. First, bounded support is managerially plausible; decision makers may be unable to specify a full density of rival beliefs but can often articulate a plausible range. Second, the three regions of the equilibrium (pooling entry, no entry, and an interior cutoff) are analytically clean under the uniform specification. Across the running examples, a low-variance franchised-retail environment has small $\kappa\sqrt{v}$ and tightly clustered priors;

a high-variance early-cloud environment has large $\kappa\sqrt{v}$ and an enormous range of rational prior means.

4.2 The closed-form equilibrium

Define $a(v) \equiv \kappa\sqrt{v}$. Then we can show the following:

Proposition 3 (Closed-form anonymous equilibrium under Assumption 2)**.** *Under Assumption 2, the unique anonymous cutoff equilibrium is as follows.*

(a) ***Pooling entry:*** *if*

$$S \leq \mu - a(v), \tag{5}$$

then $c(v) \leq -a(v)$ and every entrepreneur enters.

(b) ***No entry:*** *if*

$$S \geq X + a(v), \tag{6}$$

then $c(v) \geq a(v)$ and no entrepreneur enters.

(c) ***Interior cutoff:*** *if*

$$\mu - a(v) < S < X + a(v), \tag{7}$$

then the equilibrium cutoff lies inside the support and is

$$c(v) = \frac{2a(v)}{2a(v) + \Delta}(S - M). \tag{8}$$

Entrepreneur i enters if and only if $m_i \geq c(v)$.

Each of the three regions corresponds to an economically recognisable situation. In the pooling region, the sunk cost S is low enough that even a duopolist with the most pessimistic prior in the support ($m = -a(v)$) prefers entry to staying out. Copycat rollouts in low-cost, low-legibility niches fit this description: many founders already agree that the opportunity is broadly real and entry is cheap. In the no-entry region, the sunk cost is high enough that even the most optimistic rational type in the support cannot cover S as a monopolist; the environment simply does not support profitable entry. In the interior region (the generic case), the equilibrium cutoff is a convex combination of $S - M$ weighted by the range of rival priors. The weighting parameter $2a(v)/(2a(v) + \Delta)$ approaches one as $a(v)$ becomes large, so the cutoff approaches $S - M$ in high-variance environments; it approaches zero as $a(v)$ becomes small, so the cutoff approaches zero in low-variance environments.

4.3 How variance shifts the cutoff

The central comparative static of the static game can now be stated.

Proposition 4 (How common variance shifts the entry cutoff)**.** *Under Assumption 2, the derivative of the interior cutoff satisfies*

$$\operatorname{sign}\left(\frac{dc(v)}{dv}\right) = \operatorname{sign}(S - M). \tag{9}$$

Hence: (i) if $S < M$, higher variance lowers the cutoff and makes entry easier; (ii) if $S > M$, higher variance raises the cutoff and makes entry harder; (iii) if $S = M$, the cutoff is zero for all v in the interior region.

The intuition is a counting argument about the two tails of the rival-prior distribution. When $S < M$, entry is cheap relative to the average of monopoly and duplication returns. At the equilibrium cutoff the marginal entrepreneur is worried principally about the density of rivals sitting just above the cutoff: if a rival's prior mean is even slightly above $c(v)$, the rival enters, and duplication takes a bite out of the marginal entrant's rent. Widening the support of rival priors pulls some rivals down into the pessimistic tail, where they stay out, and it thins the mass of rivals sitting just above the cutoff. Expected crowding falls, so the cutoff falls. The franchised-retail example fits this case: unit economics are broadly understood, sunk costs are moderate relative to monopoly rents, and what the entrant fears is nearby copycats rather than idiosyncratic market failure. In such an environment, more dispersion in rational priors actually helps entry because it means some potential rivals will misjudge the opportunity downward and stay out.

When $S > M$, entry is expensive and only very optimistic types find it worthwhile. At the margin, what matters is the *upper* tail of rival priors, because these are the rivals who enter. Widening the support now enlarges the upper tail of rival optimism, raises the probability of crowded entry, and lifts the cutoff. A moonshot environment fits this case: only rare believers in the theory enter, and widening rational disagreement thickens the upper tail of fellow believers rather than the lower tail of abstainers.

At the knife-edge $S = M$, the two tail effects cancel and the cutoff is zero regardless of v. The closed-form entry probability in the interior region,

$$1 - F_v(c(v)) = \frac{1}{2} - \frac{S - M}{2a(v) + \Delta}, \tag{10}$$

makes the same point in probability language: the probability of crowded entry inherits its monotonicity in v directly from the sign of $S - M$.

4.4 Which variance levels support entry by an optimistic entrepreneur?

For any entrepreneur with prior mean $m \in \mathbb{R}$, define the *variance-entry correspondence*

$$\mathcal{V}(m) \equiv \{v > 0 : m \text{ is feasible under } F_v \text{ and } m \geq c(v)\}.$$

This is the set of variance levels for which an entrepreneur with prior mean m enters in the anonymous equilibrium.

Proposition 5 (Variance levels consistent with entry by an optimistic entrepreneur)**.** *Suppose Assumption 2 holds and fix $m > 0$.*

(a) If

$$m \leq S - X, \tag{11}$$

then $\mathcal{V}(m) = \varnothing$. Even monopoly entry has non-positive expected value.

(b) If $m > S - X$ and $S \leq M$, then

$$\mathcal{V}(m) = \left[\frac{m^2}{\kappa^2}, \infty\right). \tag{12}$$

Any feasible optimistic entrepreneur enters.

(c) If $m > S - X$, $S > M$, and

$$m \geq S - M, \tag{13}$$

then again

$$\mathcal{V}(m) = \left[\frac{m^2}{\kappa^2}, \infty\right). \tag{14}$$

(d) If

$$S - X < m < S - M, \tag{15}$$

then the variance-entry correspondence is the bounded interval

$$\mathcal{V}(m) = \left[\frac{m^2}{\kappa^2}, \left(\frac{m\Delta}{2\kappa(S - M - m)}\right)^2\right]. \tag{16}$$

Proposition 5 is the static game's central result: depending on the founder's own prior mean and the economic environment, the set of variance levels that support entry can be empty, a half-line, or a bounded interval. Each case corresponds to a distinct entrepreneurial situation. The general cutoff logic comes from Proposition 2; the bounded interval in part (d), and the associated intermediate-variance observer pattern, are specific to the bounded-support uniform family in Assumption 2. It also has an observer-side implication. An outside observer typically does not see m directly and instead observes whether entry occurs. For that observer, the relationship between variance and observed entry need not be monotone even when the equilibrium cutoff itself is monotone in v, because the feasibility of optimistic types and the competition-adjusted threshold move differently with variance.

Case (a) is the monopoly-infeasible case; see Figure 1(a). Even alone in the market, an entrepreneur with prior mean below $S - X$ does not expect to recover its sunk cost. No environment supports entry. The point is not that variance is at the wrong level; the point is that conviction is too low to justify sinking S even in the best possible competitive outcome. A moderate-belief entrepreneur in a high-cost moonshot environment like that surrounding Theranos or Webvan is in this case.

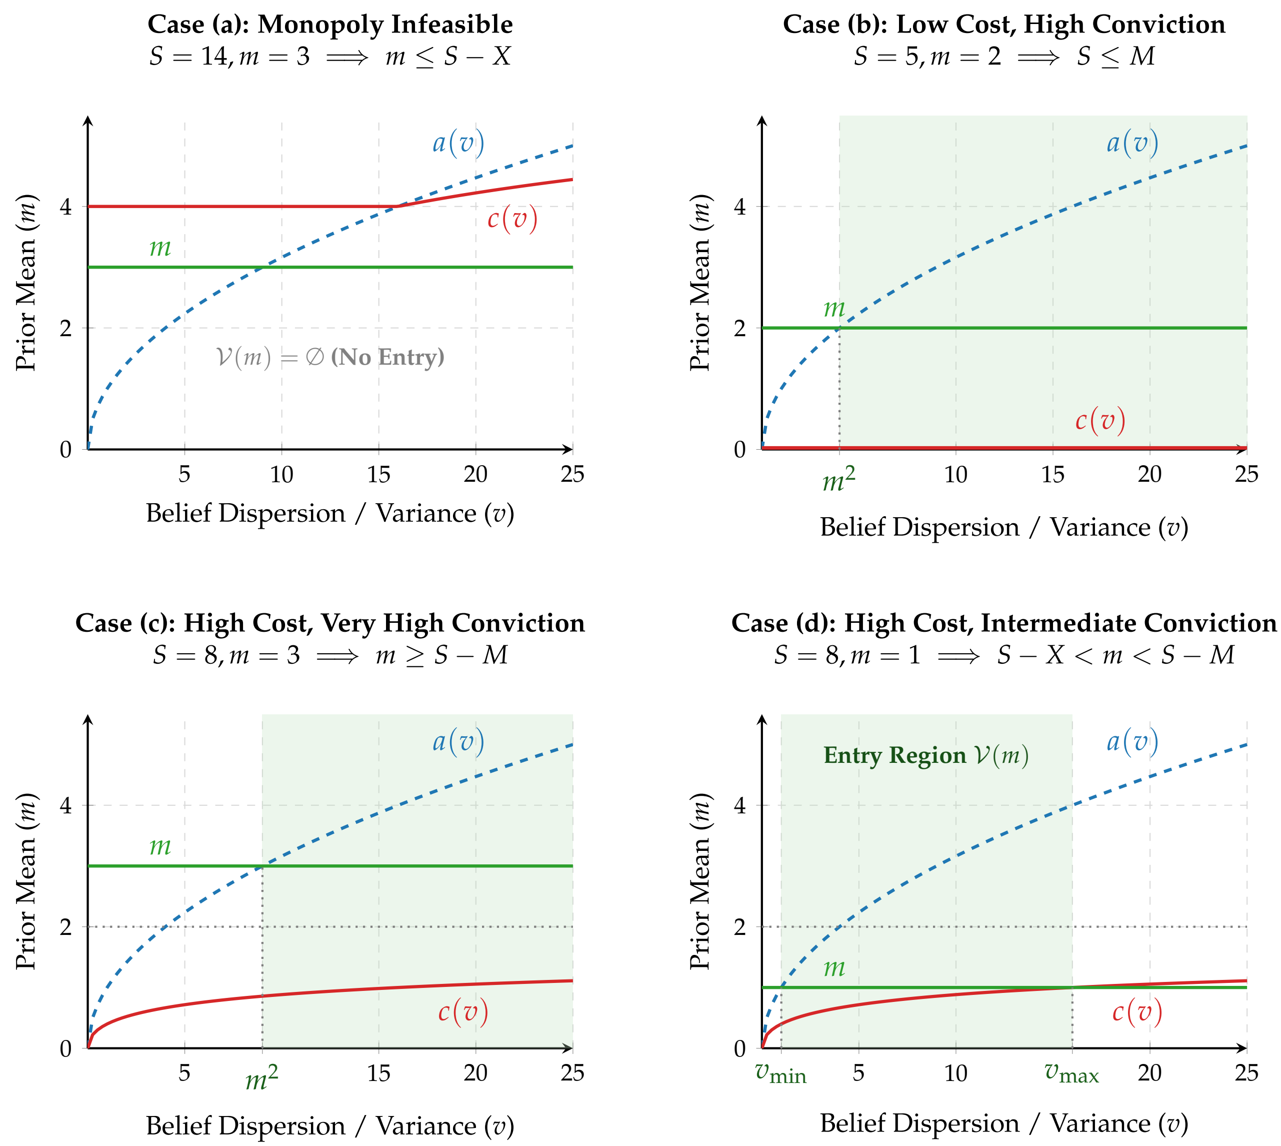

Figure 1: Entry sets, $\mathcal{V}(m)$, under the four cases of Proposition 5. The shaded region shows founder types who enter. The red curve is the optimism tax cutoff $c(v)$; the dashed blue curve is the feasibility bound $a(v) = \sqrt{v}$.

Case (b) is the low-cost, high-conviction case; see Figure 1(b). When entry is cheap ($S \le M$), the cutoff in the interior region is non-positive, so any optimistic entrepreneur with $m > 0$ clears it. As soon as m is feasible, that is, as soon as $v \ge m^2/\kappa^2$ so that m lies in the support of F_v, the entrepreneur enters. The half-line $\mathcal{V}(m) = [m^2/\kappa^2, \infty)$ describes an environment in which the founder's only concern is whether its own conviction is a type the environment rationalises at all.

Case (c) is the high-cost, very-high-conviction case; see Figure 1(c). When entry is expensive ($S > M$) but the founder's own conviction is very high ($m \ge S - M$), the cutoff can never rise above m because its limit as $v \to \infty$ is exactly $S - M$. Again $\mathcal{V}(m) = [m^2/\kappa^2, \infty)$. Reed Hastings in the streaming case, or Jeff Bezos in the cloud-infrastructure case, fit this description: the founder's own conviction is so much stronger than the likely conviction of rivals that no feasible variance level rules out entry.

Case (d), the bounded-interval case, is the most revealing; see Figure 1(d). When the founder's own conviction is intermediate ($S - X < m < S - M$), variance plays a two-sided role. The *lower* bound m^2/κ^2 is a pure feasibility constraint: if variance is too low, an entrepreneur with prior mean m does not belong to the environment's admissible range of priors. The *upper* bound reflects competition: when variance becomes too high, the range of rival prior means is so wide that the founder's own optimism no longer clears the cutoff, because the optimism tax has grown faster than the founder's conviction. Between these bounds, entry is supported. This is the only case in which an interior interval of variance values sustains entry. A founder holding a modestly optimistic but not extreme prior in an expensive, contrarian-dominated market is in case (d).

From an observer's perspective, case (d) delivers the sharpest comparative-static implication. If a population of potential entrants contains substantial mass in the intermediate-conviction region $S - X < m < S - M$, observed entry may be relatively low in very low-variance environments, higher in intermediate-variance environments, and lower again in very high-variance environments. At low variance, the support of priors is tight and strongly optimistic founders are often not feasible types under F_v. At high variance, optimism is feasible, but the optimistic tail of rival priors is also thicker, so the competition-adjusted cutoff rises, and immediate entry becomes harder. Entry is then most likely to be observed at intermediate variance, where optimism is feasible but not yet so widely shared that the optimism tax becomes too large. This is not a claim that $c(v)$ itself is non-monotone. Rather, it is a claim about the frequency with which entry is observed across environments when many potential entrants are of the intermediate-conviction type.[3]

This distinction also clarifies how variance should be measured empirically from an observer's perspective. In the model, v is not realised payoff volatility. It is a measure of the width of plausible prior means, or equivalently of how much rational disagreement the environment admits. Natural empirical proxies are, therefore, ex ante disagreement measures: the dispersion of investor,

[3]Section 5 (below) qualifies this observer-side prediction in a useful way. In the dynamic game, the upper bound in case (d) becomes a bound on immediate entry rather than on eventual entry. An observer using a short measurement window may therefore see the inverse-U pattern most clearly in contemporaneous entry rates. Over longer windows, some of the apparently missing entry in high-variance environments should reappear as delayed entry, because silence reveals that rivals are less optimistic than initially feared. High variance may therefore reduce observed *immediate* entry more strongly than it reduces observed *eventual* entry.

judge, or expert scores; the spread of forecasts of market size, adoption, or technical feasibility; disagreement in comparable-company valuations; heterogeneity in due-diligence assessments; or other measures of how differently informed observers assess the same opportunity before entry occurs. Such proxies are closer to the model's object than ex post outcome volatility, which mixes belief dispersion with realised shocks after the entrepreneurial decision has already been made.

The broader implication is that environments with the most observed entry need not be those with the highest realised uncertainty. They may instead be those with intermediate ex ante disagreement: environments in which sufficiently optimistic founders can exist, but in which optimism is not so widespread that expected rival entry drives the cutoff sharply upward.

4.5 Optimal variance is a boundary or flat-region object

The next proposition clarifies what an entrepreneur can (and cannot) gain by "choosing" the variance of the environment: under Assumption 2, payoffs are weakly monotone in v on the feasible set, so any optimum occurs at a boundary point (or on a flat region).

Proposition 6 (Optimal variance is a boundary or flat-region object)**.** *Fix $m \in \mathbb{R}$ and let $\mathcal{W} \subseteq \mathbb{R}_+$ be any non-empty closed feasible set of variance levels. Under Assumption 2, entrepreneur i's equilibrium payoff*

$$\Pi(m, v) = \max\{m - c(v), 0\}$$

is weakly monotone in v on $\mathcal{W} \cap \mathcal{V}(m)$. Consequently, if $\Pi(m, \cdot)$ attains a maximum on $\mathcal{W} \cap \mathcal{V}(m)$, then at least one maximiser is a boundary point of $\mathcal{W} \cap \mathcal{V}(m)$. If $\Pi(m, \cdot)$ is constant on $\mathcal{W} \cap \mathcal{V}(m)$, then every feasible variance level in that set is optimal.

Proposition 6 has a direct interpretation. Holding the entrepreneur's own prior mean fixed, variance affects the entrepreneur's payoff only through the cutoff $c(v)$, which is monotone in v under Assumption 2. The entrepreneur's problem of choosing an environment's variance is, therefore, not a smooth interior optimisation; it generates a weakly monotone payoff in v. Any interior maximiser would have to arise from a flat payoff region, not from a strict first-order condition. In particular, the static model does not predict that there is a universal moderate-variance "sweet spot" for payoff. Section 5 does not overturn this monotonicity logic; what it changes is the interpretation of the upper bound in Proposition 5(d). With timing, that upper bound becomes a bound on *immediate* entry rather than a definitive no-entry region for *eventual* entry.

4.6 More than two potential entrants

The two-entrepreneur benchmark isolates the threshold logic cleanly, but the same mechanism extends in a straightforward way to an industry with $N \geq 2$ ex ante identical potential entrants. Keep the information structure unchanged: entrepreneur i privately observes m_i, the common variance v is public, and each entrepreneur believes that the prior means of the other $N - 1$ potential entrants are independently drawn from F_v.

A convenient extension that preserves the linear structure of the baseline is to suppose that if $K \in \{1, \ldots, N\}$ entrepreneurs enter, each entrant receives

$$\theta + X - (K-1)\Delta - S,$$

where $\Delta > 0$ is the per-rival duplication loss. When $N = 2$, this reduces to the baseline specification because $X - \Delta = \mu$. More general congestion schedules in which per-entrant payoff is decreasing in the total number of entrants would preserve the cutoff logic below, but this pairwise-additive form keeps the algebra transparent.

Let

$$\mu_N \equiv X - (N-1)\Delta, \qquad M_N \equiv \frac{X + \mu_N}{2} = X - \frac{(N-1)\Delta}{2}.$$

The term μ_N is the per-entrant payoff increment when all N entrepreneurs enter, while M_N is the average of monopoly and full-crowding increments. Note that $M_2 = M$.

Suppose all rivals use a common cutoff c, and write

$$q(c, v) \equiv 1 - F_v(c)$$

for the probability that a given rival enters. Then the number of rival entrants is

$$B \sim \text{Binomial}(N-1, q(c, v)),$$

so the expected payoff from entry for type m_i is

$$\Pi_E^N(m_i; c, v) = m_i + X - S - \Delta \mathbb{E}[B] = m_i + X - S - (N-1)\Delta q(c, v).$$

As in the two-entrepreneur case, this payoff is strictly increasing in m_i, so the anonymous best response is again a cutoff rule. The anonymous monotone equilibrium cutoff $c_N(v)$ is therefore the unique solution to

$$c_N(v) = S - X + (N-1)\Delta\big[1 - F_v(c_N(v))\big].$$

The interpretation is unchanged. Common variance still has no direct effect on expected payoff. It matters only through the induced distribution of rival optimism. The only difference is quantitative: with more potential entrants, a given upper-tail probability of rival optimism is multiplied across more rivals, so the competition-adjusted cutoff rises.

Under Assumption 2, with $a(v) = \kappa\sqrt{v}$, the same three regions as in the two-entrepreneur case obtain. If

$$S \leq \mu_N - a(v),$$

all types enter. If

$$S \geq X + a(v),$$

no type enters. The no-entry condition is unchanged because it is pinned down by monopoly profitability of the highest type, whereas the pooling-entry condition becomes stricter as N rises because the lowest type must be willing to enter even when all other potential entrants also enter. In the interior region,

$$\mu_N - a(v) < S < X + a(v),$$

the cutoff is

$$c_N(v) = \frac{2a(v)}{2a(v) + (N-1)\Delta}(S - M_N).$$

Hence

$$\operatorname{sign}\left(\frac{dc_N(v)}{dv}\right) = \operatorname{sign}(S - M_N),$$

so the comparative static with respect to variance has exactly the same form as before, with M replaced by M_N.

The multi-entrant extension also sharpens the crowding interpretation of the variance-entry correspondence. For an optimistic entrepreneur with prior mean $m > 0$, the set of variance levels consistent with entry again can be empty, a half-line, or a bounded interval. In the analogue of the bounded-interval case, namely when

$$S - X < m < S - M_N,$$

the variance-entry correspondence becomes

$$V_N(m) = \left[\frac{m^2}{\kappa^2}, \left(\frac{m(N-1)\Delta}{2\kappa(S - M_N - m)}\right)^2\right].$$

Thus the lower bound remains a feasibility condition, while the upper bound reflects competition. As N rises, M_N falls and the upper bound shrinks, so the set of variance levels consistent with immediate entry by an intermediate optimistic type becomes narrower. Economically, the reason is simple: when there are more potential entrants, it becomes more likely that at least one rival is also sufficiently optimistic, so a founder needs more to justify entry.[4]

5 Dynamic Entry with Irreversibility and Delay

Entrepreneurial rivalry rarely resolves in a single, simultaneous move. A founder who is unsure whether a well-resourced rival will also enter can wait to see whether the rival moves first. Waiting is costly: monopoly rents are delayed through discounting, and may be preempted if the rival

[4] A possible extension of the Section 5 timing game to $N > 2$ entrants is conceptually similar but notationally less tidy. Silence is then informative about several rivals at once rather than about a single rival, so the relevant state is the truncated distribution of the set of remaining rivals. The same logic suggests that waiting should still raise the date-1 cutoff relative to the simultaneous benchmark, while informative delay should continue to restore eventual entry for sufficiently optimistic types with positive long-run monopoly surplus. The two-entrepreneur timing game is, therefore, best viewed as the cleanest case for isolating the option value of waiting, while the N-entrant simultaneous-move extension is the natural place to capture richer industry crowding.

enters first. Yet waiting is also valuable: if the rival has not entered by date t, Bayes' rule implies that the rival's prior mean lies below the most recent cutoff, and this information reduces the perceived risk of duplication. The section's objective is to keep all primitives of Section 2 and ask how the option to wait changes the variance-entry conclusions of Section 4. The model is close in spirit to Bolton and Farrell (1990), who study decentralisation, duplication and delay under private information, and extends the timing logic of that paper to a setting in which the private information is a prior mean rather than a private cost or value.

5.1 The infinite-horizon timing game

Time is discrete, $t = 1, 2, \ldots$, and payoffs are discounted by $\delta \in (0, 1)$. If no entry has yet occurred, each entrepreneur simultaneously chooses *Enter* or *Wait*. Entry is irreversible. If exactly one entrepreneur enters at date t, the entrant receives $\delta^{t-1}(\theta + X - S)$ and the rival receives 0. If both entrepreneurs enter at date t, each receives $\delta^{t-1}(\theta + \mu - S)$. If both wait, the game moves to $t + 1$. Once at least one entrepreneur enters, production takes place and the game ends. This is a deliberately sharp preemption benchmark: a sole entrant captures the opportunity and later stand-alone entry is ruled out. The purpose is to isolate the informational value of silence rather than to model richer industry evolution with capacity expansion, obsolescence, or post-entry sequential competition.

Public history consists only of whether previous dates were silent. As in Bolton and Farrell (1990), label-asymmetric timing equilibria can be constructed, but they are not the relevant benchmark when neither entrepreneur has an ex ante strategic label. The paper therefore focuses on the perfect-Bayesian timing analogue of Definition 1. Such an equilibrium is described by a non-increasing sequence of cutoffs $\{c_t(v, \delta)\}_{t \geq 1}$: after any history with no entry through date $t - 1$, entrepreneur i enters at date t if and only if $m_i \geq c_t(v, \delta)$. It is also useful to view a pure strategy in normal form as a planned entry date $a_i \in \mathbb{N} \cup \{\infty\}$, where $a_i = \infty$ means never enter. Anonymous monotone cutoff strategies correspond one-for-one to nondecreasing planned-date strategies in the entrepreneur's own prior mean. For notational convenience, define $c_0(v, \delta) \equiv \overline{m}(v)$, so that $F_v(c_0(v, \delta)) = 1$.

Before proceeding to the formal results, a short conceptual remark may help the reader. Silence in this model is not merely the absence of action; it is a public signal. If no entry has occurred through date $t - 1$ and the equilibrium cutoff at date $t - 1$ was c_{t-1}, then the rival's prior mean must lie below c_{t-1}. Bayesian updating then truncates F_v from above at c_{t-1}. Each period of silence peels another slice off the top of the rival-belief distribution, and the longer the silence, the more confident a surviving type becomes that the rival is not very optimistic.

5.2 Existence of an anonymous monotone timing equilibrium

We begin by showing the existence of an anonymous monotone equilibrium for this game.

Proposition 7 (Existence)**.** *Suppose Assumption 1 holds. Fix $v > 0$ and $\delta \in (0, 1)$. Then the infinite-horizon dynamic game admits at least one anonymous monotone perfect Bayesian equilibrium in pure cutoff strategies. Equivalently, there exists a non-increasing sequence $\{c_t(v,\delta)\}_{t\geq 1} \subseteq [\underline{m}(v), \overline{m}(v)]$ such that, after any history with no entry through date $t-1$, entrepreneur i enters at date t if and only if $m_i \geq c_t(v,\delta)$, with beliefs given by Bayes' rule on the equilibrium path.*

The proof in the Appendix has three steps. First, it represents pure strategies in the infinite-horizon game as planned entry dates, converting the timing game to a normal-form game in which each strategy chooses when first to attempt entry. Second, it truncates the set of planned dates to $\{1, \ldots, T, \infty\}$ and proves that the truncated anonymous best-response correspondence is non-empty, compact, convex-valued, and closed, so Kakutani's theorem yields an anonymous monotone equilibrium in each finite-truncation game. Third, it passes to the infinite-horizon limit via a selection argument that extracts a nondecreasing planned-date strategy from the truncation-indexed sequence and shows that the limiting strategy is a best response to itself in the untruncated game. Sequential rationality after on-path silent histories then follows because each continuation payoff from a feasible later date is the corresponding ex ante payoff multiplied by a fixed positive constant, so the ranking of continuation actions is unchanged.

The paper does not claim uniqueness of the infinite-horizon cutoff path. What it claims, and uses, is the existence of at least one anonymous monotone cutoff equilibrium, together with properties that hold in every such equilibrium. The timing extension should therefore be read as a benchmark partial baseline: the robust results are those that hold throughout the anonymous monotone class, not a unique canonical cutoff path.

5.3 Dynamic cutoff characterisation

If no entry has occurred through date $t-1$, Bayes' rule implies that the rival's prior mean is distributed according to the original F_v truncated from above at $c_{t-1}(v,\delta)$:

$$\Pr\big(m_j \leq m \mid \text{no entry through } t-1\big) = \frac{F_v(m)}{F_v(c_{t-1}(v,\delta))}, \qquad m \leq c_{t-1}(v,\delta). \tag{17}$$

Silence is informative: it reveals that the rival's prior mean lies below the most recent entry cutoff.

Proposition 8 (Dynamic cutoff characterisation)**.** *Let $\{c_t(v,\delta)\}_{t\geq 1}$ denote the cutoff sequence of any anonymous monotone perfect Bayesian equilibrium. At any on-path silent history through date $t-1$ with $F_v(c_{t-1}(v,\delta)) > 0$, the probability that the rival enters at date t is*

$$h_t(v,\delta) = 1 - \frac{F_v(c_t(v,\delta))}{F_v(c_{t-1}(v,\delta))}. \tag{18}$$

Let $\widetilde{V}_{t+1}(m; v,\delta)$ denote type m's continuation value at date $t+1$, conditional on no entry at date t. Then *equilibrium cutoff behaviour is characterised by*

$$m + X - S - \Delta h_t(v,\delta) \geq \delta\big(1 - h_t(v,\delta)\big)\widetilde{V}_{t+1}(m; v,\delta), \qquad m \geq c_t(v,\delta), \tag{19}$$

$$m + X - S - \Delta h_t(v,\delta) \leq \delta\big(1 - h_t(v,\delta)\big)\widetilde{V}_{t+1}(m; v,\delta), \qquad m < c_t(v,\delta). \tag{20}$$

If $\underline{m}(v) < c_t(v,\delta) < c_{t-1}(v,\delta)$, so that the cutoff is interior in the remaining support, then the marginal type is indifferent:

$$c_t(v,\delta) + X - S - \Delta h_t(v,\delta) = \delta\big(1 - h_t(v,\delta)\big)\widetilde{V}_{t+1}(c_t(v,\delta); v,\delta). \tag{21}$$

Hence more optimistic entrepreneurs enter weakly earlier.

Proposition 8 is the dynamic analogue of the simultaneous-move cutoff result. The first part records the hazard generated by silence. The second part says that, at every on-path silent history, immediate entry is compared with the value of waiting one more period. When the cutoff is interior, the threshold type is indifferent and the one-step recursive representation is exact; when the cutoff hits a boundary, only the one-sided cutoff inequalities are guaranteed.

Two features of this dynamic structure are worth emphasising. First, the sequence of hazards $\{h_t\}$ generated by silence is the new object relative to the static game. Where the static cutoff was pinned down by a single probability $1 - F_v(c)$ of rival entry, the dynamic cutoff is pinned down by the current hazard together with the continuation value of waiting. Second, silence is self-reinforcing: as t grows the rival's belief distribution is increasingly truncated, and the probability that the surviving rival is an extreme optimist falls. In any equilibrium, the sequence $\{c_t\}$ is non-increasing, and hazards typically decline over time.

Proposition 9 (Delay weakly raises the date-1 entry cutoff)**.** *Fix an anonymous monotone perfect Bayesian equilibrium of the dynamic game and let $c(v)$ denote the simultaneous-move cutoff from Proposition 2. Then*

$$c_1(v,\delta) \geq c(v). \tag{22}$$

The inequality is strict whenever the marginal date-1 type has strictly positive continuation value from waiting.

Proposition 9 says that immediate entry is weakly harder than in the static game, and strictly harder whenever waiting has strictly positive continuation value at the margin. The simultaneous-move benchmark is nested as the case in which waiting carries no marginal value, but Proposition 9 itself establishes only the weak comparison together with a one-sided strictness condition. In the timing game, a founder trades earlier entry against the information revealed by silence, and the marginal date-1 type is therefore weakly more selective than the static marginal type.

Return to the cloud-infrastructure example. At Amazon's own prior mean, the immediate-entry cutoff $c_1(v,\delta)$ may have been strictly higher than the static cutoff $c(v)$. Entering in 2003 meant foregoing the information content of incumbents' subsequent silence. Some of that information was eventually revealed: once IBM, HP and Dell continued to frame the market as hardware plus services, the updated belief about the optimistic tail of rival priors became more favourable to entry. The point is not that Amazon should have waited longer; it is that, mechanically, the option to wait would have lifted the threshold for immediate entry relative to the static benchmark, and any equilibrium account of Amazon's entry would have to accommodate this gap between c_1 and c.

For a given prior mean m, define the *immediate-entry correspondence*

$$\mathcal{V}_1(m;\delta) \equiv \{v > 0 : m \text{ is feasible under } F_v \text{ and } m \geq c_1(v,\delta)\}.$$

Proposition 9 implies that $\mathcal{V}_1(m;\delta) \subseteq \mathcal{V}(m)$, with strict inclusion whenever delay has strictly positive value at the margin.

5.4 Uniform dynamic cutoffs and eventual entry

Under Assumption 2, conditioning on silence preserves uniformity. If no entry has occurred through date $t-1$, the rival's prior mean is uniformly distributed on $[-a(v), c_{t-1}(v,\delta)]$. Hence the conditional rival-entry hazard is

$$h_t(v,\delta) = \frac{c_{t-1}(v,\delta) - c_t(v,\delta)}{c_{t-1}(v,\delta) + a(v)} \tag{23}$$

whenever $c_{t-1}(v,\delta) > -a(v)$. Whenever the date-$t$ cutoff is interior in the remaining support, that is, whenever

$$-a(v) < c_t(v,\delta) < c_{t-1}(v,\delta),$$

Proposition 8's indifference condition applies. Since $c_{t+1}(v,\delta) \leq c_t(v,\delta)$, the threshold type enters at date $t+1$ after another silent history, so

$$\widetilde{V}_{t+1}(c_t(v,\delta);v,\delta) = c_t(v,\delta) + X - S - \Delta h_{t+1}(v,\delta).$$

Substituting this and (23) into (21) gives the local recursion

$$c_t + X - S - \Delta\frac{c_{t-1} - c_t}{c_{t-1} + a(v)} = \delta\frac{c_t + a(v)}{c_{t-1} + a(v)}\left[c_t + X - S - \Delta\frac{c_t - c_{t+1}}{c_t + a(v)}\right], \tag{24}$$

where, to ease notation, the arguments (v,δ) are suppressed.

The useful object in the dynamic game is not only whether entry occurs at date 1, but whether it occurs at all.

Proposition 10 (Long-run cutoff and eventual entry under Assumption 2)**.** *Suppose Assumption 2 holds and fix any anonymous monotone timing equilibrium. Then*

$$\lim_{t\to\infty} c_t(v,\delta) = \max\{S - X, -a(v)\}. \tag{25}$$

Define the entry time of type m by

$$\tau(m,v,\delta) \equiv \inf\{t \geq 1 : m \geq c_t(v,\delta)\}, \tag{26}$$

with $\tau(m,v,\delta) = \infty$ *if the set is empty, and define the* eventual-entry correspondence

$$\mathcal{V}_\infty(m;\delta) \equiv \{v > 0 : m \text{ is feasible under } F_v \text{ and } \tau(m,v,\delta) < \infty\}.$$

Then, for any feasible type m:

(a) *if* $m > \max\{S - X, -a(v)\}$, *then* $\tau(m,v,\delta) < \infty$;

(b) *if* $m < \max\{S - X, -a(v)\}$, *then* $\tau(m,v,\delta) = \infty$.

In particular, for every optimistic type $m > 0$ *with* $m > S - X$,

$$\mathcal{V}_\infty(m;\delta) = \left[\frac{m^2}{\kappa^2}, \infty\right). \tag{27}$$

If $m < S - X$, *then* $\mathcal{V}_\infty(m;\delta) = \varnothing$. *At the knife-edge* $m = \max\{S - X, -a(v)\}$, *the entrepreneur is indifferent in the limit; finite-time entry depends on whether the cutoff path reaches its limit in finite time.*

Proposition 10 is the key dynamic result. The long-run cutoff is the larger of two absolute floors: the monopoly-breakeven prior $S - X$, below which entry never pays even as a monopolist, and the lowest feasible prior $-a(v)$, below which no rival exists. Silence gradually screens out the optimistic rivals, so duplication risk falls over time, and in the limit the entrepreneur's problem reduces to the monopolist's problem. Any entrepreneur whose own prior exceeds $S - X$ eventually enters, at every variance level for which it is a feasible type.

The contrast with the static game is especially sharp in the bounded-interval case of Proposition 5(d).

Corollary 2 (How timing modifies the baseline variance result)**.** *Suppose Assumption 2 holds and* $m > 0$ *with* $S - X < m < S - M$. *Then*

$$\mathcal{V}_1(m;\delta) \subseteq \mathcal{V}(m) = \left[\frac{m^2}{\kappa^2}, \left(\frac{m\Delta}{2\kappa(S - M - m)}\right)^2\right] \subsetneq \mathcal{V}_\infty(m;\delta) = \left[\frac{m^2}{\kappa^2}, \infty\right). \tag{28}$$

The upper variance bound in Proposition 5(d) becomes an upper bound on immediate *entry rather than on eventual entry.*

Corollary 2 gives the cleanest interpretation of the timing extension. In the static model, the entrepreneur compares its prior mean to a one-shot cutoff and may stay out permanently when variance is too high. In the dynamic model, the same optimistic entrepreneur can wait, observe silence, and later enter when silence has revealed that the rival is less optimistic than initially feared. High variance delays entry; it does not forbid it.

The role of common variance in the timing game is therefore different from its role in the static benchmark. In Section 4, variance changed entry only through a one-shot crowding probability. In the timing game variance also changes the informational content of silence. A larger v widens the initial support of rival priors, which can raise current duplication risk but can also make a silent history more informative about how the tail of rival optimism will unfold. For that reason the clean sign result in Proposition 4 for the simultaneous cutoff does not automatically extend to the date-1 cutoff $c_1(v, \delta)$. What survives cleanly, and what is the paper's main substantive dynamic claim, is the separation between immediate entry and eventual entry.

6 A Contrarian Perspective

6.1 Contrarian surplus

An economical way to read the results of Sections 3–5 is as a theory of *contrarian conviction*. The central entrepreneurial question is not whether the founder holds a positive view of the opportunity in isolation, but whether that view is sufficiently strong relative to the optimism the same environment is likely to generate in rivals. The paper's cutoff $c(v)$ is the optimism tax levied by expected competition; the founder's useful object is the distance between its own conviction and that tax.

Definition 2 (Contrarian surplus)**.** For an entrepreneur with prior mean m in the simultaneous-move game, define the *contrarian surplus* by

$$\mathcal{T}(m, v) \equiv m - c(v). \tag{29}$$

For the dynamic game, define the *date-1 contrarian surplus* by

$$\mathcal{T}_1(m, v, \delta) \equiv m - c_1(v, \delta). \tag{30}$$

Under Assumption 2, define the *long-run monopoly surplus* by

$$\mathcal{T}_\infty(m, v) \equiv m - \max\{S - X, -a(v)\}. \tag{31}$$

Contrarian surplus is a language or perspective, not a new result. It renames the output of the paper's cutoff machinery in a way that makes the managerial content transparent. The equilibrium cutoff is the tax expected competition levies on a founder's conviction; contrarian surplus is what remains after that tax is paid. This reading connects the model directly to Agrawal et al. (2025) and to the contrarian-conviction intuition that animates the Bayesian-entrepreneurship program: founders are systematically the agents with the highest conviction about the opportunities they pursue, but whether that conviction translates into economic rents depends on how much of it is shared by the rivals the environment can generate.

Corollary 3 (Contrarian surplus and entry)**.**

(i) In the simultaneous-move game, entrepreneur m enters if and only if $\mathcal{T}(m,v) \geq 0$, and the equilibrium payoff is $\max\{\mathcal{T}(m,v), 0\}$.

(ii) In any anonymous monotone timing equilibrium, immediate entry at date 1 *occurs if and only if $\mathcal{T}_1(m,v,\delta) \geq 0$, and*

$$\mathcal{T}_1(m,v,\delta) \leq \mathcal{T}(m,v). \tag{32}$$

(iii) Under Assumption 2, if m is feasible under F_v, then $\mathcal{T}_\infty(m,v) > 0$ guarantees eventual entry and $\mathcal{T}_\infty(m,v) < 0$ rules it out. For optimistic types $m > 0$, the knife-edge case is $\mathcal{T}_\infty(m,v) = 0$: there the entrepreneur is indifferent in the limit, and finite-time entry depends on whether the cutoff path reaches its limit in finite time.

Corollary 3 maps the paper's three formal results onto a single decision criterion. A positive prior mean is not enough. Entry requires a non-negative contrarian surplus. In the dynamic game, the relevant object for acting immediately is the tighter date-1 surplus $\mathcal{T}_1$, because the option to wait has value. Yet eventual entry can still occur even when $\mathcal{T}_1 < 0$, provided the entrepreneur's long-run monopoly surplus $\mathcal{T}_\infty$ is positive, and silence gradually screens out sufficiently optimistic rivals.

6.2 Worked cases

Return to the numerical vignette of Section 3. Take $X = 10$, $\mu = 2$, $S = 7$ and $\kappa = 1$. Then $\Delta = 8$, $M = 6$, $S - M = 1$, and $S - X = -3$. Consider two optimistic types.

First, type $m = 1.5$. Because $S = 7 > M = 6$ and $m = 1.5 \geq S - M = 1$, this type satisfies case (c) of Proposition 5. The static variance-entry correspondence is $\mathcal{V}(m) = [m^2/\kappa^2, \infty) = [2.25, \infty)$. Once this type is feasible (equivalently, once $v \geq 2.25$), it enters in every variance level in the simultaneous game. The dynamic long-run correspondence is the same half-line, because any optimistic type with $m > S - X = -3$ eventually enters. The timing extension can still matter for the date-1 decision: Proposition 9 only implies that immediate entry is weakly harder than in the static benchmark. What case (c) rules out is a long-run upper variance bound on entry, not a delay region in general.

Second, type $m = 0.5$. Because $S - X = -3 < m = 0.5 < S - M = 1$, this type is in case (d). The static variance-entry correspondence is the bounded interval

$$\mathcal{V}(m) = \left[\frac{m^2}{\kappa^2}, \left(\frac{m\Delta}{2\kappa(S-M-m)}\right)^2\right] = \left[0.25, \left(\frac{0.5 \cdot 8}{2 \cdot 0.5}\right)^2\right] = [0.25, 16].$$

Above $v = 16$ the static cutoff exceeds this type's prior mean, and the entrepreneur stays out permanently in the simultaneous game. In the dynamic game, by contrast, the eventual-entry correspondence is $\mathcal{V}_\infty(m;\delta) = [0.25, \infty)$ for every $\delta \in (0,1)$. High variance above $v = 16$ does not force the type to stay out forever; it delays entry until silence has screened out sufficiently optimistic rivals. The exact boundary between *immediate* and *delayed* entry depends on the value of waiting and therefore on δ, but the qualitative message is that the static upper bound is not a boundary of the timing game's feasible set.

These two worked cases illustrate the three regimes of Corollary 3. A high-conviction case (c) entrepreneur enters in the simultaneous game at every feasible variance and, in the dynamic game, eventually enters at every feasible variance; delay may still matter for the date-1 decision. A case (d) entrepreneur enters immediately on the subset $\mathcal{V}_1(m;\delta) \subseteq \mathcal{V}(m)$, may wait on higher feasible variance levels, and eventually enters whenever its prior mean exceeds monopoly break-even. A case (a) entrepreneur (with $m \leq S - X$) never enters, because even monopoly value is insufficient to cover sunk costs.

6.3 Empirical support

The paper's central theoretical claim, that value accrues to founders whose conviction is not merely positive but high *relative to* the distribution of rival conviction the same environment generates, has a close empirical analogue in Gius (2025). The connection is close enough to deserve a dedicated discussion before the validation extension, because it bridges the model's abstractions to an observable object without claiming a one-to-one structural identification.

Gius (2025) studies 67 venture competitions and 2,650 startups in which each venture is scored by at least three judges on a standardised rubric. The central finding is that judge *disagreement* (measured as the standard deviation of scores received by a venture) positively predicts future startup success, where success is measured by follow-on funding raised, revenue, and probability of a successful exit. The relationship survives when controlling for the mean score, which is essential: two ventures with the same average evaluation differ systematically in downstream outcomes depending on how polarising they were. Polarising startups are not more likely to fail, so disagreement is not a proxy for riskiness. The effect is strongest for ventures with distinctive business-model descriptions, and former-founder judges are disproportionately the ones who "stand out" and disagree with other experts.

Gius (2025) interprets the findings through the theory-based view of Felin and Zenger (2009, 2017) and the Bayesian-entrepreneurship framework of Agrawal et al. (2025), which are precisely the two research programs in which this paper is situated. The logic is that common opinion cannot be a source of competitive advantage (Barney, 1991; Felin and Zenger, 2009), so value is disproportionately created and captured by founders with atypical theories that spark disagreement among evaluators. This is the observational counterpart of the model's central object. Under heterogeneous priors, the distribution of conviction across evaluators carries information that the mean alone does not. A wide score distribution on a venture corresponds to a wide empirical analogue of F_v with non-trivial mass in both tails, which is precisely the environment in which a founder whose own prior sits in the upper tail has positive contrarian surplus. The mapping is interpretive rather than structural: evaluator-score dispersion is an observable proxy for the latent distribution of rival priors, not a direct measure of F_v.

Two specific mappings to the paper's objects are worth drawing out. First, Gius (2025)'s measured within-venture score dispersion can be read as an empirical proxy for the rival-prior range induced by v in this model. Ventures that elicit high judge dispersion are those for which the environment admits a wide range of rational priors. The model is consistent with, but does not by itself imply, a monotone "more disagreement is better" conclusion. What it does imply is that disagreement contains information beyond the mean about whether the founder sits sufficiently far into the upper tail of the relevant belief distribution. For high-conviction founders, wider disagreement can generate positive contrarian surplus; for intermediate-conviction founders in high-cost environments, the same widening can eventually raise the competitive cutoff enough to block immediate entry. Second, Gius (2025) frames the results as a partial remedy to the "uniqueness paradox" (Litov et al., 2012; Benner and Zenger, 2016; Zuckerman, 1999): unique strategies are essential for sustained advantage but are discounted because they are hard to evaluate. The next subsection develops the public-validation result, which gives a game-theoretic expression of the same tension from the founder's side. Gius (2025) shows the discount empirically; this paper's model shows why a founder should anticipate it and shape its validation strategy accordingly.

The two contributions are complementary. Gius (2025) identifies disagreement as a measurable marker of value. The present paper gives a structural interpretation of that marker: disagreement is informative about contrarian surplus, but its sign depends on where the founder's own conviction sits relative to the induced cutoff. The paper's founder-side cutoff can, in principle, be taken to data of the Gius (2025) kind by using judge-score dispersion as an empirical analogue of the rival-belief range and founders' own-assessment data as an analogue of m_i.

6.4 Private validation versus public validation

The baseline model treats prior means as primitives. That is appropriate for the threshold and timing results above, but it suppresses an important entrepreneurial distinction: some actions improve the founder's own conviction without educating rivals, while other actions improve conviction in a way that also makes the opportunity more legible to the market. This subsection adds the smallest extension needed to formalise the difference without changing the baseline game.

Take the simultaneous-move environment as the benchmark. Start from an entrepreneur with prior mean m facing rival meta-distribution F_v. A *private validation* increment $q \in \mathbb{R}$ shifts only the entrepreneur's own prior mean from m to $m + q$, leaving the rival meta-distribution unchanged. Examples include proprietary diligence, internal experimentation, private customer conversations, and learning by doing under confidentiality. A *public validation* increment $z \in \mathbb{R}$ shifts every entrepreneur's prior mean by the same amount. Equivalently, the rival meta-distribution becomes

$$F_v^z(x) \equiv F_v(x - z).$$

Examples include public traction milestones, visible regulatory approvals, industry-wide disclosures, and any broadly legible signal that moves beliefs across the market rather than only inside the venture.

Proposition 11 (Private validation versus public validation)**.** *Fix $v > 0$ and a type m.*

(a) ***Private validation.*** *If an increment q raises only entrepreneur i's prior mean from m to $m+q$, leaving the rival meta-distribution F_v unchanged, then the anonymous-equilibrium payoff is*

$$\Pi^{\mathrm{priv}}(m,q;v) = \max\{m+q-c(v),0\}. \tag{33}$$

Whenever $m+q > c(v)$,

$$\frac{\partial \Pi^{\mathrm{priv}}(m,q;v)}{\partial q} = 1. \tag{34}$$

(b) ***Public validation.*** *If an increment z shifts every entrepreneur's prior mean by z, so the rival distribution becomes $F_v^z(x) = F_v(x-z)$, then there exists a unique anonymous cutoff $c(v,z)$ satisfying*

$$c(v,z) = S - X + \Delta\big(1 - F_v(c(v,z)-z)\big). \tag{35}$$

If F_v is continuously differentiable with density f_v and $c(v,z)-z$ lies in the interior of the support of F_v, then

$$\frac{\partial c(v,z)}{\partial z} = \frac{\Delta f_v(c(v,z)-z)}{1+\Delta f_v(c(v,z)-z)} \in [0,1), \tag{36}$$

and the anonymous-equilibrium payoff of type m is

$$\Pi^{\mathrm{pub}}(m,z;v) = \max\{m+z-c(v,z),0\}. \tag{37}$$

Whenever $m+z > c(v,z)$,

$$\frac{\partial \Pi^{\mathrm{pub}}(m,z;v)}{\partial z} = \frac{1}{1+\Delta f_v(c(v,z)-z)} \in (0,1]. \tag{38}$$

Therefore, an equally sized increment of private validation weakly increases the entrepreneur's equilibrium payoff more than an equally sized increment of public validation, with strict inequality whenever $f_v(c(v,z)-z) > 0$.

Proposition 11 gives a formal expression to a practical entrepreneurial distinction. Private validation improves the founder's position one-for-one because it raises conviction without changing the competitive threshold. Public validation raises conviction too, but it partly dissipates that gain by making the opportunity more compelling to rivals and therefore lifting the cutoff. The derivative of the cutoff with respect to a public increment, $\Delta f_v(c(v,z)-z)/(1+\Delta f_v(c(v,z)-z))$, is strictly less than one but strictly positive whenever the density at the shifted cutoff is positive.

Under the bounded-range family in Assumption 2, the same point becomes especially transparent.

Corollary 4 (Public validation under uniform belief dispersion)**.** *Suppose Assumption 2 holds and the public-validation cutoff lies in the interior of the support of F_v^z. Then*

$$c(v,z) = c(v) + \frac{\Delta}{2a(v) + \Delta} z, \tag{39}$$

and the entrepreneur's payoff becomes

$$\Pi^{\text{pub}}(m,z;v) = \max\left\{m - c(v) + \frac{2a(v)}{2a(v)+\Delta} z,\ 0\right\}. \tag{40}$$

Hence, the entrepreneur retains only the fraction

$$\frac{2a(v)}{2a(v)+\Delta} \tag{41}$$

of a public validation increment as additional contrarian surplus, while the fraction

$$\frac{\Delta}{2a(v)+\Delta}$$

is dissipated through the higher competitive cutoff. The retained fraction is strictly decreasing in the duplication wedge Δ.

The retained-fraction formula has a clean economic reading. Markets with large duplication wedges punish public legibility severely, because a public increment of conviction is partly handed to rivals whose entry would then dissipate most of the rent. Markets with small duplication wedges are relatively tolerant of public validation, because simultaneous entry destroys little value. Low-variance environments (small $a(v)$) also penalise public legibility more than high-variance environments, because in a tightly clustered belief distribution, a small shift moves many rivals across the threshold.

A useful illustration contrasts two AI-enabled ventures that each run a pilot study of a new technique on customer data. The first runs the pilot privately and reports results only to existing investors under non-disclosure. The second publishes an open benchmark and gives a conference presentation on the pilot's early results. Both ventures learn the same thing about the underlying mechanism, so the conviction increment is the same. The first retains the increment fully as a contrarian surplus. The second retains only the fraction $2a(v)/(2a(v)+\Delta)$: rivals read the benchmark, update their own priors by z, and the competitive cutoff shifts up, eating into the founder's rent. The managerial implication is that the relevant instrument is not information per se but *legibility*. Diligence that is privately conducted strictly dominates equally informative public milestones, holding everything else fixed.

The public-validation exercise is not just a managerial refinement. It is a reduced-form probe of what happens once a founder's mean acquires a shared component. A public signal z adds the same amount to every founder's prior mean, but conditional on that common shift, the idiosyncratic residual components can still remain independent, exactly as in Section 2.4. Corollary 4 should, therefore, be read as a common-shift extension rather than as a full conditional-correlation model of priors. What the founder loses relative to the baseline is not the independence of the residual component per se; rather, some of the increase in own conviction is mechanically shared by rivals and therefore feeds back into the competitive cutoff. A full common-source model, in which observing a high own mean changes beliefs about the conditional distribution $F_v(\cdot \mid m_i)$ of rival means, remains outside the present setup.

6.5 What the model says, and does not say, to entrepreneurs

Three implications of Corollary 3 and Proposition 11 are worth stating compactly. First, contrarianism is *relative*, not absolute. The right question is not whether the founder likes the opportunity; it is whether the founder likes it enough, relative to the distribution of rival conviction implied by the same environment. The paper's threshold $c(v)$ formalises a competition-adjusted version of this intuition. Valuable opportunities are not merely those that look good, but those for which the founder's conviction remains high after accounting for how many others the environment can generate who also find the idea compelling.

Second, high variance can mean *wait* rather than *do not enter*. In the simultaneous-move benchmark, sufficiently high variance may push an intermediate type below the competitive threshold and therefore rule out entry. In the dynamic game, the same type may enter later because silence reveals that the rival is less optimistic than initially feared. Some opportunities are not bad opportunities; they are opportunities for which the founder's current contrarian surplus is negative but can become positive after delay.

Third, private conviction is worth more, at the margin, than public conviction. A unit of private validation raises contrarian surplus one-for-one. A unit of public validation raises it only by the fraction $2a(v)/(2a(v) + \Delta)$. Founders who have the option to stage validation should favour mechanisms that strengthen their own conviction without simultaneously educating rivals, and should recognise that the public-legibility penalty is strictly increasing in the duplication wedge.

The model's message is, therefore, not "seek uncertain ideas" or "avoid uncertain ideas." It is: seek ideas for which the founder's conviction is strong enough relative to the likely rival conviction, recognise that some environments make immediate entry unattractive while still rewarding patience, and prefer private validation to equally informative public validation.

7 Conclusion

When entrepreneurs hold heterogeneous prior means but share a common environmental variance, entry reduces to comparing conviction with a competition-adjusted threshold. The simultaneous-move benchmark admits a unique anonymous monotone cutoff equilibrium, and under a bounded-range belief-dispersion family the cutoff is closed form. The set of variance levels consistent with entry by a given optimistic entrepreneur is either empty, a half-line, or a bounded interval, and the sign of the comparative static of variance is pinned down by whether the sunk cost of entry is above or below the average of monopoly and duplication returns.

Adding timing changes the picture in a specific benchmark way. The infinite-horizon dynamic game with irreversible entry admits at least one anonymous monotone perfect Bayesian equilibrium in pure cutoff strategies. Within the anonymous monotone class, silence becomes informative, the date-1 entry cutoff is weakly above the static cutoff, and high variance that blocks immediate entry in the static game may only delay entry in the dynamic game. Under the bounded-range family, every feasible variance level eventually supports entry for any optimistic entrepreneur whose monopoly value is strictly positive.

Three extensions are natural avenues for future work. First, endogenising the entrepreneur's choice of validation intensity or disclosure policy inside the dynamic game would add a theory of how founders invest in conviction and legibility, rather than taking both as exogenous. Second, asymmetric belief-dispersion families would allow for environments in which the support of rational priors is skewed (heavy optimistic or pessimistic tails), which seems especially relevant for frontier technologies where most opinion is clustered in one tail. Third, and most importantly, relaxing the independence-of-priors assumption of Section 2.4 is a natural direction. A signal-based model in which each entrepreneur's prior mean is generated by a private interpretation of a partially shared signal would induce positive correlation between m_i and m_j conditional on the environment. The cutoff condition would then become a fixed point in the conditional distribution $F_v(\cdot \mid m_i)$, and the comparative statics of variance would need to disentangle two channels: variance as the width of the rival's belief range, and variance as a modulator of the correlation between own and rival priors. The present paper should be read as a clean benchmark for the idiosyncratic-priors regime that such an extension would nest as a special case.

In sum, contrarianism is relative, not absolute. The right question is whether the founder's conviction is high enough relative to the conviction that the same environment is likely to generate in rivals. Some environments reward immediate entry, some reward patience, and some do not justify entry at all. Private validation raises contrarian surplus one-for-one; public validation only partially, and the penalty grows with the duplication wedge. The model gives a formal version of a contrarian-conviction intuition while remaining silent on whether prior heterogeneity is rational or behavioural, and it provides a theoretical counterpart to the recent empirical finding that evaluator disagreement is a marker, rather than a warning, about startup success.

A Appendix: Proofs

Throughout the Appendix, the arguments (v, δ) are suppressed when no confusion can arise.

Proof of Proposition 1. If entrepreneur i enters and assigns probability p to rival entry, then the expected payoff is

$$(1-p)(m_i + X - S) + p(m_i + \mu - S) = m_i + X - S - \Delta p.$$

The variance v does not appear directly. Any effect of v must therefore operate through the entrepreneur's belief p about rival entry. □

Proof of Proposition 2. Fix $v > 0$ and conjecture that the rival uses cutoff c. Then the rival enters if and only if $m_j \geq c$, which occurs with probability $1 - F_v(c)$. From (2), type m_i prefers entering to staying out if and only if

$$m_i + X - S - \Delta\big(1 - F_v(c)\big) \geq 0,$$

which is equivalent to

$$m_i \geq S - X + \Delta\big(1 - F_v(c)\big).$$

Thus the best response to a conjectured cutoff c is itself a cutoff, namely

$$B_v(c) \equiv S - X + \Delta\big(1 - F_v(c)\big).$$

An anonymous cutoff equilibrium therefore requires a fixed point $c = B_v(c)$.

Define

$$G_v(c) \equiv c - B_v(c) = c - S + X - \Delta\big(1 - F_v(c)\big).$$

Because F_v is non-decreasing and continuous, G_v is continuous and strictly increasing in c. Moreover,

$$\lim_{c \to -\infty} G_v(c) = \lim_{c \to -\infty} (c - S + \mu) = -\infty,$$

while

$$\lim_{c \to +\infty} G_v(c) = \lim_{c \to +\infty} (c - S + X) = +\infty.$$

By continuity, a fixed point exists. By strict monotonicity, it is unique. Call it $c(v)$. The anonymous best response is then to enter if and only if $m_i \geq c(v)$, which completes the proof. □

Proof of Corollary 1. At the equilibrium cutoff $c(v)$, equation (3) implies

$$X - S - \Delta\big(1 - F_v(c(v))\big) = -c(v).$$

Substituting this into (2) gives the expected payoff from entry for type m_i:

$$\Pi_E(m_i; c(v), v) = m_i - c(v).$$

Types with $m_i \geq c(v)$ enter and obtain $m_i - c(v)$; types with $m_i < c(v)$ stay out and obtain 0. Hence

$$\Pi(m_i, v) = \max\{m_i - c(v), 0\}.$$

$\square$

Proof of Proposition 3. Under Assumption 2, write $a \equiv a(v) = \kappa\sqrt{v}$.

If $c \leq -a$, then every type enters. This is an equilibrium if and only if the lowest type is willing to enter, given certain duplication. The lowest type's payoff is $-a + \mu - S$, so pooling entry occurs if and only if

$$-a + \mu - S \geq 0,$$

which is condition (5).

If $c \geq a$, then no type enters. This is an equilibrium if and only if the highest type does not wish to enter even as a monopolist. The highest type's monopoly payoff is $a + X - S$, so no entry occurs if and only if

$$a + X - S \leq 0,$$

which is condition (6).

In the remaining case, $c \in (-a, a)$ and

$$F_v(c) = \frac{c + a}{2a}.$$

Substituting this into the fixed-point equation (3) yields

$$c = S - X + \Delta\left(1 - \frac{c + a}{2a}\right) = S - X + \Delta\frac{a - c}{2a}.$$

Rearranging gives

$$c\left(1 + \frac{\Delta}{2a}\right) = S - X + \frac{\Delta}{2} = S - \frac{X + \mu}{2} = S - M,$$

so

$$c = \frac{2a}{2a + \Delta}(S - M),$$

which is (8). The equilibrium action is then to enter if and only if $m_i \geq c(v)$. $\square$

Proof of Proposition 4. In the interior region,

$$c(v) = \frac{2a(v)}{2a(v) + \Delta}(S - M), \qquad a(v) = \kappa\sqrt{v}.$$

Differentiating with respect to a gives

$$\frac{dc}{da} = \frac{2\Delta}{(2a + \Delta)^2}(S - M).$$

Since $da/dv = \kappa/(2\sqrt{v}) > 0$, the sign of dc/dv is the same as the sign of $S - M$. This proves (9) and the three cases stated in the proposition. □

Proof of Proposition 5. Fix $m > 0$.

Part (a). If $m \leq S - X$, then even monopoly entry has expected payoff

$$m + X - S \leq 0.$$

Since duplication can only reduce the payoff further, entry is never optimal. Hence $\mathcal{V}(m) = \varnothing$.

Now suppose $m > S - X$. A necessary feasibility condition is that m lies in the support of F_v, which under Assumption 2 requires

$$m \leq a(v) = \kappa\sqrt{v},$$

equivalently,

$$v \geq \frac{m^2}{\kappa^2}.$$

Part (b). If $S \leq M$, then in the interior region Proposition 3 gives

$$c(v) = \frac{2a(v)}{2a(v) + \Delta}(S - M) \leq 0.$$

Because $m > 0$, every feasible optimistic type satisfies $m \geq c(v)$. In the pooling-entry region, all types enter by definition. Since $S \leq M < X$, the no-entry region cannot occur for any feasible v. Therefore

$$\mathcal{V}(m) = \left[\frac{m^2}{\kappa^2}, \infty\right).$$

Part (c). Suppose $S > M$ and $m \geq S - M$. In the relevant interior region the cutoff is increasing in v by Proposition 4 and satisfies

$$\lim_{v \to \infty} c(v) = S - M.$$

Hence $c(v) \leq S - M \leq m$ for every $v > 0$. Pooling entry cannot arise when $S > M > \mu$, and once m is feasible, the no-entry region is excluded because $m > S - X$ implies $a(v) \geq m > S - X$. Therefore, the entrepreneur enters for every feasible variance level, so

$$\mathcal{V}(m) = \left[\frac{m^2}{\kappa^2}, \infty\right).$$

Part (d). Suppose $S - X < m < S - M$. Since $S > M > \mu$, pooling entry cannot arise, and once m is feasible, the no-entry region is excluded by the same argument as above. Entry, therefore, requires only that $m \geq c(v)$ in the interior region. Using (8), this is

$$m \geq \frac{2a(v)}{2a(v) + \Delta}(S - M).$$

Because $S - M - m > 0$, rearranging gives

$$2a(v)\big(S - M - m\big) \leq m\Delta,$$

so

$$a(v) \leq \frac{m\Delta}{2(S - M - m)}.$$

Substituting $a(v) = \kappa\sqrt{v}$ yields

$$v \leq \left(\frac{m\Delta}{2\kappa(S - M - m)}\right)^2.$$

Combining this with the feasibility condition $v \geq m^2/\kappa^2$ proves

$$\mathcal{V}(m) = \left[\frac{m^2}{\kappa^2}, \left(\frac{m\Delta}{2\kappa(S - M - m)}\right)^2\right].$$

□

Proof of Proposition 6. By Corollary 1, the entrepreneur's equilibrium payoff is

$$\Pi(m, v) = \max\{m - c(v), 0\}.$$

Under Assumption 2, Proposition 4 shows that $c(v)$ is weakly monotone in v on the interior region, with sign determined by $S - M$. The corner regions preserve the same monotonicity because they correspond to $c(v)$ lying weakly outside the support. Therefore $\Pi(m, v)$ is weakly monotone in v on $A \equiv \mathcal{W} \cap \mathcal{V}(m)$.

If A is empty, there is nothing to prove. Suppose $\Pi(m, \cdot)$ attains a maximum on A. If it is constant on A, then every point of A is optimal. If it is not constant, then a monotone function on a subset of the real line attains any maximum it has at an extreme point of that subset: at sup A when it is weakly increasing, and at inf A when it is weakly decreasing. Such extreme points are boundary points of A. Hence, at least one maximiser lies on the boundary of $\mathcal{W} \cap \mathcal{V}(m)$. □

Lemma 1 (Monotone best responses to a monotone cutoff path)**.** *Fix $v > 0$, $\delta \in (0, 1)$, and a non-increasing rival cutoff path $c = \{c_t\}_{t \geq 1} \subseteq [\underline{m}(v), \overline{m}(v)]$. Let $c_0 = \overline{m}(v)$. If a type m plans to enter for the first time at date t, its expected payoff is*

$$U_t(m; c) = \delta^{t-1}\Big[F_v(c_{t-1})(m + X - S) - \Delta\big(F_v(c_{t-1}) - F_v(c_t)\big)\Big]. \tag{42}$$

The payoff from never entering is $U_\infty(m; c) = 0$. Moreover, $U_t(\cdot; c)$ is affine in m with slope

$$\alpha_t(c) \equiv \delta^{t-1}F_v(c_{t-1}),$$

and the sequence $\{\alpha_t(c)\}_{t\geq 1}$ is non-increasing. Hence the latest optimal entry date

$$\tau^*(m;c) \equiv \max \arg \max_{a\in\mathbb{N}\cup\{\infty\}} U_a(m;c)$$

is weakly decreasing in m. In particular, every monotone rival cutoff path admits a monotone pure cutoff best response.

Proof. If type m plans to enter at date t, then rival non-entry through date $t-1$ occurs exactly when the rival's type lies below c_{t-1}, which has probability $F_v(c_{t-1})$. Simultaneous entry at date t occurs when the rival's type lies in $[c_t, c_{t-1})$, which has probability $F_v(c_{t-1}) - F_v(c_t)$. This yields (42). The payoff from never entering is 0.

Because Assumption 1 gives compact support, there is a constant $K < \infty$ such that $|m + X - S| \leq K$ for all $m \in [\underline{m}(v), \overline{m}(v)]$. Hence

$$|U_t(m;c)| \leq \delta^{t-1}(K + \Delta)$$

for every feasible (m, c), so $U_t(m;c) \to 0 = U_\infty(m;c)$ uniformly as $t \to \infty$. Therefore the maximisation problem over $\mathbb{N} \cup \{\infty\}$ is well defined and the latest maximiser exists.

For $t < s$, the difference $U_t(m;c) - U_s(m;c)$ is affine in m with slope $\alpha_t(c) - \alpha_s(c) \geq 0$, because c is non-increasing, F_v is non-decreasing, and $\delta \in (0,1)$. Thus, once date t weakly dominates date s at some type m, it weakly dominates date s for every higher type. This single-crossing property implies that the latest optimal date $\tau^*(m;c)$ is weakly decreasing in m. The induced best response can therefore be represented by a non-increasing cutoff path. □

Lemma 2 (Finite-truncation anonymous best-response correspondence)**.** *Fix $T \in \mathbb{N}$ and let K_T denote the compact convex set of cutoff vectors*

$$K_T \equiv \Big\{(c_1, \ldots, c_T) \in [\underline{m}(v), \overline{m}(v)]^T : c_1 \geq \cdots \geq c_T\Big\}.$$

For $c \in K_T$, let $\mathcal{B}_T(c) \subseteq K_T$ denote the set of cutoff vectors induced by monotone pure best replies to an opponent using c in the T-action planned-entry game. Then $\mathcal{B}_T(c)$ is non-empty, compact, and convex for every $c \in K_T$, and the correspondence $\mathcal{B}_T : K_T \rightrightarrows K_T$ has closed graph.

Proof. By Lemma 1, for every pair of actions in the finite action set, the payoff difference is affine and weakly increasing in the entrepreneur's type. The T-action planned-entry game therefore satisfies the single-crossing condition in the sense of Athey (2001, Theorem 1). The jump-point construction in that theorem identifies monotone pure best responses with ordered vectors of cutoff points. Under the standard embedding of jump points into K_T, the set of monotone pure best responses is a non-empty complete interval of cutoff vectors and is therefore compact and convex. Continuity of the payoff array $U_t(m;c)$ in c implies that the best-response inequalities defining that interval are closed, so the graph of $\mathcal{B}_T$ is closed. □

Proof of Proposition 7. Fix $v > 0$ and $\delta \in (0,1)$, and abbreviate $\underline{m} \equiv \underline{m}(v)$ and $\overline{m} \equiv \overline{m}(v)$.

Step 1: a planned-entry normal form. Represent a pure action by a planned entry date $a \in A \equiv \{0\} \cup \{1/t : t \in \mathbb{N}\}$, where $1/t$ means "enter at date t" and 0 means "never enter." Earlier dates are higher actions in the order on A. For type m, the stage payoff against rival action a' is

$$u(m, a, a') = \begin{cases} \delta^{t-1}(m + X - S), & a = 1/t > a', \\ \delta^{t-1}(m + \mu - S), & a = a' = 1/t, \\ 0, & a < a' \text{ or } a = 0. \end{cases}$$

Because A has only one accumulation point, namely 0, and $\delta^{t-1} \to 0$, the function u is bounded and continuous on $[\underline{m}, \overline{m}] \times A \times A$. A monotone pure strategy in this normal-form game is a nondecreasing map $\sigma : [\underline{m}, \overline{m}] \to A$. Such a strategy is equivalent to a non-increasing cutoff path $\{c_t\}_{t \geq 1}$, with type m entering at date t exactly when $\sigma(m) = 1/t$.

Step 2: finite-action truncations. For each $T \in \mathbb{N}$, let

$$A_T \equiv \{0, 1, 1/2, \ldots, 1/T\}.$$

Let K_T be as in Lemma 2. For $c \in K_T$, let $\mathcal{B}_T(c) \subseteq K_T$ denote the set of anonymous monotone best responses to an opponent using c in the T-action planned-entry game. Lemma 2 implies that $\mathcal{B}_T(c)$ is non-empty, compact, and convex for every c, and that the correspondence $\mathcal{B}_T$ has a closed graph. Kakutani's fixed-point theorem therefore yields $c^T \in K_T$ with $c^T \in \mathcal{B}_T(c^T)$. Hence, the truncated game with action set A_T admits an anonymous monotone Bayesian equilibrium; let σ^T denote the associated strategy.

Step 3: passage to the infinite horizon. Each σ^T is nondecreasing and takes values in the compact ordered set $A \subset [0, 1]$. By Helly's Selection Theorem, there exists a subsequence $\{T_n\}$ and a nondecreasing function $\sigma : [\underline{m}, \overline{m}] \to A$ such that $\sigma^{T_n}(m) \to \sigma(m)$ at every continuity point of σ. Because a monotone real-valued function has at most countably many discontinuities, we may refine the subsequence by a diagonal argument so that the convergence also holds at each discontinuity of σ. Thus

$$\sigma^{T_n}(m) \longrightarrow \sigma(m) \qquad \text{for every } m \in [\underline{m}, \overline{m}].$$

Fix any type m and any action $a \in A$. If $a = 1/t$, then $a \in A_{T_n}$ for all sufficiently large n; if $a = 0$, it belongs to every A_{T_n}. Since σ^{T_n} is an equilibrium of the T_n-action game,

$$U(m, \sigma^{T_n}(m); \sigma^{T_n}) \geq U(m, a; \sigma^{T_n})$$

for all large n, where U denotes expected payoff in the planned-entry normal form. Because u is bounded and continuous and $\sigma^{T_n} \to \sigma$ pointwise, dominated convergence implies

$$U(m, \sigma^{T_n}(m); \sigma^{T_n}) \to U(m, \sigma(m); \sigma), \qquad U(m, a; \sigma^{T_n}) \to U(m, a; \sigma).$$

Taking limits gives

$$U(m, \sigma(m); \sigma) \geq U(m, a; \sigma) \qquad \text{for every } a \in A.$$

Thus σ is a best response to itself in the full planned-entry game.

Step 4: from a normal-form equilibrium to a perfect Bayesian equilibrium. Because σ is nondecreasing, it can be represented by a non-increasing cutoff path $\{c_t(v,\delta)\}_{t\geq 1}$. Consider any on-path silent history through date $s-1$, so $F_v(c_{s-1}(v,\delta)) > 0$. If a remaining type m plans to enter at some date $r \geq s$, the continuation payoff from date s onward equals

$$\delta^{r-s}\left[\frac{F_v(c_{r-1})}{F_v(c_{s-1})}(m + X - S) - \Delta\frac{F_v(c_{r-1}) - F_v(c_r)}{F_v(c_{s-1})}\right] = \frac{U_r(m;c)}{\delta^{s-1}F_v(c_{s-1})}.$$

Hence, every continuation payoff from a feasible later date is the corresponding date-1 payoff multiplied by the same positive factor $[\delta^{s-1}F_v(c_{s-1})]^{-1}$. The ranking of continuation actions is therefore unchanged after any on-path silent history, so the planned-date equilibrium is sequentially rational. Bayes' rule gives beliefs after on-path silent histories as in (17); off-path beliefs may be chosen arbitrarily. The induced cutoff path is therefore an anonymous monotone perfect Bayesian equilibrium in pure cutoff strategies. □

Proof of Proposition 8. Fix any anonymous monotone timing equilibrium with cutoff sequence $\{c_t(v,\delta)\}_{t\geq 1}$, and consider a date t such that the silent history through date $t-1$ occurs with positive probability, equivalently $F_v(c_{t-1}(v,\delta)) > 0$. If no entry has occurred through date $t-1$, then equilibrium behaviour implies that the rival's type must satisfy $m_j < c_{t-1}(v,\delta)$. Bayes' rule therefore yields the truncated posterior (17). The rival enters at date t if and only if $m_j \in [c_t(v,\delta), c_{t-1}(v,\delta))$, so the conditional hazard is

$$h_t(v,\delta) = \Pr\big(c_t(v,\delta) \leq m_j < c_{t-1}(v,\delta) \mid m_j < c_{t-1}(v,\delta)\big) = 1 - \frac{F_v(c_t(v,\delta))}{F_v(c_{t-1}(v,\delta))},$$

which is (18).

Let $\widetilde{V}_{t+1}(m;v,\delta)$ denote the continuation value from date $t+1$ onward for type m, conditional on no entry at date t. Conditional on no entry through date $t-1$, immediate entry at date t yields

$$\delta^{t-1}\big[m + X - S - \Delta h_t(v,\delta)\big],$$

whereas waiting yields

$$\delta^{t}\big(1 - h_t(v,\delta)\big)\widetilde{V}_{t+1}(m;v,\delta).$$

Because the equilibrium strategy at date t is the cutoff rule $m \geq c_t(v,\delta)$, the former weakly dominates the latter for types $m \geq c_t(v,\delta)$ and is weakly dominated by it for types $m < c_t(v,\delta)$. Dividing by the positive factor δ^{t-1} proves (19) and (20).

If $\underline{m}(v) < c_t(v,\delta) < c_{t-1}(v,\delta)$, both action regions are non-empty in a neighbourhood of $c_t(v,\delta)$. The waiting-value function is the supremum of affine continuation payoffs from dates

$t+1, t+2, \ldots$, hence finite and convex in m, and therefore continuous on the interior of the type interval. The difference between the two sides of (19)–(20) is thus continuous in m. Since it is weakly non-negative above the cutoff and weakly non-positive below it, it must equal zero at the interior threshold, which gives (21). Monotonicity of entry dates is part of the equilibrium class under study and, equivalently, follows from the same single-crossing property used in Lemma 1: immediate-entry payoff is strictly increasing in m, and payoff differences across entry dates are weakly increasing in m. $\square$

Proof of Proposition 9. Define

$$L(c) \equiv c + X - S - \Delta\big(1 - F_v(c)\big).$$

Because F_v is continuous and non-decreasing, $L(c)$ is continuous and strictly increasing in c. By Proposition 2, the simultaneous-move cutoff $c(v)$ is the unique solution to $L(c(v)) = 0$.

Let $\widetilde{W}_1(m; v, \delta)$ denote the equilibrium value of waiting at date 1 for type m. Since $c_0(v, \delta) = \overline{m}(v)$ and $F_v(c_0(v, \delta)) = 1$, the date-1 hazard is $h_1(v, \delta) = 1 - F_v(c_1(v, \delta))$. The equilibrium strategy prescribes entry at date 1 for type $c_1(v, \delta)$, so

$$L\big(c_1(v,\delta)\big) = c_1(v,\delta) + X - S - \Delta h_1(v,\delta) \geq \widetilde{W}_1\big(c_1(v,\delta); v, \delta\big).$$

The waiting value is weakly non-negative because the entrepreneur can always continue to wait forever and obtain 0. Hence $L(c_1(v, \delta)) \geq 0$. Since L is strictly increasing and $L(c(v)) = 0$, it follows that $c_1(v, \delta) \geq c(v)$.

If the continuation value from waiting is strictly positive for the marginal date-1 type, then $\widetilde{W}_1(c_1(v, \delta); v, \delta) > 0$, so the inequality above is strict. Hence $L(c_1(v, \delta)) > 0$ and therefore $c_1(v, \delta) > c(v)$. $\square$

Proof of Proposition 10. Fix any anonymous monotone timing equilibrium under Assumption 2. The cutoff sequence $\{c_t(v, \delta)\}_{t \geq 1}$ is non-increasing by the equilibrium class under study. It is bounded below by $-a(v)$ because types below the support do not exist. It is also bounded below by $S - X$. To see this, note that if $m < S - X$, then at any on-path silent history through date $t - 1$, the payoff from entering immediately is at most

$$\delta^{t-1}(m + X - S) < 0,$$

because duplication can only reduce the payoff further. Waiting forever yields 0, so such types never enter in equilibrium. Therefore

$$c_t(v, \delta) \geq \max\{S - X, -a(v)\}$$

for every t, and the monotone sequence has a limit, call it $c_\infty(v, \delta)$.

Suppose, for contradiction, that $c_\infty(v,\delta) > \max\{S - X, -a(v)\}$. Choose any type

$$m \in \big(\max\{S - X, -a(v)\},\, c_\infty(v,\delta)\big).$$

Then m is feasible, satisfies $m + X - S > 0$, and also satisfies $m < c_t(v,\delta)$ for every t, so equilibrium prescribes that this type never enters.

Because $c_\infty(v,\delta) > -a(v)$, the on-path silent history through date $t-1$ has positive probability for all sufficiently large t, and the uniform hazard formula (23) is well defined there. Since $c_t(v,\delta) \downarrow c_\infty(v,\delta)$, we have

$$h_t(v,\delta) = \frac{c_{t-1}(v,\delta) - c_t(v,\delta)}{c_{t-1}(v,\delta) + a(v)} \longrightarrow 0.$$

Hence there exists T such that, for all $t \geq T$,

$$m + X - S - \Delta h_t(v,\delta) > 0.$$

At any on-path silent history through date $t-1$ with $t \geq T$, type m can therefore deviate and enter immediately, obtaining the strictly positive payoff

$$\delta^{t-1}\big[m + X - S - \Delta h_t(v,\delta)\big] > 0.$$

But under the equilibrium cutoff path, the same type waits forever from that history onward, because $m < c_s(v,\delta)$ for every $s \geq t$, and so obtains continuation value 0. This contradicts sequential rationality. Therefore

$$c_\infty(v,\delta) = \max\{S - X, -a(v)\},$$

which proves (25).

Now fix any feasible type m. If $m > c_\infty(v,\delta)$, then $m > \max\{S - X, -a(v)\}$. Since $c_t(v,\delta) \downarrow c_\infty(v,\delta)$, there exists a finite date t such that $m \geq c_t(v,\delta)$. Hence $\tau(m,v,\delta) < \infty$.

If $m < c_\infty(v,\delta)$, then $m < \max\{S - X, -a(v)\} \leq c_t(v,\delta)$ for every t, so $\tau(m,v,\delta) = \infty$.

For the optimistic case $m > 0$, feasibility under Assumption 2 requires $m \leq a(v)$, equivalently $v \geq m^2/\kappa^2$. Because $m > 0 > -a(v)$, we have

$$m > \max\{S - X, -a(v)\} \quad \Longleftrightarrow \quad m > S - X.$$

Therefore, for every optimistic type $m > 0$ with $m > S - X$,

$$\mathcal{V}_\infty(m;\delta) = \left[\frac{m^2}{\kappa^2}, \infty\right),$$

which is (27). If $m < S - X$, the earlier argument shows that entry at any finite date is strictly dominated by waiting forever, so $\mathcal{V}_\infty(m;\delta) = \varnothing$.

At the knife-edge $m = c_\infty(v,\delta) = \max\{S - X, -a(v)\}$, the entrepreneur is indifferent in the limit because the limiting net monopoly payoff is zero. Whether finite-time entry occurs depends on whether the cutoff path reaches its limit in finite time. □

Proof of Corollary 2. For $m > 0$ with $S - X < m < S - M$, Proposition 5(d) gives the simultaneous-move variance-entry correspondence

$$\mathcal{V}(m) = \left[\frac{m^2}{\kappa^2}, \left(\frac{m\Delta}{2\kappa(S - M - m)}\right)^2\right].$$

Proposition 9 implies $c_1(v,\delta) \geq c(v)$ for every v, so immediate entry in the dynamic game requires static entry. Therefore $\mathcal{V}_1(m;\delta) \subseteq \mathcal{V}(m)$.

By Proposition 10, every feasible optimistic type with $m > S - X$ eventually enters, so

$$\mathcal{V}_\infty(m;\delta) = \left[\frac{m^2}{\kappa^2}, \infty\right).$$

Since $m < S - M$, the upper bound in $\mathcal{V}(m)$ is finite. Hence $\mathcal{V}(m)$ is a proper subset of $\mathcal{V}_\infty(m;\delta)$, which proves (28). □

Proof of Corollary 3. Part (i) follows immediately from Definition 2 and Corollary 1: in the simultaneous-move game, entry occurs if and only if $m \geq c(v)$, which is equivalent to $\mathcal{T}(m,v) \geq 0$, and the equilibrium payoff is $\max\{m - c(v), 0\} = \max\{\mathcal{T}(m,v), 0\}$.

Part (ii) follows from Definition 2 and Proposition 9. Immediate entry at date 1 occurs if and only if $m \geq c_1(v,\delta)$, equivalently $\mathcal{T}_1(m,v,\delta) \geq 0$. Proposition 9 gives $c_1(v,\delta) \geq c(v)$, so

$$\mathcal{T}_1(m,v,\delta) = m - c_1(v,\delta) \leq m - c(v) = \mathcal{T}(m,v).$$

For part (iii), Proposition 10 implies that, conditional on feasibility under F_v,

$$\mathcal{T}_\infty(m,v) = m - \max\{S - X, -a(v)\} > 0$$

guarantees eventual entry, while

$$\mathcal{T}_\infty(m,v) < 0$$

rules it out. If $m > 0$, then $-a(v) < 0 < m$, so

$$\mathcal{T}_\infty(m,v) > 0 \iff m > S - X,$$

Thus, for optimistic types, Proposition 10 implies that $\mathcal{T}_\infty(m,v) > 0$ guarantees eventual entry, $\mathcal{T}_\infty(m,v) < 0$ rules it out, and $\mathcal{T}_\infty(m,v) = 0$ leaves a knife-edge case in which the entrepreneur is indifferent in the limit and there is no uniform finite-time conclusion. □

Proof of Proposition 11. Part (a) is immediate. Private validation changes only entrepreneur i's own prior mean. The rival still uses the cutoff $c(v)$ from Proposition 2. Therefore, the entrepreneur's equilibrium payoff is

$$\Pi^{\text{priv}}(m, q; v) = \max\{m + q - c(v), 0\},$$

which is (33). On the region where $m + q > c(v)$, the max operator is locally inactive and differentiation yields (34).

For part (b), define

$$G(c; v, z) \equiv c - S + X - \Delta\big(1 - F_v(c - z)\big).$$

Because F_v is continuous and non-decreasing, $G(\cdot; v, z)$ is continuous and strictly increasing in c. Moreover,

$$\lim_{c \to -\infty} G(c; v, z) = -\infty, \qquad \lim_{c \to +\infty} G(c; v, z) = +\infty.$$

Hence, there exists a unique cutoff $c(v, z)$ such that $G(c(v, z); v, z) = 0$, which is exactly (35).

Now suppose F_v is continuously differentiable with density f_v and that $c(v, z) - z$ lies in the interior of the support of F_v. The implicit function theorem applied to $G(c(v, z); v, z) = 0$ gives

$$\frac{\partial c(v, z)}{\partial z} = -\frac{\partial G / \partial z}{\partial G / \partial c} = -\frac{-\Delta f_v(c(v, z) - z)}{1 + \Delta f_v(c(v, z) - z)} = \frac{\Delta f_v(c(v, z) - z)}{1 + \Delta f_v(c(v, z) - z)},$$

which proves (36). Since $f_v \geq 0$, the derivative lies in $[0, 1)$.

The payoff expression (37) follows from Corollary 1 applied to the shifted cutoff $c(v, z)$. On the region where $m + z > c(v, z)$, differentiation yields

$$\frac{\partial \Pi^{\text{pub}}(m, z; v)}{\partial z} = 1 - \frac{\partial c(v, z)}{\partial z} = \frac{1}{1 + \Delta f_v(c(v, z) - z)},$$

which is (38). Comparing (34) and (38) shows that a unit of private validation weakly dominates a unit of public validation at the margin, with strict inequality whenever $f_v(c(v, z) - z) > 0$. $\square$

Proof of Corollary 4. Under Assumption 2, write $a \equiv a(v) = \kappa\sqrt{v}$. If the public-validation cutoff lies in the interior of the support of F_v^z, then

$$F_v(c - z) = \frac{c - z + a}{2a}.$$

Substituting this into (35) gives

$$c = S - X + \Delta\left(1 - \frac{c - z + a}{2a}\right) = S - X + \frac{\Delta}{2} - \frac{\Delta}{2a}c + \frac{\Delta}{2a}z.$$

Since $S - X + \Delta/2 = S - (X + \mu)/2 = S - M$, rearranging yields

$$c\left(1 + \frac{\Delta}{2a}\right) = S - M + \frac{\Delta}{2a}z,$$

and therefore

$$c(v,z) = \frac{2a}{2a+\Delta}(S-M) + \frac{\Delta}{2a+\Delta}z.$$

Using (8), this is exactly (39).

Substituting (39) into (37) gives

$$\Pi^{\text{pub}}(m,z;v) = \max\left\{m - c(v) + \left(1 - \frac{\Delta}{2a+\Delta}\right)z,\, 0\right\} = \max\left\{m - c(v) + \frac{2a}{2a+\Delta}z,\, 0\right\},$$

which is (40). The retained fraction is therefore $2a/(2a+\Delta)$, while the dissipated fraction is $\Delta/(2a+\Delta)$. Differentiating the retained fraction with respect to Δ gives

$$\frac{\partial}{\partial\Delta}\left(\frac{2a}{2a+\Delta}\right) = -\frac{2a}{(2a+\Delta)^2} < 0,$$

so it is strictly decreasing in Δ. □

References

Agrawal, A., Camuffo, A., Gambardella, A., Gans, J. S., Scott, E. L., and Stern, S. (2025). The foundations of Bayesian entrepreneurship. In *Bayesian Entrepreneurship*. MIT Press, Cambridge, MA. Forthcoming.

Athey, S. (2001). Single crossing properties and the existence of pure strategy equilibria in games of incomplete information. *Econometrica*, 69(4):861–889.

Barney, J. (1991). Firm resources and sustained competitive advantage. *Journal of Management*, 17(1):99–120.

Benner, M. J. and Zenger, T. R. (2016). The lemons problem in markets for strategy. *Strategy Science*, 1(2):71–89.

Bolton, P. and Farrell, J. (1990). Decentralization, duplication, and delay. *Journal of Political Economy*, 98(4):803–826.

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

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

Gans, J. S., Scott, E. L., and Stern, S. (2026). Theory-based differences and the origins of heterogeneous priors in entrepreneurship. Working paper, Massachusetts Institute of Technology and University of Toronto.

Gius, L. (2025). Disagreement predicts startup success: Evidence from venture competitions. *Strategy Science*. Forthcoming.

Litov, L. P., Moreton, P., and Zenger, T. R. (2012). Corporate strategy, analyst coverage, and the uniqueness paradox. *Management Science*, 58(10):1797–1815.

Thiel, P. and Masters, B. (2014). *Zero to One: Notes on Startups, or How to Build the Future*. Crown Business, New York.

Zuckerman, E. W. (1999). The categorical imperative: Securities analysts and the illegitimacy discount. *American Journal of Sociology*, 104(5):1398–1438.