NBER WORKING PAPER SERIES

COMPETITIVE EXPOSURE AND ENTREPRENEURIAL EXPERIMENTATION

Joshua S. Gans

Working Paper 35172
http://www.nber.org/papers/w35172

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

Thanks to Refine.ink, ChatGPT 5.5 and Claude Opus 4.7 for valuable research assistance. 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.

Competitive Exposure and Entrepreneurial Experimentation
Joshua S. Gans
NBER Working Paper No. 35172
May 2026
JEL No. O30, O33

ABSTRACT

Entrepreneurs learn by experimenting, but experiment choices are often public. A closed beta, private pilot, or public launch not only generates evidence; it also reveals what kind of entrepreneur would choose that action. We develop a dynamic model in which a founder chooses between stealthy and public experiments while potential entrants infer from both actions and outcomes. Public outcomes are modelled as garblings of the founder's private experimental evidence, so public leakage informs outsiders without giving the founder information beyond the private signal already observed. The key state variable is competitive exposure: the public runway before entry becomes attractive. Exposure is depleted by two forces, leakage burn from public outcomes, and action burn from public inference about experiment choice. This distinction implies that competition can distort experiment design without forcing earlier scale: lower-confidence founders choose stealthier tests, while higher-confidence founders spend exposure to obtain faster private learning through more public tests. Scale is accelerated only when exposure reduces the value of waiting more than it reduces the value of scaling. Finally, exogenous funding gates make observable scale more selective and, therefore, more informative to entrants. The analysis shows that entrepreneurial experimentation is not merely private learning under uncertainty; it is public action under competitive inference.

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

1 Introduction

Consider a founder choosing between a quiet pilot and a public traction launch. A quiet pilot may teach whether customers value the product, but it leaves little evidence for outsiders. A public launch may generate sharper evidence and attract follow-on resources, but it is also a visible act. Rivals can observe not only the customer response but also the fact that the founder was willing to expose the venture publicly. The public launch may reveal a favourable customer response. The decision to launch publicly may already reveal that the founder is optimistic.

In this paper, we ask how entrepreneurs should experiment when experiment choices are public actions. The question is not simply whether the founder should disclose information. In many entrepreneurial settings, information is produced through the same actions that reveal it. A closed beta, conference demonstration, open waitlist, and public usage milestone are all decisions that generate evidence and are themselves interpreted as evidence. Financing rounds can play a similar empirical role when they are publicly visible certificates of quality, but the formal model below treats finance more narrowly: an investor funding standard is an exogenous feasibility gate for scale, not a separate contracting or announcement game.

We build a dynamic model of this problem. The state of the venture is binary: the opportunity is viable or not viable. The entrepreneur privately holds evidence about this state and chooses in each period whether to abandon, run a low-publicity experiment, run a high-publicity experiment, or scale. Experiments produce private signals observed by the entrepreneur and public signals observed by outsiders. The public signal is a garbling of the private signal. This makes the signal structure coherent: outsiders learn from public evidence, while the entrepreneur's posterior is updated by the private signal because the public signal contains no additional information conditional on what the entrepreneur already observes.[1] A potential entrant observes public actions and public signals. The entrant enters once its public belief is high enough. An investor imposes a funding gate: scale is feasible only after the entrepreneur's evidence is sufficiently strong. The gate is used to study certification-through-scale, not investor screening, bargaining, or staged finance.[2]

The main object is *competitive exposure*. Let m be the entrant's public log-odds belief

[1]The path here may not be strictly public but may leak through other relationships, as explored by Cox Pahnke et al. (2015).

[2]The baseline model is deliberately parsimonious, and the appendices show that the same exposure logic survives richer experiment menus, heterogeneous priors, incremental public information, long-lived entrants, and a general lock-in buffer.

that the venture is viable and let $\bar{m}$ be the entry threshold. Exposure (s) is

$$s = \bar{m} - m.$$

This is the entrepreneur's *exposure budget*. It is the public runway before entry becomes attractive. High exposure means that the entrepreneur can afford additional public evidence before inducing entry. Low exposure means that a small public move may make imitation worthwhile.

The main result is an exposure-accounting representation. Exposure is depleted by two forces. The first is *leakage burn*: public outcomes move the entrant's belief directly. The second is *action burn*: public actions reveal which hidden entrepreneur types are willing to take those actions. A public experiment can, therefore, consume exposure even before its outcome is observed. The resulting law of motion is

$$\text{next exposure} = \text{current exposure} - \text{action burn} - \text{leakage burn}.$$

This equation is the spine of the paper.

Understanding the intuition behind this law of motion helps clarify the economic results. A founder with weak private evidence may prefer a quiet test because the main value of experimentation is preserving the option to learn again. A founder with stronger private evidence may accept a public test because faster learning is more valuable and the founder can better afford the exposure cost of being visible. Hence, competition can generate an endogenous stealth-to-traction pattern. The pattern is not assumed through a fixed sequence of stages. It arises because types sort over public actions.

The model delivers three implications. First, public experiment choices generate action burn. This is significant because it means the public state cannot generally be summarised by a scalar posterior over viability: two public histories can leave the entrant with the same belief today yet imply different distributions over hidden entrepreneur types, so the same action tomorrow can trigger different posterior updates and entry risks. Secondly, action burn changes experiment design. Less confident continuation types conserve exposure by choosing low-publicity experiments, while more confident continuation types spend exposure on high-publicity experiments. Thirdly, competition changes entrepreneurial scale-up only on a specific margin. It induces earlier scaling when exposure destroys the value of waiting more than it destroys the value of scaling. It does not mechanically accelerate commercialisation.

Financing creates a further implication within this reduced-form structure. A stricter funding standard means fewer ventures are able to scale. When scaling still occurs, it is therefore a more selective public signal and more informative to entrants. Financing conditions thus affect competition not through a separately modelled funding announcement, but by changing what funding-gated scale tells rivals.

The baseline model uses a short-lived entrant. This benchmark captures settings in which entry opportunities are time-sensitive or capacity-constrained, so a potential rival either acts now or loses the opportunity. Examples include a platform deciding whether to copy a feature during the current product cycle, a corporate rival with a limited budget window, a startup choosing which of several adjacent opportunities to pursue this quarter, or an enterprise channel partner whose adoption decision creates a temporary opening for imitation. In such settings, declining to enter today may mean that the opportunity is reassigned, budgeted away, or preempted by another project. Appendix E studies a long-lived entrant that can wait. The public belief accounting is unchanged, but the now-or-never entry threshold is replaced by an option-value boundary that moves with the entrant's value of waiting.

The paper contributes to several literatures. It relates to persuasion and dynamic information design (Kamenica and Gentzkow, 2011; Bergemann and Morris, 2019; Ely et al., 2015; Au, 2015; Rayo and Segal, 2010), but differs because the entrepreneur does not merely choose what to disclose. The entrepreneur chooses experiments that both generate information and reveal confidence through action. It also relates to strategic experimentation (Bolton and Harris, 1999; Keller et al., 2005), voluntary disclosure and entry (Verrecchia, 1983; Darrough and Stoughton, 1990), disciplined entrepreneurial experimentation (Thomke, 1998, 2003; Camuffo et al., 2020; Agrawal et al., 2021; Gans, 2023), venture finance (Bergemann and Hege, 2005; Gans, 2026), and real-options views of investment (McDonald and Siegel, 1986; Dixit and Pindyck, 1994). In entrepreneurship, the paper is closest in spirit to work that treats experimentation as central to strategy and notes that experimentation can foreclose on commercialisation options (Gans et al., 2019).[3] The distinctive force here is that outsiders learn from entrepreneurial choices as well as from experimental outcomes.

The rest of the paper is organised as follows. Section 2 presents the baseline model. Section 3 develops public inference and exposure accounting. Section 4 studies the no-competition benchmark. Section 5 characterises competitive experiment design. Section 6 studies scaling. Section 7 analyses financing as competitive certification through scale. Section 8 discusses empirical implications, and Section 9 concludes. The appendices contain a closed finite-state equilibrium model, continuous-baseline regularity conditions,

[3] There is also a related literature that examines experimental design when different agents, including the entrepreneur, have heterogeneous priors; see (Van den Steen, 2004; Gans, 2025; Agrawal et al., 2026)

heterogeneous priors, an incremental public-information extension, the long-lived entrant, a general lock-in buffer, and auxiliary results.

2 Baseline Model

2.1 State, Types, and Beliefs

Time is discrete and all agents discount future payoffs by $\delta \in (0,1)$. The unknown state is

$$\omega \in \{0,1\},$$

where $\omega = 1$ denotes a viable opportunity. The entrepreneur's private type is their posterior log odds

$$\ell = \log\left(\frac{\mathbb{P}(\omega = 1 \mid \text{private evidence})}{\mathbb{P}(\omega = 0 \mid \text{private evidence})}\right).$$

Let

$$\mu(\ell) = \frac{e^{\ell}}{1+e^{\ell}}$$

denote the corresponding posterior probability that the opportunity is viable.

The entrant does not observe ℓ. At any public history, the entrant holds a public distribution Π over entrepreneur types. The induced public posterior over viability is

$$\nu(\Pi) = \int \mu(\ell)\,\Pi(d\ell), \qquad m(\Pi) = \operatorname{logit}(\nu(\Pi)). \tag{1}$$

All agents have a common-prior.[4]

2.2 Actions and Timing

In each period, before stopping, the entrepreneur chooses one of four public actions:

$$a \in \{A, L, H, S\}.$$

Here A denotes abandonment, L a low-publicity experiment, H a high-publicity experiment, and S scale. The low-publicity experiment is a stealthy test such as a private pilot.

[4] Appendix C considers heterogeneous priors for the entrepreneur, investor, and entrant following the arguments in the Bayesian entrepreneurship literature (Agrawal et al., 2026). The exposure accounting (below) is unchanged once the entrant aggregates type distributions using its own posterior map.

The high-publicity experiment is a traction-oriented test such as a public launch or visible usage milestone.[5]

The timing within a period is as follows.

(i) The public state is Π, and the entrepreneur chooses A, L, H, or S.

(ii) Outsiders observe the action and update beliefs. If the action-stage public belief reaches the entry threshold, entry occurs immediately. Entry does not erase the entrepreneur's private learning from an experiment already chosen, but it moves the venture to the post-entry continuation value.

(iii) If $j \in \{L, H\}$ was chosen, the experiment produces a private signal z observed by the entrepreneur and a public signal y observed by outsiders. If entry has not already occurred, outsiders update on y and enter if the signal-stage public belief reaches the threshold.

(iv) If no entry occurs, continuation is pre-entry. If entry occurs at either public checkpoint, continuation is post-entry. Abandonment and scale are terminal actions.

Scale is subject to the investor gate specified below. The gate is a feasibility constraint, not a zero payoff from an unavailable action. Each experiment costs a flow amount $c > 0$, common across L and H in the baseline so that the experiment-design trade-off isolates private informativeness and public exposure rather than direct resource-cost differences.

2.3 Experiment Technology

If the entrepreneur chooses experiment $j \in \{L, H\}$, the experiment produces a private signal $z \in \{-1, +1\}$ and a public signal $y \in \{-1, +1\}$. The private signal satisfies

$$\mathbb{P}(z = +1 \mid \omega = 1, j) = \frac{1 + \alpha_j}{2}, \qquad \mathbb{P}(z = +1 \mid \omega = 0, j) = \frac{1 - \alpha_j}{2},$$

where

$$0 \leq \alpha_L < \alpha_H < 1.$$

The parameter α_j governs private learning. A higher α_j means that the experiment is more informative for the entrepreneur.

[5] The two-mode baseline should be read as an adjacent pair on a richer experiment frontier. Appendix A gives a closed finite-state version with a finite ordered experiment menu.

The public signal is a noisy garbling of the private signal. Let

$$0 \leq r_L < r_H < 1, \qquad r_j \leq \alpha_j \quad \text{for } j \in \{L, H\}.$$

When $\alpha_j > 0$, define $\beta_j = r_j/\alpha_j \in [0, 1]$; when $\alpha_j = 0$, set $r_j = \beta_j = 0$. Conditional on z, the public signal is generated by

$$\mathbb{P}(y = z \mid z, j) = \frac{1 + \beta_j}{2}, \qquad \mathbb{P}(y = -z \mid z, j) = \frac{1 - \beta_j}{2}.$$

Equivalently, for $z, y \in \{-1, +1\}$ and $\omega \in \{0, 1\}$,

$$\mathbb{P}(z, y \mid \omega, j) = \frac{1 + z(2\omega - 1)\alpha_j}{2} \frac{1 + yz\beta_j}{2}. \tag{2}$$

The public marginal is

$$\mathbb{P}(y = +1 \mid \omega = 1, j) = \frac{1 + r_j}{2}, \qquad \mathbb{P}(y = +1 \mid \omega = 0, j) = \frac{1 - r_j}{2}. \tag{3}$$

Thus r_j governs public leakage. A high-publicity experiment is more informative privately and more revealing publicly.

Define the log-likelihood increment

$$\varphi(x) = \log\left(\frac{1 + x}{1 - x}\right).$$

After private signal z from experiment j, the entrepreneur's log odds become

$$\ell' = \ell + z\varphi(\alpha_j). \tag{4}$$

The public signal does not enter (4) because it is conditionally independent of the state given the private signal. The entrepreneur may observe the public signal, but it contains no additional information about viability once z is known. Outsiders do not observe z, so y still changes their beliefs.

For later use, define the entrepreneur's subjective probability of a signal pair as

$$q_{zy}(\ell, j) = \mu(\ell)\mathbb{P}(z, y \mid \omega = 1, j) + \big(1 - \mu(\ell)\big)\mathbb{P}(z, y \mid \omega = 0, j). \tag{5}$$

The marginal probability of private signal z is

$$p_z(\ell, j) = \sum_y q_{zy}(\ell, j).$$

Assumption 1 (Ordered experiment pair)**.** *The high-publicity experiment is more informative and more public:*

$$\alpha_H > \alpha_L, \qquad r_H > r_L, \qquad r_j \le \alpha_j \text{ for } j \in \{L, H\}.$$

The last inequality is the coherence restriction implied by the garbling structure: the public signal cannot be more informative than the private signal from which it is generated. The assumption captures a simple trade-off. A public traction launch can teach more than a quiet pilot, but it also creates more evidence for rivals. The assumption does not say that ventures must begin with L and later move to H. The entrepreneur chooses freely between the two modes whenever continuation is chosen.

2.4 Entrant and Competitive Exposure

Entrants in the baseline model are short-lived. In each period, a potential entrant observes the public state and faces a now-or-never entry opportunity. Entry yields payoff $w > 0$ when $\omega = 1$ and $-\lambda < 0$ when $\omega = 0$. Staying out yields zero. The entrant enters when

$$\nu(\Pi)w - \big(1 - \nu(\Pi)\big)\lambda \ge 0.$$

Equivalently, entry occurs when

$$\nu(\Pi) \ge \bar{\nu} \equiv \frac{\lambda}{w + \lambda}, \qquad m(\Pi) \ge \bar{m} \equiv \operatorname{logit}(\bar{\nu}).$$

Definition 1 (Competitive exposure)**.** *At public state Π, competitive exposure is*

$$s(\Pi) = \bar{m} - m(\Pi).$$

Entry is attractive whenever $s(\Pi) \le 0$.

Exposure is the public runway before entry. The short-lived entrant is a benchmark, not a claim that all rivals are short-lived. It isolates the belief movements caused by actions and outcomes. Appendix E shows how the same belief accounting extends when a single entrant can wait.

2.5 Investor Gate

It is assumed that an investor funds scale only when the entrepreneur's private evidence is strong enough:

$$\ell \geq \bar{\ell}_I.$$

A founder below the gate cannot scale; they must abandon or continue experimenting. Thus, passing the gate makes observed scale a certificate of private evidence.[6]

2.6 Scale Payoff and Post-Entry Value

If the entrepreneur scales and the opportunity is viable, monopoly profit is $V > 0$ and duopoly profit is $V - D$, where $D \in (0, V)$. If the opportunity is not viable, scale yields $-K$ with $K > 0$.

The private monopoly-profitability threshold absent entry and financing is

$$\bar{\ell}_M = \text{logit}\left(\frac{K}{V + K}\right).$$[7]

Scale may secure monopoly if the venture has enough residual exposure after the scale action. We model this as a reduced-form lock-in buffer. The interpretation is not that the same short-lived entrant who declined entry before scale waits forever. Rather, after scale there is an execution interval during which operational lock-in, customer contracting, integration, or installed-base advantages may be completed before any effective rival

[6] One way to motivate this assumption (outside the baseline model here) is to consider a heterogeneous-prior interpretation. Suppose the entrepreneur, investor, and entrant begin with priors $p_E > p_I > p_C$, and suppose the investor funds when its posterior reaches ρ_I. Normalising ℓ to the entrepreneur's evidence state gives fixed prior shifts

$$\Delta_I = \text{logit}(p_I) - \text{logit}(p_E), \qquad \Delta_C = \text{logit}(p_C) - \text{logit}(p_E).$$

The investor gate is then

$$\bar{\ell}_I = \text{logit}(\rho_I) - \Delta_I = \text{logit}(\rho_I) + \text{logit}(p_E) - \text{logit}(p_I).$$

At the initial history, $\ell_0 = \text{logit}(p_E)$ and initial exposure is $s_0 = \bar{m} - \text{logit}(p_C)$. The distance from the entrepreneur's initial evidence state to the funding gate is

$$\bar{\ell}_I - \ell_0 = \text{logit}(\rho_I) - \text{logit}(p_I).$$

Thus entrant scepticism determines initial public runway, while investor scepticism determines the distance to the funding gate. Appendix C gives the full heterogeneous-prior notation.

[7] Under the heterogeneous-prior interpretation above, the entrepreneur's initial distance to monopoly profitability is $\bar{\ell}_M - \text{logit}(p_E)$. The three priors, therefore, move three distinct margins: p_C moves public exposure, p_I moves the funding gate, and p_E moves the founder's private stopping margin.

response materialises. Residual exposure is a sufficient statistic for the probability that this interval is long enough. Here, we use the uniform-buffer specification

$$F_B(s) = \min\left\{\frac{\max\{s,0\}}{B}, 1\right\}, \qquad B > 0. \tag{6}$$

Thus, $F_B(s)$ is the probability that scale locks in monopoly when residual exposure is s, while $1 - F_B(s)$ is the probability that post-scale rivalry becomes effective and the venture receives the duopoly payoff. If $s \leq 0$, scale has no immediate lock-in protection. If $0 < s < B$, the lock-in probability rises linearly with exposure. If $s \geq B$, the venture has enough public runway for lock-in with probability one. Appendix F generalises this reduced-form execution buffer.

If observing scale induces the entrant to condition on the scale region $\Gamma_S(\Pi)$, namely the set of private evidence states at public history Π for which the entrepreneur chooses scale, residual exposure (or the *exposure budget*) after scale is

$$s^S(\Pi) = \bar{m} - m^S(\Pi),$$

where $m^S(\Pi)$ is the entrant's action-conditioned log odds after scale. The feasible scale payoff is

$$R^S(\ell, \Pi) = \begin{cases} \mu(\ell)\left[(V - D) + DF_B(s^S(\Pi))\right] - \big(1 - \mu(\ell)\big)K, & \ell \geq \bar{\ell}_I, \\ -\infty, & \ell < \bar{\ell}_I. \end{cases} \tag{7}$$

The $-\infty$ notation simply removes infeasible scale from the maximisation problem.

If entry has already occurred, protected monopoly is no longer available. Let $V^D(\ell)$ be the entrepreneur's post-entry value. The duopoly scale payoff is

$$R^S_D(\ell) = \begin{cases} \mu(\ell)(V - D) - \big(1 - \mu(\ell)\big)K, & \ell \geq \bar{\ell}_I, \\ -\infty, & \ell < \bar{\ell}_I. \end{cases} \tag{8}$$

The post-entry value solves

$$V^D(\ell) = \max\left\{0, R^S_D(\ell), \max_{j \in \{L,H\}} \left[-c + \delta \sum_{z \in \{-1,+1\}} p_z(\ell, j) V^D\big(\ell + z\varphi(\alpha_j)\big)\right]\right\}. \tag{9}$$

Because scale and abandonment are terminal, payoffs are bounded. The Bellman operator in (9) is a contraction on bounded functions, so V^D is well defined and unique.

Scope of the baseline. The baseline deliberately keeps four simplifications in the main text. First, it uses two experiments, L and H, to isolate the stealth-versus-public margin; Appendix A gives the finite ordered-menu version. Second, it uses common-prior notation so that action burn is not obscured by prior-shift bookkeeping; Appendix C restores heterogeneous priors. Third, it uses a short-lived entrant to isolate belief movements; Appendix E adds a long-lived entrant and the resulting boundary drift. Fourth, it uses a uniform lock-in buffer for transparent scale payoffs; Appendix F gives the general buffer and clarifies the post-scale execution interpretation. These simplifications buy a clean statement of the exposure mechanism while preserving the economic margins of experiment design, scale timing, and certification-through-scale.

3 Public Inference and Exposure Accounting

This section develops the central accounting system. The key economic point is that rivals respond not only to public outcomes, but also to what a visible entrepreneurial action reveals about the founder's hidden evidence. For that reason, the entrant's current posterior over viability is generally not enough by itself. Once the action is observed, exposure can change before any experimental outcome is realised.

3.1 Why Public History Matters Beyond the Current Posterior

The reason to care is substantive, not merely technical. The paper asks how public experiments and scale decisions shape rival entry. If outsiders cared only about the venture's current average probability of success, then a single number, the entrant's current log odds $m(\Pi)$, would be enough. But once outsiders also learn from what the entrepreneur chooses next, the composition of hidden entrepreneur types matters. Two public histories can imply the same current belief about viability and yet make the same visible action mean different things. That difference is exactly what generates action burn in the next subsection.

To fix ideas, consider three types

$$\ell_L < \ell_M < \ell_H$$

with
$$\mu(\ell_M) = \frac{1}{2}\mu(\ell_L) + \frac{1}{2}\mu(\ell_H).$$

Now compare two public distributions:

$$\Pi = \frac{1}{2}\delta_{\ell_L} + \frac{1}{2}\delta_{\ell_H}, \qquad \widetilde{\Pi} = \delta_{\ell_M}.$$

They imply the same entrant posterior over viability because $\nu(\Pi) = \nu(\widetilde{\Pi})$. Now suppose scale is chosen by types at least ℓ_M. Under Π, observing scale reveals type ℓ_H, because ℓ_L would not scale. Under $\widetilde{\Pi}$, observing scale reveals type ℓ_M. Thus the current public belief is the same in the two histories, but the inference from the next public action is different.

Proposition 1 (Current public belief alone is not sufficient)**.** *When public actions are informative, two public histories can induce the same entrant posterior m and different action-conditioned posteriors. Hence a scalar public posterior over viability is not generally sufficient for predicting action burn or successor public states.*

Proof. The construction above gives $\nu(\Pi) = \nu(\widetilde{\Pi})$ and hence $m(\Pi) = m(\widetilde{\Pi})$. If scale is chosen by the upper tail $\{\ell : \ell \geq \ell_M\}$, then observing scale conditions Π on $\{\ell_H\}$ and conditions $\widetilde{\Pi}$ on $\{\ell_M\}$. The action-conditioned public log odds are

$$\operatorname{logit}(\mu(\ell_H)) = \ell_H \qquad \text{and} \qquad \operatorname{logit}(\mu(\ell_M)) = \ell_M.$$

Since $\ell_H > \ell_M$, the same scalar public posterior produces different action-conditioned posteriors. The scalar posterior is therefore not sufficient. $\square$

The next result gives the technical counterpart on a finite ordered type space. It says that this is not just an artefact of the three-type example: if one wants a state variable that predicts all monotone action probabilities and Bayes updates, no coarser object than the full public distribution will do.

Proposition 2 (Finite-lattice minimality of the public distribution)**.** *Suppose the entrepreneur type space is the finite ordered lattice*

$$\mathcal{L} = \{\ell_1 < \ell_2 < \cdots < \ell_n\}.$$

A public statistic $\tau(\Pi)$ is monotone-action sufficient if, whenever $\tau(\Pi) = \tau(\widetilde{\Pi})$, it gives the same Bayesian prediction under every monotone upper-tail action rule. That is, for each upper tail

$$T_k = \{\ell_i : i \geq k\},$$

the probability of observing an action chosen exactly by T_k is the same under Π and $\widetilde{\Pi}$, and the action-conditioned posterior after that observation is the same whenever the tail has positive probability. Any monotone-action sufficient statistic identifies the full public distribution over entrepreneur types. Equivalently, if $\tau(\Pi) = \tau(\widetilde{\Pi})$, then $\Pi = \widetilde{\Pi}$.

Proof. For any k, the probability of observing the upper-tail action T_k is $\Pi(T_k)$. Monotone-action sufficiency therefore implies

$$\Pi(T_k) = \widetilde{\Pi}(T_k)$$

for every $k = 1, \ldots, n$. Tail probabilities determine atom probabilities on a finite ordered lattice because

$$\Pi(\{\ell_k\}) = \Pi(T_k) - \Pi(T_{k+1})$$

for $k < n$, and $\Pi(\{\ell_n\}) = \Pi(T_n)$, with the same identities for $\widetilde{\Pi}$. Hence every atom has the same probability under Π and $\widetilde{\Pi}$, so $\Pi = \widetilde{\Pi}$. $\square$

Proposition 1 gives the economic intuition, while Proposition 2 gives the technical minimality result. The substantive message for what follows is simple: competitive exposure depends not only on what outsiders currently believe, but on what they think a visible action says about the founder. That is why the public state below is Π, a distribution over hidden entrepreneur types, rather than a single posterior over viability. The finite-lattice result formalises this claim by showing that, in a monotone signalling environment, any statistic rich enough to recover all action probabilities and post-action beliefs must recover Π itself.

3.2 Action Burn

Fix a public state Π and an equilibrium action rule. Let

$$\Gamma_a(\Pi) = \{\ell : \text{type } \ell \text{ chooses action } a \text{ at } \Pi\}$$

be the set of types choosing action $a \in \{A, L, H, S\}$. If $\Pi(\Gamma_a(\Pi)) > 0$, Bayes' rule gives the action-conditioned public distribution

$$\Pi^a(B \mid \Pi) = \frac{\Pi(B \cap \Gamma_a(\Pi))}{\Pi(\Gamma_a(\Pi))}.$$

If the action is off path, we use the neutral rule $\Pi^a = \Pi$ to make deviations well defined. The substantive results below use on-path action conditioning; different off-path refinements can select different equilibria but do not alter the accounting identities.

The entrant's action-conditioned log odds are

$$m^a(\Pi) = \operatorname{logit}\left(\int \mu(\ell)\, \Pi^a(d\ell \mid \Pi)\right).$$

Definition 2 (Action burn)**.** *The action burn from observing action a at public state* Π *is*

$$\Delta^a(\Pi) = m^a(\Pi) - m(\Pi).$$

Residual exposure after the action is

$$s^a(\Pi) = s(\Pi) - \Delta^a(\Pi) = \bar{m} - m^a(\Pi).$$

Action burn is the exposure cost of the public action itself. If only high types scale, scale burns exposure even before the market observes whether scale succeeds. If only low types abandon, abandonment can move beliefs in the opposite direction.

3.3 Public-Signal Updating

Suppose action $j \in \{L, H\}$ is observed and the action-conditioned public distribution is Π^j. The public signal $y \in \{-1, +1\}$ has likelihood ratio

$$\frac{\mathbb{P}(y \mid \omega = 1, j)}{\mathbb{P}(y \mid \omega = 0, j)} = \exp\{y\varphi(r_j)\}.$$

Therefore, after action inference, the public signal shifts the entrant's log odds by $y\varphi(r_j)$. This term is *leakage burn*. A favourable public signal, $y = +1$, consumes exposure; an unfavourable public signal, $y = -1$, replenishes exposure mechanically by making entry less attractive.

The full public successor distribution is also Bayesian. Let

$$\ell_j^z = \ell + z\varphi(\alpha_j).$$

After observing public signal y, outsiders know that the entrepreneur's type has moved to ℓ_j^z for some unobserved private signal z. Define

$$T_y^j(\Pi^j)(B) = \frac{\int \sum_{z\in\{-1,+1\}} \mathbf{1}\{\ell_j^z \in B\} q_{zy}(\ell, j)\, \Pi^j(d\ell)}{\int \sum_{z\in\{-1,+1\}} q_{zy}(\ell, j)\, \Pi^j(d\ell)}. \tag{10}$$

The denominator is the public probability of signal y after action j.

Lemma 1 (Bayesian consistency of the public successor)**.** *For every on-path action-conditioned public state* Π^j *and public signal* y*,*

$$\operatorname{logit}\left(\int \mu(\ell')\, T_y^j(\Pi^j)(d\ell')\right) = m^j(\Pi) + y\varphi(r_j).$$

Proof. The distribution $T_y^j(\Pi^j)$ is the conditional law of the entrepreneur's next private posterior type after outsiders observe y but not z. By the law of iterated expectations,

$$\int \mu(\ell')\, T_y^j(\Pi^j)(d\ell') = \mathbb{P}(\omega = 1 \mid \Pi^j, y, j).$$

Before observing y, the entrant's log odds are $m^j(\Pi)$. By (3), the likelihood-ratio increment from y is $y\varphi(r_j)$. Hence the displayed equality follows. $\square$

3.4 Exposure Recursion

The two channels combine additively in log-odds space. Suppose the entrepreneur chooses experiment $j \in \{L, H\}$ at public state Π. The entrant first updates from the action. If

$$s(\Pi) - \Delta^j(\Pi) \leq 0,$$

entry occurs immediately at the action stage. If the action does not trigger entry, the public signal is realised and exposure evolves as follows.

Proposition 3 (Exposure budgeting)**.** *After experiment* $j \in \{L, H\}$ *and public signal* $y \in \{-1, +1\}$*, exposure under short-lived entry satisfies*

$$s' = s(\Pi) - \underbrace{\Delta^j(\Pi)}_{\text{action burn}} - \underbrace{y\varphi(r_j)}_{\text{leakage burn}}. \tag{11}$$

The next state is pre-entry if and only if

$$s(\Pi) - \Delta^j(\Pi) > 0 \qquad \text{and} \qquad s' > 0.$$

Proof. After observing action j, the entrant's log odds are $m^j(\Pi)$. Conditional on the public signal, Lemma 1 gives the post-signal log odds

$$m' = m^j(\Pi) + y\varphi(r_j).$$

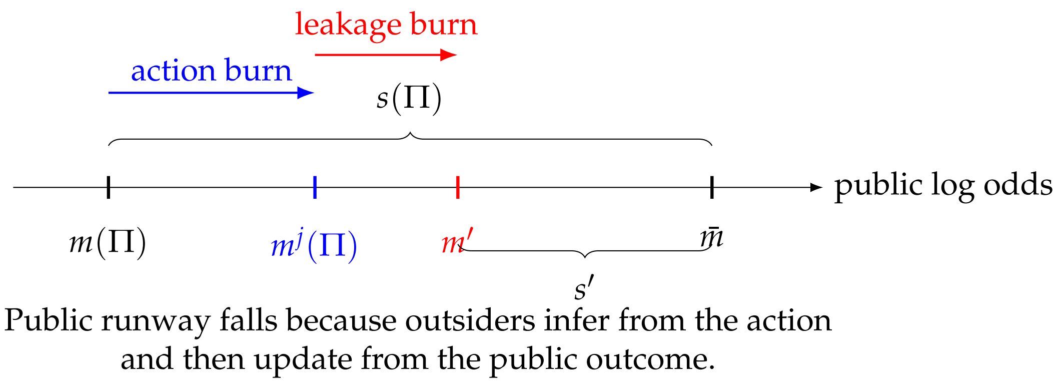

Figure 1: Exposure budgeting with action burn and leakage burn.

Using $s' = \bar{m} - m'$, $s(\Pi) = \bar{m} - m(\Pi)$, and $\Delta^j(\Pi) = m^j(\Pi) - m(\Pi)$ gives (11). The two entry inequalities are exactly the action-stage and signal-stage requirements that the entrant's log odds remain below $\bar{m}$. $\square$

Figure 1 illustrates the recursion. The public action first moves the entrant's belief from $m(\Pi)$ to $m^j(\Pi)$. The public outcome then moves it from $m^j(\Pi)$ to m'. Entry occurs whenever either step reaches the boundary $\bar{m}$.

3.5 Leakage-Only Benchmark

A useful benchmark sets action burn to zero. If actions are uninformative, then

$$\Pi^a(\cdot \mid \Pi) = \Pi \qquad \text{for all } a,$$

so $\Delta^a(\Pi) = 0$. Exposure evolves as

$$s' = s(\Pi) - y\varphi(r_j).$$

Pure leakage models capture public learning from realised outcomes. They miss the fact that experiment choice itself is informative. The difference is important because financing and scale become competitively meaningful only when actions carry information. Table 1 summarises the contrast between the leakage-only benchmark and the full model with action burn.

4 Experimentation Without Competitive Exposure

To isolate the impact of competition, here we identify what the entrepreneur would do if experiments taught the founder but did not expose the venture to entry.

4.1 No-Entry Entrepreneur Problem

Let $V^{OV}(\ell)$ denote the no-competition value. The feasible monopoly scale payoff is

$$R_M^S(\ell) = \begin{cases} \mu(\ell)V - \big(1-\mu(\ell)\big)K, & \ell \geq \bar{\ell}_I, \\ -\infty, & \ell < \bar{\ell}_I. \end{cases}$$

The no-competition value solves

$$V^{OV}(\ell) = \max\left\{0,\, R_M^S(\ell),\, \max_{j\in\{L,H\}}\left[-c + \delta \sum_{z\in\{-1,+1\}} p_z(\ell,j) V^{OV}\big(\ell + z\varphi(\alpha_j)\big)\right]\right\}. \tag{12}$$

Let

$$C^{OV}(\ell) = \max_{j\in\{L,H\}}\left[-c + \delta \sum_z p_z(\ell,j) V^{OV}\big(\ell + z\varphi(\alpha_j)\big)\right].$$

When the scale region is nonempty and upper, define the no-competition scale cutoff

$$\ell^{OV} = \inf\{\ell : R_M^S(\ell) \geq \max\{0, C^{OV}(\ell)\}\}.$$

4.2 Most-Informative Experiment Benchmark

Consider the following regularity condition.

Assumption 2 (Convex no-exposure benchmark)**.** *On the no-competition continuation region, the value function expressed as a function of posterior probability is increasing and convex. The*

Table 1: Leakage-only learning versus action burn

Mechanism	With leakage only	With action burn and leakage
Public outcomes affect entry	Yes	Yes
Experiment choice reveals type	No	Yes
Stealth/public sorting	Not generally	Yes, under single crossing
Financing as certification to rivals	No	Yes, when scale remains on path
Scalar posterior sufficient	Often yes	Generally no

scale region is nonempty and upper.

Assumption 2 is satisfied in the usual Bayesian option-value environment in which the stopping payoff is increasing and convex in posterior beliefs and the investor gate does not introduce a binding nonconvex feasibility kink on the relevant continuation domain. If the investor gate is tighter, the no-competition benchmark remains (12), but the more public experiment, H, need not be optimal at every continuation state. Later comparisons then use the actual benchmark continuation value C^{OV} rather than imposing that H is always selected.

Proposition 4 (No-exposure benchmark)**.** *Absent competitive exposure, the no-competition problem has a unique bounded value V^{OV}. Under Assumption 2, the high-information experiment H is optimal at all continuation states:*

$$C^{OV}(\ell) = -c + \delta \sum_z p_z(\ell, H) V^{OV}\big(\ell + z\varphi(\alpha_H)\big)$$

whenever continuation is chosen. Scale occurs at the option-value cutoff ℓ^{OV}.

Proof. Define the Bellman operator associated with (12) on bounded functions. Terminal payoffs are bounded above by V and abandonment yields zero. For any two bounded functions U and $\widetilde{U}$, the continuation part of the operator differs by at most $\delta\|U - \widetilde{U}\|_\infty$, and the maximum operation preserves this Lipschitz bound. Hence the operator is a contraction and has a unique bounded fixed point.

Because $\alpha_H > \alpha_L$, the private signal from H Blackwell-dominates the private signal from L. Indeed, a signal with accuracy α_L can be generated from a signal with accuracy α_H by flipping the latter with the appropriate independent noise. Under Assumption 2, the expected continuation value of a more informative signal is weakly higher for every prior. Since the two experiments have the same cost in the benchmark, H is optimal whenever continuation is chosen. The scale cutoff statement follows from the maintained assumption that the scale region is nonempty and upper. □

The interpretation is direct. Without exposure costs, the entrepreneur values information for its option value. Since H is more informative than L, it is optimal whenever the continuation value is convex in beliefs. The competitive model changes this comparison by adding action burn and leakage burn.

5 Competitive Experiment Design

The main behavioural prediction of the model concerns experiment design rather than scale timing. Scale is a terminal margin, but experimentation is the repeated margin on which exposure is spent or conserved. The results in this section, therefore, ask when competition changes the composition of experimentation, holding fixed the possibility that scale timing may not move.

The entrepreneur now evaluates experiments by trading off faster private learning against the public exposure generated by actions and outcomes.

5.1 Competitive Bellman Problem

Let $V(\ell, \Pi)$ be the entrepreneur's pre-entry value. For $j \in \{L, H\}$, define

$$\kappa_j(y, \Pi) = \mathbf{1}\{s(\Pi) - \Delta^j(\Pi) > 0\}\mathbf{1}\{s(\Pi) - \Delta^j(\Pi) - y\varphi(r_j) > 0\}.$$

This indicator equals one exactly when the experiment remains pre-entry after both the action-stage and signal-stage public checks. The continuation value from experiment j is

$$Q_j(\ell, \Pi) = -c + \delta \sum_{z,y} q_{zy}(\ell, j)\Big[\kappa_j(y, \Pi) V\Big(\ell + z\varphi(\alpha_j), T^j_y(\Pi^j)\Big) \tag{13}$$

$$+ \big(1 - \kappa_j(y, \Pi)\big) V^D\big(\ell + z\varphi(\alpha_j)\big)\Big]. \tag{14}$$

Here $q_{zy}(\ell, j)$ is given by (5), Π^j is the public distribution after action inference, and $T^j_y(\Pi^j)$ is the Bayesian successor public distribution in (10). If entry occurs after the action or after the public signal, continuation is governed by the post-entry value V^D defined in (9). When action-stage entry occurs, the public signal no longer affects the entry decision, but the experiment chosen for the period still produces the entrepreneur's private signal. The summation over y in this case integrates over the joint signal process; because V^D depends only on the entrepreneur's updated private type, the public signal has no additional payoff relevance after action-stage entry.

The pre-entry value is

$$V(\ell, \Pi) = \max\Big\{0,\ R^S(\ell, \Pi),\ Q_L(\ell, \Pi),\ Q_H(\ell, \Pi)\Big\}. \tag{15}$$

We focus on monotone threshold equilibria in which actions are ordered

$$A < L < H < S.$$

Appendix A gives sufficient conditions for existence in a closed finite-state version of the model. The condition does not require that the entrepreneur is restricted to monotone deviations but that the monotone strategies can be supported as an equilibrium. Given beliefs, the entrepreneur compares all four primitive actions.

For a selected monotone threshold equilibrium with strategy σ, let the selected scale region be

$$\mathcal{C}_S^{\sigma}(\Pi) = \{\ell : \sigma(\ell, \Pi) = S\}.$$

The selected scale cutoff is

$$\theta_S(\Pi) = \inf \mathcal{C}_S^{\sigma}(\Pi),$$

with the convention $\theta_S(\Pi) = +\infty$ if the scale region is empty. Under a fixed tie-breaking rule that selects scale at weak scale optima, this lower boundary can be written as

$$\theta_S(\Pi) = \inf \left\{ \ell : R^S(\ell, \Pi) \geq \max\{0, Q_L(\ell, \Pi), Q_H(\ell, \Pi)\} \right\},$$

with ties interpreted under the maintained equilibrium tie-breaking rule. Thus, scale is chosen by types in the selected upper region $\{\ell : \ell \geq \theta_S(\Pi)\}$, subject to the investor gate.

Remark 1 (Continuous single-crossing regularity)**.** *The continuous log-odds baseline does not assume that signal garbling alone preserves dynamic single crossing. The comparative statics below are therefore stated as conditional results under explicit payoff-gap restrictions: the experiment gap* Δ_{HL} *must be single crossing where Proposition 5 is used, and the scale-stopping gap* G_S *must have the stated monotonicity where scale timing is compared. Appendix B gives a sufficient operator-level route for these restrictions in continuous specifications, while Appendix A proves existence for a closed finite-state analogue. The role of the continuous baseline is to make the Bayesian exposure accounting transparent; the behavioural propositions use the stated regularity conditions rather than claiming that every continuous specification satisfies them automatically.*

5.2 Experiment Sorting

Define the experiment gap

$$\Delta_{HL}(\ell, \Pi) = Q_H(\ell, \Pi) - Q_L(\ell, \Pi).$$

The gap compares the value of spending additional public exposure on the high-publicity experiment with the value of conserving exposure through the low-publicity experiment.

Let

$$\mathcal{C}^{w}(\Pi) = \{\ell : \max\{Q_L(\ell, \Pi), Q_H(\ell, \Pi)\} \geq \max\{0, R^S(\ell, \Pi)\}\}$$

be the weak continuation-candidate region. For a selected equilibrium strategy σ, let

$$\mathcal{C}^{\sigma}(\Pi) = \{\ell : \sigma(\ell, \Pi) \in \{L, H\}\}$$

be the selected continuation region. When actions tie, the maintained tie-breaking rule determines whether a type actually continues, scales, or abandons. The sorting result here is, therefore, a statement about experiment choices among types selected to continue.

Proposition 5 (Experiment sorting)**.** *Fix a public state Π and the beliefs induced by a selected equilibrium strategy σ. Suppose the selected continuation region $\mathcal{C}^{\sigma}(\Pi)$ is an interval and $\Delta_{HL}(\ell, \Pi)$ is weakly increasing in ℓ on an interval containing $\mathcal{C}^{\sigma}(\Pi)$. Then the optimal experiment preference among selected continuation types is single crossing. In particular, there is a domain-normalised cutoff*

$$\ell^{E}(\Pi) = \inf\{\ell \in \mathcal{C}^{\sigma}(\Pi) : \Delta_{HL}(\ell, \Pi) \geq 0\},$$

with $\ell^{E}(\Pi) = +\infty$ if the displayed set is empty, such that selected continuation types below the cutoff strictly prefer L to H, while selected continuation types above the cutoff weakly prefer H to L. If every selected continuation type weakly prefers H, the cutoff is $\inf \mathcal{C}^{\sigma}(\Pi)$, possibly $-\infty$. The equilibrium action follows these preferences with ties assigned according to the maintained tie-breaking rule.

Proof. The set $\{\ell \in \mathcal{C}^{\sigma}(\Pi) : \Delta_{HL}(\ell, \Pi) \geq 0\}$ is an upper set in the interval $\mathcal{C}^{\sigma}(\Pi)$ because Δ_{HL} is weakly increasing. If the set is empty, the high-publicity experiment is never weakly preferred among selected continuation types and the cutoff is $+\infty$. If the set is all of $\mathcal{C}^{\sigma}(\Pi)$, the lower boundary is $\inf \mathcal{C}^{\sigma}(\Pi)$. Otherwise its lower boundary is the cutoff. Types below that boundary strictly prefer L among the two experiments, while types above it weakly prefer H. Equilibrium choices among continuation actions then follow from sequential rationality and the fixed tie-breaking rule. □

The interpretation of Proposition 5 is that experiment design becomes a signalling allocation. Lower continuation types conserve exposure because their main value is keeping the option alive. Higher continuation types spend exposure because they benefit more from faster learning and can better absorb the exposure cost of public traction. Thus, a stealth-to-traction pattern arises endogenously when the experiment gap has the stated single-crossing property.

5.3 Exposure and the Public-Experiment Cutoff

The next result links the accounting system to experiment design. More public runway makes the high-publicity experiment attractive for a larger set of types when the relative value of public experimentation is monotone in usable exposure.

Define an experiment-useful exposure order $\succeq_E$ over public states by

$$\Pi' \succeq_E \Pi \quad \text{if} \quad \Delta_{HL}(\ell, \Pi') \geq \Delta_{HL}(\ell, \Pi) \quad \text{for every relevant continuation type } \ell.$$

This order is deliberately payoff-based because the public state contains more than scalar exposure: it contains the full distribution over hidden entrepreneur types. When the public state affects experiment choice only through scalar exposure, $\Pi' \succeq_E \Pi$ is implied by $s(\Pi') \geq s(\Pi)$.

Proposition 6 (More exposure permits more public experimentation)**.** *Suppose $\Pi' \succeq_E \Pi$, and compare the domain-normalised experiment cutoffs defined in Proposition 5. Also suppose the comparison does not remove the old high-experiment tail from continuation:*

$$\{\ell \in \mathcal{C}^\sigma(\Pi) : \ell \geq \ell^E(\Pi)\} \subseteq \mathcal{C}^\sigma(\Pi').$$

Then the selected public-experiment cutoff is weakly lower at Π':

$$\ell^E(\Pi') \leq \ell^E(\Pi).$$

Proof. By definition of $\succeq_E$, any relevant continuation type that weakly prefers H at Π also weakly prefers H at Π'. The maintained continuation-domain condition ensures that the old high-experiment tail remains in the selected continuation region at Π'. Hence the selected upper set of types weakly preferring H at Π is contained in the corresponding selected upper set at Π', while any lower types newly admitted to continuation lie below the relevant lower boundary and are handled by the domain-normalised cutoff convention. The lower boundary of the latter upper set is weakly below the lower boundary of the former, with $+\infty$ covering the case in which H is never selected. □

Proposition 6 is a conditional monotone-comparative-static result. It gives a direct empirical prediction in environments where the maintained exposure monotonicity is plausible: ventures with more public runway should move to public traction experiments earlier, conditional on continuing. Ventures close to the entry boundary should rely more on quiet tests.

This comparison is not a statement about arbitrary changes in the continuation domain. It applies to the selected continuation intervals relevant for the two public states, using the convention that, when all selected continuation types choose H, the cutoff is $\inf \mathcal{C}^\sigma(\Pi)$ rather than a separate value $-\infty$. Thus, adding lower continuation types who choose L does not mechanically overturn the comparison. In the special case where the public state affects the gap only through scalar exposure and the continuation-domain condition above holds,

$$s(\Pi') \geq s(\Pi) \quad \Rightarrow \quad \ell^E(\Pi') \leq \ell^E(\Pi).$$

5.4 Design Distortion Without Premature Scale

Competition can change how the entrepreneur experiments without making scale earlier. This distinction is important because scale timing is only one behavioural margin: in some parameter regions the scale cutoff is unchanged, while in others competition can delay scale by making protected scale less valuable.

Proposition 7 (A sufficient condition for design distortion without earlier scale)**.** *Fix a public state* Π*. Suppose:*

(i) *H is the no-competition continuation experiment under Proposition 4;*

(ii) *the selected competitive continuation region is an interval and $\Delta_{HL}(\ell, \Pi)$ crosses zero from below to above on that interval;*

(iii) *the competitive scale-stopping gap*

$$G_S(\ell, \Pi) = R^S(\ell, \Pi) - \max\{0, Q_L(\ell, \Pi), Q_H(\ell, \Pi)\}$$

is strictly increasing in ℓ on the relevant interval below and including ℓ^{OV}; and

(iv) $G_S(\ell^{OV}, \Pi) \leq 0$.

Then competition shifts some continuation types from H to L, while the competitive scale cutoff is weakly above the no-competition cutoff ℓ^{OV} whenever the competitive scale region is nonempty. If condition (iv) holds strictly, the competitive scale cutoff is strictly above ℓ^{OV}.

Proof. By condition (i), absent exposure the relevant continuation types choose H. By conditions (ii) and Proposition 5, the competitive problem has a nondegenerate L/H cutoff in the continuation region, so lower continuation types choose L rather than the no-competition experiment H. This is a design distortion.

Condition (iv) says that scale is not strictly optimal at the no-competition scale cutoff in the competitive problem. By strict monotonicity in condition (iii), for every relevant $\ell < \ell^{OV}$ we have $G_S(\ell, \Pi) < G_S(\ell^{OV}, \Pi) \leq 0$. Hence scale is not a weak optimum below ℓ^{OV}, so the tie-breaking rule that selects scale at weak scale optima cannot generate a scale region below ℓ^{OV}. The upper scale region, if nonempty, therefore begins weakly above ℓ^{OV}. If condition (iv) holds strictly, then $G_S(\ell^{OV}, \Pi) < 0$ as well, so scale is not selected at ℓ^{OV} and the cutoff is strictly above ℓ^{OV}. □

Proposition 7 shows why looking only at launch dates, funding dates, or scale dates can miss the main effect of competition. Competitive pressure may first appear in the composition of experiments: quieter pilots, smaller test cohorts, or delayed public metrics. The formal conclusion is deliberately weaker than an unchanged-scale claim: under the stated conditions, competition distorts design without inducing earlier commercial commitment.

Illustration. Table 2 gives a schematic five-type example. In the no-competition benchmark, the continuing types use the high-information experiment and scale begins only at the highest type. With high exposure, competition introduces some caution but leaves the scale cutoff unchanged. With lower exposure, the additional action burn and leakage burn from H make one more continuation type switch to L. The scale cutoff is unchanged; the experiment cutoff moves.

6 Scaling Under Competitive Exposure

Experiment design is the first margin. Scaling is the second. Competition induces premature scaling only when it erodes the value of waiting more than the value of scaling.

Table 2: Design distortion without earlier scale: a schematic five-type example

Environment	Type 1	Type 2	Type 3	Type 4	Type 5
No competition	A	H	H	H	S
High exposure	A	L	H	H	S
Low exposure	A	L	L	H	S

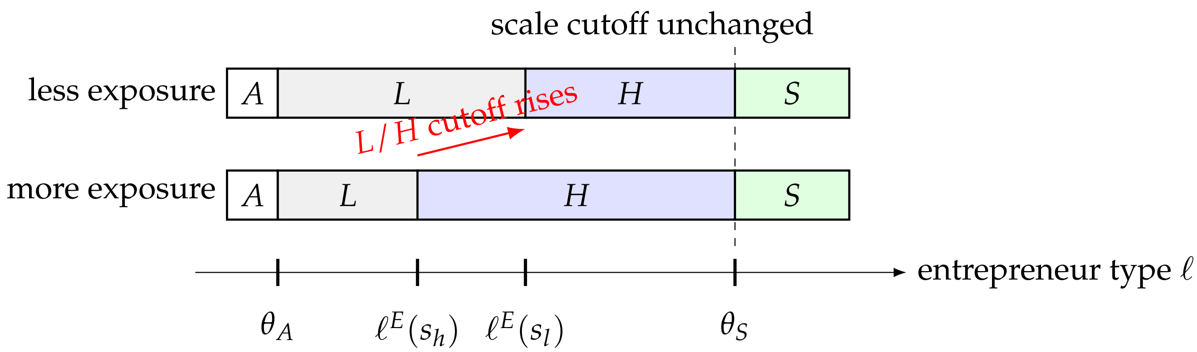

Figure 2: Competition can shift experiment design without moving scale timing. Lower exposure raises the cutoff for public experimentation while the scale cutoff can remain unchanged.

6.1 Competitive Scale Payoff

The competitive scale payoff is $R^S(\ell, \Pi)$ from (7). Relative to the no-competition payoff $R^S_M(\ell)$, scale loses expected value because monopoly is not guaranteed:

$$R^S_M(\ell) - R^S(\ell, \Pi) = \mu(\ell) D \left[1 - F_B(s^S(\Pi))\right]$$

whenever scale is feasible. The loss is larger when residual exposure after scale is low.

Waiting also loses value under competition. Continuation experiments burn exposure through action inference and public outcomes, and may move the entrepreneur into the post-entry value. The premature-scaling question is, therefore, a comparison of two losses.

The premature-scaling force is not that scale hides favourable information from the entrant. Scale is itself a public action, and when scale is chosen only by sufficiently optimistic types, observing scale raises the entrant's belief and consumes exposure. The reason competition can nevertheless accelerate scale is that waiting has also become competitively costly. By continuing to experiment, the entrepreneur preserves the private option value of learning, but does so through actions and outcomes that may reveal enough information to make entry attractive before the venture has locked in a protected position. Thus, the relevant comparison is not simply between scaling now and scaling later. It is between the exposure-adjusted value of scaling and the exposure-adjusted value of waiting.

At the no-competition cutoff ℓ^{OV}, the entrepreneur is just indifferent between scaling and preserving the option to learn. Competition changes both sides of this indifference. Scaling loses value because the action may induce entry and because monopoly is protected only with probability $F_B(s^S(\Pi))$. Waiting loses value because continued experimentation burns exposure through action inference and public outcomes, and may move the venture into the post-entry continuation value. Premature scaling arises precisely when the latter loss is larger than the former: the entrepreneur scales with less private confidence not because scale is uninformative, but because further learning would reveal too much before protection can be secured.

6.2 Premature Scaling Condition

At public state Π, define the competitive scale-stopping gap

$$G_S(\ell, \Pi) = R^S(\ell, \Pi) - \max\{0, Q_L(\ell, \Pi), Q_H(\ell, \Pi)\}.$$

Proposition 8 (Premature scaling)**.** *Fix a public state Π and suppose $G_S(\ell, \Pi)$ is increasing in type. If, at the no-competition scale cutoff ℓ^{OV},*

$$R^S(\ell^{OV}, \Pi) > \max\{0, Q_L(\ell^{OV}, \Pi), Q_H(\ell^{OV}, \Pi)\},$$

then the competitive scale cutoff satisfies

$$\theta_S(\Pi) \leq \ell^{OV}.$$

If the inequality is reversed, this sufficient condition for premature scaling fails; competition may still affect experiment design without implying earlier scale.

Proof. At ℓ^{OV}, the displayed inequality says that competitive scale strictly dominates abandonment and both experiments. Since the scale-stopping gap is increasing in type, every higher type also prefers scale to the best alternative. The equilibrium scale region is therefore an upper set that contains ℓ^{OV}, so its cutoff is weakly below ℓ^{OV}. Reversing the inequality removes this sufficient implication because scale does not beat the best alternative at the benchmark margin. □

When the no-competition cutoff is a scale-continuation indifference point, the condition can be read as

$$\underbrace{C^{OV}(\ell^{OV}) - \max\{Q_L(\ell^{OV}, \Pi), Q_H(\ell^{OV}, \Pi), 0\}}_{\text{loss of waiting value}} > \underbrace{\mu(\ell^{OV}) D\left[1 - F_B(s^S(\Pi))\right]}_{\text{loss of scale value}}.$$

Premature scaling occurs when competition erodes waiting value more than scale value. This is why competition need not accelerate commercialisation. It may instead make the entrepreneur experiment more quietly while leaving the stopping margin unchanged.

7 Financing as Competitive Certification

Financing matters in the model because it changes both feasibility and inference. The formal object is an exogenous gate on scale, not a separate financing announcement or investor-screening game. A higher funding gate directly restricts scale to a smaller set of types. If scale remains on path, it also changes what outsiders infer from observing funding-gated scale.

7.1 Funding Gate

Scale is feasible only if

$$\ell \geq \bar{\ell}_I.$$

In a monotone threshold equilibrium, observed scale is interpreted as coming from the scale region

$$\Gamma_S(\Pi; \bar{\ell}_I) = \{\ell : \ell \geq \theta_S(\Pi; \bar{\ell}_I)\}.$$

The entrant's posterior after scale is therefore increasing in the selected scale cutoff, whenever scale occurs with positive public probability.

There are two forces behind the cutoff comparison. The direct feasibility force is mechanical: scale cannot be chosen below $\bar{\ell}_I$. The indirect equilibrium force is that changing the gate can also change continuation values and action-conditioned beliefs. The proposition below assumes a monotone equilibrium selection for this indirect component, but the direct component can be derived without an equilibrium-selection assumption.

Lemma 2 (Direct effect of a funding gate)**.** *Fix a public state Π and hold fixed the scale interpretation, continuation values, and action-conditioned beliefs. Let $G_S^0(\ell, \Pi)$ be the resulting scale-stopping gap ignoring the feasibility restriction, and suppose $\{\ell : G_S^0(\ell, \Pi) \geq 0\}$ is an upper set. The constrained scale cutoff*

$$\widehat{\theta}_S(\Pi; \bar{\ell}_I) = \inf\{\ell : \ell \geq \bar{\ell}_I \text{ and } G_S^0(\ell, \Pi) \geq 0\}$$

is weakly increasing in $\bar{\ell}_I$.

Proof. Let $\theta^0(\Pi) = \inf\{\ell : G_S^0(\ell, \Pi) \geq 0\}$. Since the unconstrained scale set is an upper set, the constrained cutoff is $\widehat{\theta}_S(\Pi; \bar{\ell}_I) = \max\{\bar{\ell}_I, \theta^0(\Pi)\}$, with the usual $+\infty$ convention if the set is empty. This expression is weakly increasing in the gate. $\square$

7.2 Tighter Gates and Action Burn from Scale

Proposition 9 (Funding-gated scale as competitive certification)**.** *Consider two funding gates $\bar{\ell}_I' < \bar{\ell}_I''$. For any monotone equilibrium selection in which the selected scale cutoff is nondecreasing in the gate,*

$$\theta_S(\Pi; \bar{\ell}_I'') \geq \theta_S(\Pi; \bar{\ell}_I')$$

for every public state Π, if scale remains on path under both gates, then

$$m^S(\Pi; \bar{\ell}_I'') \geq m^S(\Pi; \bar{\ell}_I'), \qquad s^S(\Pi; \bar{\ell}_I'') \leq s^S(\Pi; \bar{\ell}_I').$$

Proof. Raising the funding gate removes scale from the feasible set for additional lower types, and Lemma 2 gives the direct monotone effect holding the continuation environment fixed. The maintained equilibrium-selection condition rules out an indirect feedback effect that lowers the selected scale cutoff. If scale is on path under both gates, observing scale conditions the public type distribution on a weakly higher upper tail under the tighter gate. This higher upper-tail conditioning first-order stochastically dominates the conditional distribution after scale under the looser gate. Since $\mu(\ell)$ is increasing in ℓ, the entrant's expected posterior after scale weakly rises under this FOSD shift. Residual exposure falls because $s^S = \bar{m} - m^S$. □

Proposition 9 says that financing standards affect competition by changing what scale means to rivals when scale is funding-gated. A selective funding environment makes observed scale a stronger certificate of private evidence and can, therefore, invite faster rival response. The claim is about certification-through-scale under an exogenous gate; a separate theory of investor screening or funding announcements would require additional structure.

Corollary 1 (Finance can deter protected entrepreneurship)**.** *If the funding gate is sufficiently demanding that no reachable pre-entry public state supports financeable, profitable, and protected scale, and if $V^D(\ell) = 0$ at all reachable post-entry types, then protected entrepreneurship has zero value at all reachable pre-entry states.*

Proof. Consider the Bellman operator for the protected pre-entry problem on the reachable state space. Let $W(\ell, \Pi) \equiv 0$. At any reachable pre-entry state, abandonment yields zero. By assumption, scale is infeasible, unprofitable, or unprotected, so it contributes no positive protected value. If experimentation triggers entry, the continuation value is $V^D = 0$. If experimentation does not trigger entry, the successor state is again reachable and has candidate value $W = 0$. Since experimentation costs $c > 0$, each continuation action yields at most $-c + \delta \cdot 0 < 0$. Thus the Bellman operator maps the zero function into itself. By contraction, the unique protected-value fixed point is zero on all reachable pre-entry states. $\square$

This corollary is the financing version of the exposure logic. A demanding gate may make scale impossible for moderate types. If only very high types can scale, the scale action may also burn so much exposure that protected scale disappears. The conclusion is about protected entrepreneurship. If duopoly scale or post-entry continuation is profitable, entrepreneurship may continue after entry even though protected monopoly value has been eliminated.

8 Empirical Implications

The model's empirical content lies in separating action burn from leakage burn. Public outcomes matter, but public choices matter too. The separation is not automatic. A comparison between announced and unannounced tests, for example, is informative only if the visibility difference is plausibly independent of project quality or if quality, selection, and disclosure choices are directly controlled. Otherwise the empirical design merely recreates the inference problem the model is about, which is a heroic way to rediscover confounding.

A useful empirical template maps each model object to a distinct source of variation:

$$\Delta^a(\Pi) = m^a(\Pi) - m(\Pi) \quad \text{captures action burn,}$$

while

$$y\varphi(r_j) \quad \text{captures leakage burn from the realised public outcome.}$$

The entry threshold $\bar{m}$ is tied to imitation costs, incumbent capacity, legal protection, and rival payoff parameters. The current public belief $m(\Pi)$ is tied to what outsiders already know about the venture. Exposure $s(\Pi) = \bar{m} - m(\Pi)$ is the gap between those two objects. The funding gate $\bar{\ell}_I$ is a feasibility threshold for scale. The lock-in function $F_B(s)$ maps residual exposure into the probability that post-scale protection is completed before rivalry becomes effective. Empirical proxies should therefore be assigned to these objects separately rather than pooled into a single "visibility" index.

8.1 Identifying Action Burn and Leakage Burn

The clean prediction is not that every more public experiment generates a stronger rival response for every realised outcome. In the model, H has both higher action burn and higher leakage, $r_H > r_L$. A favourable public outcome under H therefore combines action burn with larger positive leakage burn, while an unfavourable public outcome under H combines action burn with a larger negative leakage movement. Comparing public launches and closed betas with similar realised usage can therefore confound the two channels rather than isolate action burn.

A sharper design observes three objects around the same experiment family: the public action, the realised public outcome, and rival response. Action burn is identified from rival response to the action before, or net of, the realised outcome. Leakage burn is identified from rival response to the public outcome conditional on the action. Plausible sources of variation include disclosure-rule changes, platform design changes that make the same test more or less observable to rivals, embargo rules, randomised or rule-based metric publication, or institutional practices that publicise some otherwise similar milestones and not others. These designs are useful only to the extent that they hold fixed the private quality distribution that generated the action.

8.2 Measuring Exposure Rather Than Generic Visibility

Ventures facing high entry risk should remain stealthy longer or use lower-publicity experiments, conditional on continuing. The relevant empirical object is exposure,

$$s(\Pi) = \bar{m} - m(\Pi),$$

not visibility by itself. Measures of incumbent attention or prior public information proxy for $m(\Pi)$. Measures of imitation cost, legal protection, capacity constraints, platform dependence, and product-cycle timing proxy for $\bar{m}$. Observability of usage metrics and test publicity proxy for r_j and the action sets $\Gamma_a(\Pi)$. These are different primitives. Treating all of them as interchangeable measures of "competitive pressure" would blur the distinction between the current public belief, the entry threshold, and the leakage technology.

A credible exposure design therefore varies one component while holding the others fixed. Variation in platform disclosure rules can shift r_j without necessarily changing $\bar{m}$. Variation in incumbent capacity or imitation cost can shift $\bar{m}$ without changing the public signal technology. Variation in prior public information can shift $m(\Pi)$ while leaving the experiment menu unchanged. The prediction is that lower exposure shifts the experiment mix toward quieter tests before it necessarily changes continuation or scale decisions.

8.3 Competition Can Change Composition Without Earlier Scale

Competition need not show up first in entry dates, launch dates, or funding dates. It may appear in quieter pilots, smaller test cohorts, delayed public metrics, less public product iteration, or different sequencing of tests. This is the behavioural implication of Proposition 7: experiment design can shift without earlier scale. In the strongest empirical test, the researcher would show that the L/H composition changes while the scale cutoff, launch date, or funding-gated scale date is unchanged or delayed.

8.4 Funding-Gated Scale as Certification

The model does not contain a separate financing event. It predicts that when scale is feasible only after a stricter funding gate, observed scale is a stronger certificate of private evidence. Funding announcements can be empirical proxies for this mechanism only when they are closely tied to scale feasibility or to a publicly legible funding standard. Possible tests compare rival response to scale or scale-enabling funding across plausibly different funding standards, such as rule-based investor eligibility, time variation in investment committees, or market-level funding tightness, while controlling for capital amount and subsequent deployment. The target is the certification content of passing the gate, not the causal effect of cash alone.

8.5 Bad Public News Can Increase Runway While Being Associated with Lower Private Value

Unfavourable public signals reduce rival entry incentives. In the garbling model, a negative public signal is statistically associated with a negative private signal, although it does not add information to the entrepreneur once the private signal is observed. Thus bad public news can increase public runway while coinciding with lower private continuation value. Appendix G.2 defines the protected-future-experiment set and states the monotonicity condition under which the negative public signal weakly expands that set. This is not a claim that bad news is good for the entrepreneur. It is a claim that public runway and private value are distinct objects.

8.6 Scope of the Modelling Assumptions

The model is intentionally spare. The binary state captures viability rather than the full set of dimensions entrepreneurs learn about, such as market size, defensibility, customer acquisition cost, or regulatory feasibility. The short-lived entrant captures time-sensitive entry opportunities; persistent incumbent monitoring requires the boundary-drift extension in Appendix E. The funding gate is a reduced-form representation of finance, not a full contracting model: it captures the certification content of passing a funding standard through scale, not bargaining, staged contracts, investor monitoring, or a distinct funding announcement. Similarly, the lock-in buffer F_B is a reduced-form post-scale execution technology, not an explicit post-scale entry game. The strategic interpretation is therefore sharpest when observable experiments and scale actions are publicly legible certificates, and when residual exposure is a reasonable sufficient statistic for the chance of securing protection before imitation becomes effective. These restrictions keep the exposure mechanism transparent. They also indicate where the model is most directly applicable: settings in which public experimentation conveys viability evidence to rivals and in which early protected scale is strategically valuable.

9 Conclusion

Entrepreneurial experiments are public actions. They do not merely resolve uncertainty for the founder; they also consume a scarce strategic resource: competitive exposure. Every visible experiment, public launch, usage milestone, or scale decision spends part of an exposure budget. Some of that spending occurs through leakage burn, as outsiders observe outcomes. Some occurs through action burn, as outsiders infer what kind of entrepreneur would choose the action in the first place. The entrepreneur's problem is therefore not simply to learn as quickly as possible. It is to decide when information is worth buying publicly, when it should be bought quietly, and when the remaining exposure budget is better spent on scale.

This paper provides an accounting system for that problem. Competitive exposure measures the public runway before entry becomes attractive. Its law of motion separates the exposure cost of outcomes from the exposure cost of actions. That distinction matters because public choices can reveal confidence even before any market outcome is observed. A public experiment may teach the entrepreneur more, but it also tells rivals that the entrepreneur was willing to be seen. Funding-gated scale may certify the venture's private evidence, but it may also certify the venture to competitors. A scale decision may protect monopoly if enough exposure remains, but it may also reveal that the venture has crossed an important private threshold.

The framework, therefore, offers a way to think about entrepreneurial strategy as exposure budgeting. Founders choose not only what to learn, but how visibly to learn. They choose not only when to disclose evidence, but which actions are worth making observable. Lower-confidence continuation types conserve exposure through quieter tests, while higher-confidence types spend exposure on more informative and more visible experiments. Competitive pressure can therefore distort experiment design without necessarily accelerating scale. Premature scaling arises only on a specific margin: when competition erodes the value of waiting more than it erodes the value of scaling.

This perspective also reframes several familiar entrepreneurial choices. A closed beta, private pilot, public launch, press announcement, scale-enabling venture round, or rapid scale-up can all be understood as different ways of allocating the exposure budget. Some actions preserve runway and keep learning private. Others spend runway in exchange for faster learning, stronger certification, access to resources, or the possibility of locking in a protected position before rivals respond. The relevant question is not whether visibility is good or bad in general. It is whether the marginal public action buys enough private or strategic value to justify the exposure it burns.

There are several natural directions for future research. Empirically, the central challenge is to distinguish leakage burn from action burn. Researchers could study settings in which the publicity of an experiment varies independently from its realised outcome, such as embargoed launches, platform disclosure rules, public versus private pilots, or scale-enabling funding announcements tied to observable standards. Another direction is to measure exposure budgets more directly by separately identifying public beliefs, entry thresholds, public-signal precision, and lock-in conditions. Such work would help determine whether ventures close to the entry boundary shift first toward quieter experiments, as the model predicts, before changing their scale timing.

Theoretically, future work could enrich the model in several ways. One path is to study long-lived entrants with their own experimentation and option values, where the exposure boundary moves endogenously over time. Another is to allow multidimensional entrepreneurial learning, since founders often learn not only whether an opportunity is viable, but also which market, technology, business model, or regulatory path is most promising. A third direction is to embed the financing gate in a fuller contracting model, allowing investor selection, staged finance, monitoring, and syndication to affect both feasibility and competitive certification. Finally, future research could examine exposure budgeting in ecosystems where publicity has positive network effects, making the same public action both a source of demand creation and a signal to rivals.

The broader lesson is that entrepreneurial experimentation is neither a pure learning problem nor a pure disclosure problem. It is public action under competitive inference. Entrepreneurs spend an exposure budget whenever they experiment, publicise, finance, or scale. Understanding how that budget is spent helps explain why some ventures stay quiet, others seek traction loudly, and still others scale before the evidence would justify doing so in a world without rivals.

Appendices

A Closed Finite-State Model and Equilibrium Existence

This appendix records a closed finite-state version of the model and proves existence of monotone Markov equilibria under standard finite monotonicity conditions. The continuous log-odds model in the main text is useful for transparent Bayesian accounting. A literal finite lattice with updates $\ell \mapsto \ell \pm \varphi(\alpha_j)$ is not closed unless one adds projection, truncation, or a primitive finite transition system. This appendix takes the latter route: the finite model is a closed finite-state analogue, not an unclosed restriction of the continuous process.

Let the type space be a finite ordered lattice

$$\mathcal{L} = \{\ell_1 < \ell_2 < \cdots < \ell_n\}.$$

The experiment menu is finite and ordered,

$$\mathcal{A} = \{1, \ldots, J\}, \qquad A < 1 < \cdots < J < S.$$

Higher j denotes a more informative and more public experiment. For each experiment j, each current type ℓ_i, and each private-public signal pair (z, y), let $q_{zy}(\ell_i, j)$ be the subjective probability of that signal pair and let

$$\tau_j^z(\ell_i) \in \mathcal{L}$$

be the next private type. The transition map τ_j^z is assumed to be weakly increasing in ℓ_i for each z and j. This closure condition is the finite-state replacement for the continuous update $\ell + z\varphi(\alpha_j)$.

The public state belongs to a finite set $\mathcal{P} \subset \Delta(\mathcal{L})$. For the fixed-point argument below, closure is imposed for every admissible threshold profile, not only along the equilibrium path selected at the fixed point. Let $\mathcal{T}$ be the finite lattice of monotone threshold profiles. The set $\mathcal{P}$ is assumed to contain the initial public distribution and to satisfy the following all-profile closure property: for every $\Pi \in \mathcal{P}$, every $\theta \in \mathcal{T}$, and every action with positive public probability under θ, the action-conditioned distribution is again in $\mathcal{P}$; and for every on-path experiment and public signal generated under θ, the successor distribution defined below is again in $\mathcal{P}$. If an action has zero probability, the neutral off-path rule leaves the public state at Π.

This is a primitive finite-state closure assumption. Exact Bayesian updating of an unrestricted continuous-belief process will generally generate infinitely many reachable beliefs. Therefore, the theorem applies either to exact finite systems satisfying the all-profile closure property or to a discretised model with an explicitly specified projection $\rho : \Delta(\mathcal{L}) \to \mathcal{P}$, in which case the successor kernel below is replaced by $\rho \circ T^j_y$. It does not apply to an unclosed belief grid with missing continuation states.

For an action rule σ and $\Pi \in \mathcal{P}$, the action-conditioned public distribution is obtained by conditioning on

$$\Gamma^\sigma_a(\Pi) = \{\ell : \sigma(\ell, \Pi) = a\}.$$

The public successor after experiment j and public signal y is the Bayesian kernel

$$T^j_y(\Pi^j)(B) = \frac{\sum_{\ell \in \mathcal{L}} \sum_z \mathbf{1}\{\tau^z_j(\ell) \in B\} q_{zy}(\ell, j) \Pi^j(\ell)}{\sum_{\ell \in \mathcal{L}} \sum_z q_{zy}(\ell, j) \Pi^j(\ell)}. \tag{16}$$

Because τ^z_j maps $\mathcal{L}$ into $\mathcal{L}$ and $\mathcal{P}$ is closed under the kernel, no probability mass leaves the finite state space.

Proposition 2 applies directly to this finite lattice. Thus the public state used below is not merely convenient bookkeeping. In a monotone signalling environment, any statistic rich enough to predict all upper-tail action probabilities and Bayes posteriors must recover Π itself.

A stationary monotone Markov equilibrium consists of a pure strategy $\sigma : \mathcal{L} \times \mathcal{P} \to \{A, 1, \ldots, J, S\}$, Bayes-consistent on-path beliefs, the neutral off-path rule, and thresholds

$$-\infty = \theta_{-1}(\Pi) \leq \theta_0(\Pi) \leq \cdots \leq \theta_J(\Pi) \leq \theta_{J+1}(\Pi) = +\infty$$

such that the action regions are

$$\Gamma_A = \{\ell < \theta_0\}, \qquad \Gamma_j = \{\theta_{j-1} \leq \ell < \theta_j\}, \qquad \Gamma_S = \{\ell \geq \theta_J\}.$$

Scale is feasible only for $\ell \geq \bar{\ell}_I$.

The existence result uses the following finite sufficient conditions.

Assumption 3 (Finite monotonicity and selection)**.** *In the closed finite-state model:*

(i) payoffs are bounded and increasing in type;

(ii) private transition maps and signal likelihoods are monotone in type, and the induced public-belief successor T^j_y is isotone in the public belief under the maintained stochastic order on $\mathcal{P}$;

(iii) action payoff differences have increasing differences in type and action under the order $A < 1 < \cdots < J < S$;

(iv) the scale-stopping gap crosses lower actions from below; and

(v) ties are broken by a fixed monotone rule, so the selected threshold best-response operator is isotone on the finite threshold lattice.

These are finite analogues of the single-crossing and strategic-complementarity conditions in Milgrom and Shannon (1994) and Topkis (1998). They are sufficient, not necessary.

Theorem 1 (Existence of monotone Markov equilibria)**.** *Under the exact closed finite-state structure, bounded payoffs, discounting, and Assumption 3, the least and greatest fixed points of the selected monotone threshold best-response operator exist. Each fixed point induces a stationary monotone Markov perfect Bayesian equilibrium. With an explicitly specified projection, the same argument gives a stationary monotone Markov equilibrium of the projected finite-state game.*

Proof. Because both $\mathcal{L}$ and $\mathcal{P}$ are finite, the threshold profiles form a complete lattice under the componentwise order. Fix any threshold profile in this lattice. By the all-profile closure property, the profile determines action-conditioned beliefs, successor public states, and therefore a Bellman operator on bounded functions over the finite state space $\mathcal{L} \times \mathcal{P}$. Since $\delta < 1$, this Bellman operator is a contraction and has a unique fixed point. Assumption 3 implies that the selected best response is monotone in type and that the selected threshold best-response operator is an isotone self-map on the finite threshold lattice. By Tarski's fixed-point theorem, this self-map has least and greatest fixed points. At any fixed point, the strategy is sequentially rational by construction. In exact finite systems, beliefs are Bayes consistent on path by (16); in projected systems, they are consistent with the specified projected updating rule. Hence the induced assessment is a stationary monotone Markov perfect Bayesian equilibrium for the exact finite system, or the corresponding stationary monotone Markov equilibrium of the projected finite-state game. □

The theorem justifies the threshold representation used in the main text for finite closed versions of the model. It is not used to claim that every equilibrium of the continuous-log-odds model is monotone or that the continuous model always satisfies the finite monotonicity conditions. Its role is narrower: it supplies a closed dynamic environment in which the threshold representation used in the main text is internally consistent. The main-text comparative statics are stated directly under the relevant single-crossing and monotonicity conditions rather than being claimed as unconditional properties of every continuous-log-odds specification.

B Continuous-Baseline Regularity

The continuous baseline is useful because log-odds updating gives transparent exposure accounting. Its dynamic sorting results require more than the garbling primitives alone. This appendix records sufficient regularity conditions that make the maintained single-crossing restrictions explicit.

Fix a public state Π and order public distributions by first-order stochastic dominance over entrepreneur types. A continuous specification has the regularity needed for the experiment-design results if the following operator property holds on the relevant reachable state space:

(i) terminal payoffs 0, $R^S(\ell, \Pi)$, and $R^S_D(\ell)$ are increasing in ℓ, and the scale-stopping gap has the monotonicity stated in the main-text proposition in which it is used;

(ii) the private transition kernels generated by $\ell \mapsto \ell \pm \varphi(\alpha_j)$ are monotone in the monotone-likelihood-ratio order, and the public successor kernels T^j_y are isotone in the public state under the maintained stochastic order;

(iii) for any value function with increasing differences in private type and the relevant public-state order, the continuation payoff difference $Q_H(\ell, \Pi) - Q_L(\ell, \Pi)$ is weakly increasing in ℓ on the selected continuation interval;

(iv) the selected tie-breaking rule is monotone and selects upper action regions when payoff gaps are weakly positive.

Under these conditions, the Bellman operator maps the class of bounded functions satisfying the required increasing-differences property into itself and is a contraction. Its fixed point therefore preserves the payoff-gap monotonicities used in Propositions 5, 6, 7, and 8. In applications, these conditions can be checked directly in a parametric specification or imposed through a finite projection satisfying Assumption 3. They are sufficient, not necessary, and they separate the paper's accounting identities from the additional regularity needed for monotone behavioural predictions.

C Heterogeneous Priors

The main text uses common-prior notation. This appendix restores the heterogeneous-prior environment. Let the entrepreneur, investor, and entrant have priors

$$p_E > p_I > p_C.$$

The entrepreneur's type remains the evidence-state log odds ℓ, normalised to the entrepreneur's posterior. The posterior maps for the three audiences are

$$\mu_E(\ell) = \frac{e^{\ell}}{1+e^{\ell}}, \qquad \mu_I(\ell) = \frac{e^{\ell+\Delta_I}}{1+e^{\ell+\Delta_I}}, \qquad \mu_C(\ell) = \frac{e^{\ell+\Delta_C}}{1+e^{\ell+\Delta_C}},$$

where

$$\Delta_I = \operatorname{logit}(p_I) - \operatorname{logit}(p_E), \qquad \Delta_C = \operatorname{logit}(p_C) - \operatorname{logit}(p_E).$$

The entrant's induced public posterior is

$$\nu_C(\Pi) = \int \mu_C(\ell)\,\Pi(d\ell), \qquad m_C(\Pi) = \operatorname{logit}(\nu_C(\Pi)).$$

All exposure-accounting expressions in the main text remain valid after replacing μ by μ_C in public aggregation:

$$s' = s(\Pi) - \Delta^j(\Pi) - y\varphi(r_j).$$

The reason is that log-likelihood updates are additive once the entrant has formed its action-conditioned posterior.

The investor gate is

$$\ell \geq \bar{\ell}_I \equiv \operatorname{logit}(\rho_I) - \Delta_I = \operatorname{logit}(\rho_I) + \operatorname{logit}(p_E) - \operatorname{logit}(p_I),$$

where ρ_I is the investor's posterior funding threshold. Hence the distance from the initial evidence state $\ell_0 = \operatorname{logit}(p_E)$ to the gate is

$$\bar{\ell}_I - \ell_0 = \operatorname{logit}(\rho_I) - \operatorname{logit}(p_I).$$

The initial exposure is

$$s_0 = \bar{m} - \operatorname{logit}(p_C).$$

Thus p_C affects initial public runway, p_I affects the distance to the funding gate, and p_E affects the entrepreneur's distance to private profitability thresholds.

D Incremental Public Information

The baseline garbling assumption isolates action burn from leakage burn by making the public signal uninformative to the entrepreneur once the private signal is known. Some public tests also generate genuinely incremental market information for the founder. The exposure mechanism does not require excluding that case.

One tractable extension adds a public market signal $x \in \{-1, +1\}$ observed by both the entrepreneur and outsiders, conditionally independent of the private signal z given ω, with

$$\mathbb{P}(x = +1 \mid \omega = 1, j) = \frac{1 + \eta_j}{2}, \qquad \mathbb{P}(x = +1 \mid \omega = 0, j) = \frac{1 - \eta_j}{2}.$$

The entrepreneur's posterior after experiment j becomes

$$\ell' = \ell + z\varphi(\alpha_j) + x\varphi(\eta_j),$$

while the entrant's public log odds after action inference and the public market signal become

$$m' = m^j(\Pi) + x\varphi(\eta_j).$$

Under short-lived entry, the exposure recursion is therefore

$$s' = s(\Pi) - \Delta^j(\Pi) - x\varphi(\eta_j).$$

If the experiment also has a garbled public signal of the baseline kind, the leakage term is the log-likelihood-ratio contribution of the whole public signal vector.

The Bellman equation changes only by replacing the successor private type in Q_j with the updated type that includes the incremental public information. Action burn is unchanged because it is generated by the action-conditioned type distribution before the public outcome is realised. Consequently, Propositions 5–7 continue to hold in the incremental-information model whenever their stated single-crossing and scale-gap conditions hold for the modified continuation values. The incremental signal can make high-publicity experiments more attractive, but it does not remove the exposure cost of public action or the public-outcome leakage term.

E Long-Lived Entrant

The benchmark with a short-lived entrant assumes that entry opportunities are now-or-never. If a single entrant can wait, entry becomes an optimal-stopping problem. Fix an entrepreneur strategy σ. The entrant's current payoff from entry at public state Π is

$$G_C(\Pi) = \nu(\Pi)w - \big(1 - \nu(\Pi)\big)\lambda.$$

Let $O_\sigma(\Pi)$ be the entrant's value of waiting. The entrant's value is

$$X_\sigma(\Pi) = \max\{G_C(\Pi), O_\sigma(\Pi)\},$$

and entry occurs if and only if

$$G_C(\Pi) \geq O_\sigma(\Pi).$$

Equivalently, the entrant uses the state-dependent threshold

$$\nu(\Pi) \geq \bar{v}_\sigma(\Pi) \equiv \frac{\lambda + O_\sigma(\Pi)}{w + \lambda},$$

or, in log odds,

$$m(\Pi) \geq \bar{m}_\sigma(\Pi) \equiv \log\left(\frac{\lambda + O_\sigma(\Pi)}{w - O_\sigma(\Pi)}\right),$$

whenever $O_\sigma(\Pi) < w$.

The public belief state is unchanged. The entrant still needs Π because actions are informative. What changes is the boundary. Strategic exposure is

$$s_\sigma^C(\Pi) = \bar{m}_\sigma(\Pi) - m(\Pi).$$

After experiment j and public signal y, it satisfies

$$s_\sigma^C(\Pi') = s_\sigma^C(\Pi) - \underbrace{\Delta^j(\Pi)}_{\text{action burn}} - \underbrace{y\varphi(r_j)}_{\text{leakage burn}} + \underbrace{\left[\bar{m}_\sigma(\Pi') - \bar{m}_\sigma(\Pi)\right]}_{\text{boundary drift}}.$$

The first two forces are the same as in the main text. The third force appears because the entrant's option value of waiting changes across public states. This is why the benchmark with a short-lived entrant gives the cleanest exposure recursion, while the long-lived entrant adds a state-dependent boundary.

F General Stochastic Lock-In Buffer

The main text uses the uniform-buffer case

$$F_B(s) = \min\{\max\{s, 0\}/B, 1\}.$$

The same analysis applies to any weakly increasing cumulative distribution function F_B on $\mathbb{R}_+$, extended by $F_B(s) = 0$ for $s \le 0$. The scale payoff becomes

$$R^S(\ell, \Pi) = \begin{cases} \mu(\ell)\left[(V - D) + DF_B(s^S(\Pi))\right] - \big(1 - \mu(\ell)\big)K, & \ell \ge \bar{\ell}_I, \\ -\infty, & \ell < \bar{\ell}_I. \end{cases}$$

All results that use only monotonicity of scale value in residual exposure continue to hold. The uniform-buffer case is used in the main text because it gives a transparent marginal value of exposure, $D\mu(\ell)/B$, on the interior interval $0 < s < B$.

G Deterrence, Pivots, and Leakage-Only Benchmark

This appendix records three auxiliary implications that are useful but not central to the main text.

G.1 Deterrence

For any public state Π and any upper-tail scale interpretation T, let $s^S_T(\Pi)$ be residual exposure after scale and $U^S_T(\ell, \Pi)$ the corresponding scale payoff. Define the maximal protected set

$$\mathcal{P}^\star(\Pi) = \bigcup_T \left\{\ell : \ell \ge \bar{\ell}_I,\, U^S_T(\ell, \Pi) \ge 0,\, s^S_T(\Pi) > 0\right\},$$

where the union is over upper tails with positive public mass. If $\mathcal{P}^\star(\Pi)$ is empty at every reachable public state and $V^D = 0$ at all reachable post-entry types, then the protected entrepreneurial value is zero at every reachable state. The proof is immediate from the Bellman equation: scale is infeasible, unprofitable, or unprotected; continuation costs $c > 0$ and leads only to zero continuation value; abandonment yields zero.

G.2 Pivots and Bad Public News

Fix a public state Π and experiment j. Let Π'_+ and Π'_- be the successor public states after positive and negative public signals. The exposure recursion gives

$$s(\Pi'_-) - s(\Pi'_+) = 2\varphi(r_j) \ge 0.$$

Thus, with short-lived entry, bad public news mechanically increases exposure. This does not imply that bad news is privately valuable. In the garbling model, the negative public signal is statistically associated with a worse private signal, but the entrepreneur's posterior is updated by the private signal itself.

To make the protected-experiment claim precise, define the protected-scale set at public state Π as

$$\mathcal{P}^\star(\Pi) = \bigcup_T \left\{ \ell : \ell \geq \bar{\ell}_I,\, U_T^S(\ell, \Pi) \geq 0,\, s_T^S(\Pi) > 0 \right\},$$

where the union is over upper-tail scale interpretations with positive public mass. A type-experiment pair (ℓ, k) is a protected future experiment at public state Π if the experiment does not trigger action-stage entry and has positive probability of leading to a next state in which protected scale is feasible:

$$\mathcal{X}^\star(\Pi) = \left\{ (\ell, k) : \begin{array}{l} k \in \{L, H\}, \quad s(\Pi) - \Delta^k(\Pi) > 0, \\ \exists (z, y) \text{ with } q_{zy}(\ell, k) > 0 \text{ such that} \\ s(\Pi) - \Delta^k(\Pi) - y\varphi(r_k) > 0, \\ \ell + z\varphi(\alpha_k) \in \mathcal{P}^\star(T_y^k(\Pi^k)) \end{array} \right\}.$$

Along the two successor states Π'_+ and Π'_-, suppose protected future feasibility is monotone in the induced exposure comparison:

$$s(\Pi'_-) \geq s(\Pi'_+) \quad \Longrightarrow \quad \mathcal{X}^\star(\Pi'_+) \subseteq \mathcal{X}^\star(\Pi'_-).$$

A sufficient special case is that $\mathcal{X}^\star(\Pi')$ depends on the successor public state only through scalar exposure $s(\Pi')$ and is weakly increasing in that scalar. Under this condition, the negative public signal weakly expands the set of protected future experiments because $s(\Pi'_-) \geq s(\Pi'_+)$. The statement is about public runway, not about the private value of receiving bad news.

G.3 Leakage-Only Benchmark

If actions are uninformative, $\Delta^a(\Pi) = 0$ for all actions. Exposure evolves only through public outcomes:

$$s' = s(\Pi) - y\varphi(r_j).$$

The experiment design and financing-as-certification results depend on action burn and therefore disappear in this benchmark. Deterrence through public outcomes can remain, because leakage alone can still exhaust exposure.

References

AGRAWAL, A., A. CAMUFFO, A. GAMBARDELLA, J. GANS, E. L. SCOTT, AND S. STERN (2026): "The foundations of Bayesian entrepreneurship," in *Bayesian Entrepreneurship*, ed. by A. Agrawal, A. Camuffo, A. Gambardella, J. Gans, E. L. Scott, and S. Stern, MIT Press.

AGRAWAL, A., J. S. GANS, AND S. STERN (2021): "Enabling entrepreneurial choice," *Management Science*, 67, 5510–5524.

AU, P. H. (2015): "Dynamic Information Disclosure," *The RAND Journal of Economics*, 46, 791–823.

BERGEMANN, D. AND U. HEGE (2005): "The Financing of Innovation: Learning and Stopping," *The RAND Journal of Economics*, 36, 719–752.

BERGEMANN, D. AND S. MORRIS (2019): "Information Design: A Unified Perspective," *Journal of Economic Literature*, 57, 44–95.

BOLTON, P. AND C. HARRIS (1999): "Strategic Experimentation," *Econometrica*, 67, 349–374.

CAMUFFO, A., A. CORDOVA, A. GAMBARDELLA, AND C. SPINA (2020): "A Scientific Approach to Entrepreneurial Decision Making: Evidence from a Randomized Control Trial," *Management Science*, 66, 564–586.

COX PAHNKE, E., R. MCDONALD, D. WANG, AND B. HALLEN (2015): "Exposed: Venture capital, competitor ties, and entrepreneurial innovation," *Academy of Management Journal*, 58, 1334–1360.

DARROUGH, M. N. AND N. M. STOUGHTON (1990): "Financial Disclosure Policy in an Entry Game," *Journal of Accounting and Economics*, 12, 219–243.

DIXIT, A. K. AND R. S. PINDYCK (1994): *Investment under Uncertainty*, Princeton, NJ: Princeton University Press.

ELY, J. C., A. FRANKEL, AND E. KAMENICA (2015): "Suspense and Surprise," *Journal of Political Economy*, 123, 215–260.

GANS, J. S. (2023): "Experimental choice and disruptive technologies," *Management Science*, 69, 7044–7058.

——— (2025): "Experiments by "Visionaries"," *Strategy Science*, articles in Advance.

——— (2026): "Entrepreneurial Experimentation Design for Venture Finance," Working paper, Rotman School of Management, University of Toronto.

GANS, J. S., S. STERN, AND J. WU (2019): "Foundations of Entrepreneurial Strategy," *Strategic Management Journal*, 40, 736–756.

KAMENICA, E. AND M. GENTZKOW (2011): "Bayesian Persuasion," *American Economic Review*, 101, 2590–2615.

KELLER, G., S. RADY, AND M. CRIPPS (2005): "Strategic Experimentation with Exponential Bandits," *Econometrica*, 73, 39–68.

MCDONALD, R. L. AND D. R. SIEGEL (1986): "The Value of Waiting to Invest," *The Quarterly Journal of Economics*, 101, 707–727.

MILGROM, P. AND C. SHANNON (1994): "Monotone Comparative Statics," *Econometrica*, 62, 157–180.

RAYO, L. AND I. SEGAL (2010): "Optimal Information Disclosure," *Journal of Political Economy*, 118, 949–987.

THOMKE, S. H. (1998): "Managing Experimentation in the Design of New Products," *Management Science*, 44, 743–762.

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

TOPKIS, D. M. (1998): *Supermodularity and Complementarity*, Princeton, NJ: Princeton University Press.

VAN DEN STEEN, E. (2004): "Rational Overoptimism (and Other Biases)," *American Economic Review*, 94, 1141–1151.

VERRECCHIA, R. E. (1983): "Discretionary Disclosure," *Journal of Accounting and Economics*, 5, 179–194.