

RACING TO RUIN*

Drew Fudenberg
MIT

Andrew Koh
Columbia

July 29, 2026

Abstract

We study R&D competition in the shadow of disaster: advancing the technology frontier raises the risk of permanently ending all firms' payoffs. Under perfect monitoring and common knowledge of rationality, the equilibrium frontier is bounded below by the optimal stopping time of a monopolist, and above by that of a representative firm that persistently but mistakenly believes its rival is about to stop. We then analyze how the frontier is shaped by transparency (speed of monitoring) and trust (belief in the rationality of rival firms).

*Fudenberg: drew.fudenberg@gmail.com; Koh: andrew.koh@columbia.edu. We thank [acknowledgments to be added] and National Science Foundation grant SES-2417162 for financial support.

1 INTRODUCTION

In January 2026 at the World Economic Forum, Demis Hassabis, chief executive of Google DeepMind, was asked whether he would advocate a pause if he knew that every other competitor would pause too: “I think so.” Dario Amodei, chief executive of Anthropic, concurred: “I would prefer that. I think that would be better for the world.” (Hashim and Ford, 2026). This suggests that each firm would prefer to slow down, but only if its rivals do the same. But Amodei went on to add: “it’s very hard to have an enforceable agreement where they slow down and we slow down.” Why exactly is coordination hard? And what would it take?

To answer these questions, we develop a simple model of R&D competition between duopolists in the shadow of disaster:¹ as frontier firms scale the technology, they raise the hazard of an event that permanently drives all firms’ flow payoffs to zero. The hazard is a known function of the firms’ technology levels, and it comes from developing the technology, not from using it. Stopping therefore eliminates future disaster risk, even as firms continue to enjoy the flow payoffs from the technology they have already built; it is as if disaster strikes exactly when the technology crosses some unknown threshold. Only the leader’s state governs whether that threshold is crossed; whether their rival has technology that is nearly as advanced is irrelevant.

Each firm’s flow profit rises in its own technology and falls in that of its rival. When risk is low relative to the profits from scaling, a firm prefers to keep developing regardless of what its rival does: the competitive motive dominates. But when risk is high relative to the marginal profit from scaling, a firm prefers to stop if and only if its rival stops as well, so the problem is coordination, not competition.

We show that even under perfect observability competition still carries the technology frontier to a well-defined threshold, and that degrading the timeliness of information pushes the frontier further out. We also show that there is a tight connection between the equilibrium technology frontier and the technology chosen by a single representative agent that when both firms are active evaluates the prospect of scaling further as if the technology of their competitor is frozen at t . Thus they are persistently overoptimistic about the profits from scaling, and stop only at the time τ^R when even this overoptimism about profits does not outweigh the disaster risk.

When firms observe each other’s behavior perfectly, τ^R is an upper bound on the frontier firm’s technology. The logic is that observability facilitates cooperation: at τ^R , a firm that is contemplating stopping knows that if it stops, their rival will see the stop.

¹The model may also apply to players as nation states e.g., US-China competition.

We then develop a matching lower-bound on the technology frontier across all equilibria, and show that the two bounds are close in several parametric models of competition.

The argument above relied on both *transparency*—firms’ stopping decisions are perfectly and immediately observed—and *trust*: it is common knowledge that firms are rational. We relax each in turn. We first suppose that stopping is observed with delay: a firm learns of its rival’s stop after a random detection lag distributed exponentially with rate ω indexing the transparency/opacity of the environment. When monitoring is very noisy, news travels too slowly for a unilateral stop to be reciprocated in time, so racing forever becomes self-enforcing: there are then equilibria in which neither firm ever stops and disaster arrives with probability one, even though no such equilibrium survives under perfect transparency. When monitoring is sufficiently precise, every equilibrium stops in finite time, but a new temptation appears: each firm would like to stop second, and exits only upon confirmation that the rival has stopped. We show that under a regularity condition on the tail payoffs, the frontier overshoots τ^R by at most $O(1/\omega)$ in expectation.

We next suppose that, in addition to imperfect transparency, firms are also unsure about the rationality of their rival—believing it to be a “crazy type” that never stops in the spirit of [Kreps and Wilson \(1982\)](#). To understand the role of trust, we specialize to a stationary environment with $\tau^R = 0$, so that at each technology state, firms would prefer to stop if their rival also does so. Our main result partitions the trust (belief in their rival’s rationality) into three zones. With low trust, every equilibrium races to ruin: the disaster arrives with probability one. With intermediate trust, immediate stopping and racing to ruin are both equilibria. With high trust, in every equilibrium, the probability that two rational firms race forever vanishes quadratically in the prior odds ratio of rationality. Transparency is double-edged: faster detection makes it cheaper to wait for confirmation that a rival has stopped before stopping oneself instead of stopping unconditionally, so at intermediate trust, increasing transparency can first destroy the early-stopping equilibrium (by making this free-riding deviation attractive) before restoring it as detection becomes fast enough to make stopping self-enforcing.

2 RELATED LITERATURE

Our paper draws on several literatures: dynamic R&D and patent races, wars of attrition and exit games, contests, and the recent economics of AI and existential risk.

R&D races and preemption. A large literature studies dynamic competition to develop a new technology. The early models of [Loury \(1979\)](#), [Lee and Wilde \(1980\)](#), and [Dasgupta](#)

and Stiglitz (1980) relate market structure to the intensity and speed of R&D through a single investment choice; Reinganum (1982) introduces commitment to time-varying open-loop strategies. Fudenberg et al. (1983), Harris and Vickers (1985) and Harris and Vickers (1987) study preemption, leapfrogging, and the discouragement of laggards in models where firms observe and respond to their rivals' progress.² Preemption also drives the timing games of Fudenberg and Tirole (1985). Weeds (2002) and Hopenhayn and Squintani (2011) study strategic delay and preemption when investment is a real option or progress is privately observed. Our model shares the competitive force of these races—flow profits rise in own technology and fall in rivals'—but because the frontier generates a common disaster hazard, the object of interest is how far the frontier is carried before firms stop.

Wars of attrition and exit. When disaster risk is high relative to profits, our firms would prefer to stop if and only if their rivals stop, and under imperfect monitoring the equilibrium becomes a war of attrition. The continuous-time war of attrition originates with Maynard Smith (1974) and is developed by Hendricks et al. (1988), Bulow and Klemperer (1999), and Krishna and Morgan (1997), which relate it to the all-pay auction; the exit dynamics of declining industries in Fudenberg and Tirole (1986) are a close antecedent to our stopping game. Our extension to imperfect monitoring of stopping time is related to other work in which information about others' behavior arrives with delay: Fudenberg et al. (2014) studies repeated games with observation lags, and Daley and Green (2012) analyzes dynamics driven by the stochastic arrival of news. And our extension to “crazy” types builds on the predation model of Kreps and Wilson (1982).

Contests. Our leading example casts flow payoffs as shares of a prize determined by a logit contest success function, following Tullock (1980) and the axiomatic treatment of Skaperdas (1996). We also draw on the broader theory of contests and all-pay auctions, e.g. Baye et al. (1996), Moldovanu and Sela (2001), Siegel (2009), and Konrad (2009).

AI races, existential risk, and regulation. Our motivation comes from the economics of transformative and potentially dangerous technologies. Jones (2024) and Trammell and Aschenbrenner (2024) study the optimal growth–risk tradeoff when scaling raises existential risk, and Ord (2020) argues that a wiser society would proceed more cautiously; our results identify a strategic source of excessive risk-taking that arises even with risk-

²Fudenberg et al. (1983) introduced multistage patent races, and Aghion et al. (2001) embeds step-by-step innovation in a growth setting.

neutral, patient agents. [Armstrong et al. \(2016\)](#) studies a static race between AI developers who choose how much safety to sacrifice for speed, so the issues of observability and dynamic coordination do not arise; [Bueno de Mesquita et al. \(2026\)](#) enrich this static game with market structure. [Koh and Sanguanmoo \(2024\)](#) analyzes robust regulation of multiple AI firms with differing risk aversion but does not model competition.

3 BASELINE MODEL

Time and Technology There are two firms $\mathcal{N} = 1, 2$. Each firm advances its technology at unit speed until it stops, and stopping is irreversible. In the baseline model, firms perfectly and immediately observe their rival’s past stopping behavior, and each has access to private randomization. A strategy specifies a randomized stopping rule after every history; a strategy profile and the private randomization devices induce $\overline{\mathbb{R}}_+$ -valued stopping times $\tau = (\tau_1, \tau_2)$ adapted to the firms’ information. A firm’s time- t stopping decision cannot depend on its rival’s simultaneous time- t decision. Since stopping is irreversible, firm i ’s technology at calendar date s is $s \wedge \tau_i$, and its eventual technology is τ_i .³

Flow Payoffs For a technology profile $\mathbf{t} = (t_1, t_2)$, firm i receives flow payoff $\pi(t_i, t_j)$, where t_i is firm i ’s own technology state (= time spent in the race) and t_j is its rival’s technology state. We impose symmetry across firms: every firm has the same payoff function π , and π treats rivals anonymously. Thus relabelling the firms relabels their payoffs but does not change any payoff level. We write $D_i\pi(t_i, t_j)$ for the derivative with respect to firm i ’s own technology state and $D_j\pi(t_i, t_j)$ for the derivative with respect to firm j ’s technology state.

Disaster risk The *frontier technology scale* of a technology profile $\mathbf{t}$ is $T = \max_i t_i$. There is a risk of a market collapse that arrives at a rate $\lambda : \mathbb{R}_+ \rightarrow \mathbb{R}_+$ where $\lambda(x)$ is the hazard rate when the frontier technology scale is x . Flow payoffs are zero after the collapse. Our leading interpretation of the collapse is the extinction of the industry: after it occurs, no firm is around to enjoy flow payoffs.⁴ This follows the [Jones \(2024\)](#) and [Trammell](#)

³As shown in [Appendix B.1](#), this stopping-time formulation is without loss, because it generates exactly the outcome distributions generated by randomized stopping times in the sense of [Touzi and Vieille \(2002\)](#): adapted, right-continuous, nondecreasing $[0, 1]$ -valued processes recording cumulative stopping probability. Because payoffs are continuous in the stopping dates, the payoffs from stopping just before or just after one’s rival converge to the simultaneous-stopping payoff as the gap between the stopping dates vanishes. Thus, resolving an atom at a common stopping date does not require the extended strategy space of [Fudenberg and Tirole \(1985\)](#).

⁴The starkest examples are ‘loss of control’ scenarios in which beyond an (unknown) technology threshold, AI might be able to rapidly improve itself leading to the extinction of all players ([Bostrom, 2014](#); [Bengio](#)

and Aschenbrenner (2024) analyses of growth under existential risk, and the value-of-life literature (Rosen, 1988; Murphy and Topel, 2006; Hall and Jones, 2007), in which the level of flow utility relative to death pins down how much mortality risk an agent will accept in exchange for higher consumption.

Firms discount the future at a common discount rate of $r > 0$. For a realized stopping profile $\tau \in \overline{\mathbb{R}}_+^2$, define the instantaneous disaster risk at time s by

$$\lambda_s(\tau) := \lambda(s) \cdot \mathbf{1} \left\{ s \leq \max_{j \in \mathcal{N}} \tau_j \right\}.$$

Let $\Lambda_s(\tau) := \int_0^s \lambda_u(\tau) du$ be cumulative risk up to s . Hence, payoffs can be written as

$$U_i(\tau) := \int_0^\infty e^{-rs - \Lambda_s(\tau)} \pi(s \wedge \tau_i, s \wedge \tau_j) ds.$$

To ensure that payoffs are finite and stopping behavior is well-behaved, we will maintain the following assumption on payoffs and risk.

Assumption 1. π and λ satisfy:

- (i) **Increasing hazard:** λ is continuously differentiable and nondecreasing.
- (ii) **Regularity:** π is twice continuously differentiable and strictly positive, strictly increasing in own technology and strictly decreasing in each rival's technology: $D_i \pi > 0$ and, for every $j \neq i$, $D_j \pi < 0$.
- (iii) **Log-supermodularity:** for all firms $i \neq j$, $D_{ij} \log \pi(t_i, t_j) \geq 0$.
- (iv) **Single-crossing:** for every firm i , the map $(t_i, t_j) \mapsto D_i \log \pi(t_i, t_j) - \lambda(t_i)$ is continuous and, for every rival technology t_{-i} , strictly decreasing in t_i over $[t_j, +\infty)$.

(v) **Integrability:**

$$\int_0^\infty e^{-rs} \sup_{t \leq s} \left\{ e^{-\int_0^t \lambda(u) du} \pi(t, 0) \right\} ds < \infty.$$

Assumption 1 (i)–(v) are fairly mild. The increasing hazard assumption is consistent with λ constant, which fits the model where the unknown disaster threshold is exponentially distributed. Regularity (ii) asks only that flow payoffs be smooth, strictly positive, increasing in own technology and decreasing in each rival's. Log-supermodularity (iii) says that when a rival's technology is higher, the marginal payoff from increasing one's

et al., 2024).

own technology is weakly higher; this holds whenever payoffs are additively or multiplicatively separable. Single-crossing (iv) holds if log profits are concave in own technology while the hazard is nondecreasing, and at least one of these conditions is strict. Because λ is nondecreasing. Integrability (v) guarantees that every strategy profile generates a finite expected payoff, uniformly ([Lemma 7](#)).

A representative firm with misspecified beliefs We begin by constructing a single agent with incorrect beliefs about the other firm's strategy. Consider a history at which no disaster has arrived and both firms have continued until t . Fix a generic firm $i \in \mathcal{N}$. The representative firm compares stopping immediately with continuing alone to some $\tau \geq t$, taking the other firm's technology as frozen at t . The continuation payoff from choosing τ is thus

$$R_t(\tau) := \underbrace{\int_t^\tau e^{-\int_t^s (r+\lambda(z)) dz} \pi(s, t) ds}_{\text{Incremental payoff from scaling own tech from } t \text{ to } \tau} + \underbrace{e^{-\int_t^\tau (r+\lambda(z)) dz} \frac{\pi(\tau, t)}{r}}_{\text{Payoff from } \tau \text{ on}}$$

where we use the convention that $R_t(+\infty) := \lim_{\tau \rightarrow +\infty} R_t(\tau)$ for the payoff from never stopping. We define

$$\tau^R := \inf \left\{ t \geq 0 : R_t(t) = \sup_{\tau \geq t} R_t(\tau) \right\}.$$

as the solution to the misspecified firm's problem. The following lemma establishes that there exists a unique solution to the representative agent's stopping problem.

Lemma 1. For each $t \geq 0$, R_t has a unique maximizer $\chi(t)$ on $[t, +\infty) \cup +\infty$ and χ is continuous in t . Moreover, $R_t(t) = \sup_{\tau \geq t} R_t(\tau)$ iff $D_i \log \pi(t, t) \leq \lambda(t)$. Thus, if $\tau^R < +\infty$, then $R_{\tau^R}(\tau^R) = \sup_{\tau \geq \tau^R} R_{\tau^R}(\tau)$ and if $t < \tau^R$, then $D_i \log \pi(t, t) > \lambda(t)$.

[Lemma 1](#) reduces the representative firm's problem to a race between two rates: continuing for an instant scales all future flow profits up at the proportional rate $D_i \log \pi(t, \dots, t)$, while exposing the continuation value to catastrophe at rate $\lambda(t)$. Single-crossing ([Assumption 1](#)(iv)) ensures that, holding the rival's frozen, the former crosses the latter only once from above so τ^R is well defined. The threshold τ^R is the first date at which immediate stopping becomes optimal under the optimistic benchmark in which the rival's technology remains frozen. It will be our benchmark for how far competition scales the technology frontier.

4 PERFECT TRANSPARENCY

We start by analyzing the case of perfect transparency in which firms' actions (stopping decisions) are perfectly and immediately observed. We study subgame perfect equilibria (henceforth just equilibria), and denote equilibrium outcomes with $\text{SPE}(\pi, \lambda)$.

Upper bounds on the frontier We first bound the frontier from above, for every equilibrium, by the representative threshold τ^R .

Proposition 1 (Upper bound). For every $\tau \in \text{SPE}(\pi, \lambda)$, $\max_i \tau_i \leq \tau^R$ almost surely.

A sole survivor at τ^R prefers to stop even if its rival's technology is frozen. Hence, if one of the two active firms stops at τ^R , the remaining firm also stops then; anticipating this, no firm continues past the threshold. More broadly, in the game after crossing τ^R each player prefers to stop as long as the other player does so quickly enough. But because actions are perfectly observed, the logic of backward induction ensures that coordination to stop by τ^R succeeds.⁵ The proof is simple and instructive:

Proof of Proposition 1. If $\tau^R = +\infty$, the claim is immediate. Suppose $\tau^R < +\infty$. First note that, for every firm i , every rival technology level $x_j \leq \tau^R$ for all $j \neq i$, and every $x_i \geq \tau^R$, $D_i \log \pi(x_i, x_j) - \lambda(x_i) \leq D_i \log \pi(x_i, \tau^R) - \lambda(x_i) \leq 0$, where the first inequality uses log-supermodularity and the second uses Lemma 1 at τ^R together with Assumption 1(iv), fixing rivals at τ^R . The last inequality is strict when $x_i > \tau^R$. Hence, after any history at time τ^R in which the rival's technology is weakly below τ^R , a firm's frozen-rival payoff is uniquely maximized by stopping immediately.

We claim that, after any history at time τ^R , no SPE continuation carries the frontier above τ^R with positive probability. Consider first a history at τ^R at which only one firm is active, its rival having stopped at some $y \leq \tau^R$. The active firm's continuation payoff from stopping at $\sigma \geq \tau^R$ is its frozen-rival payoff, which by the preceding paragraph is uniquely maximized by stopping immediately; sequential rationality therefore forces it to stop at τ^R . Consider next a history at τ^R with both firms active, and suppose the continuation carries the frontier above τ^R with positive probability; then some firm i remains active past τ^R with positive probability. Let firm i deviate by stopping at τ^R . After the deviation its rival is a sole survivor whose (frozen) opponent sits at τ^R ; by the one-active-firm step, the rival stops at τ^R immediately. The deviation therefore gives firm i

⁵The argument underpinning Proposition 1 is quite general and can be extended to $n \geq 2$ firms: each firm who stops at τ^R knows that its remaining $n - 1$ rivals would observe it, and each of these rivals know that if they stop, then its remaining $n - 2$ rivals would observe their stop, and so on until a single firm remains and prefers to stop. See also Gale (1995) who studies a dynamic coordination game with irreversible investment.

the payoff of stopping at τ^R with its rival frozen at τ^R . On the equilibrium path, firm i 's payoff is no larger than the payoff from continuing while its rival is frozen at its time- τ^R technology: the rival's actual technology is weakly higher, and frontier risk lasts weakly longer. By the preceding paragraph, that frozen-rival payoff is strictly below the payoff from stopping at τ^R , contradicting sequential rationality. This proves the claim.

If an equilibrium had $\max_{i \in \mathcal{N}} \tau_i > \tau^R$ on a positive-probability event, then conditional on survival to time τ^R the public history at τ^R would have at least one active firm and any stopped firm at technology weakly below τ^R . The claim rules out any continuation from that history with frontier above τ^R , a contradiction. $\square$

Lower bounds on the frontier Note that the equilibrium is not necessarily unique and the equilibrium set need not form a lattice. We will develop simple lower bounds on the frontier technology across all equilibria.

Define

$$\underline{\tau} := \inf \left\{ t : \underbrace{D_i \log \pi(t, 0)}_{\text{Origin index}} \leq \lambda(t) \right\}$$

where the origin index is the marginal payoff to a firm from pushing the technology at t when its competitor's technology are at 0 i.e., it is a monopolist. $\underline{\tau}$ is the time at which a monopolist finds it optimal to stop.

Proposition 2 (Lower bound). For every $\tau \in \text{SPE}(\pi, \lambda)$, $\max_i \tau_i \geq \underline{\tau}$ almost surely.

Proposition 2 states that across all equilibria the technology will advance to at least $\underline{\tau}$. This is because at the moment the last firm contemplates stopping, its rival is weakly behind. Log-supermodularity (**Assumption 1**(iii)) says that rival technology raises the marginal return to the firm's own scaling. Hence the last active firm faces at least as strong an incentive to push the technology as a monopolist would. To see why, let $\chi(t)$ denote the optimal stopping time of the remaining firm once the other firm stops at t . **Assumption 1** (iii) implies that $\chi(t) \geq \underline{\tau}$, so $\max_i \tau_i < \underline{\tau}$ can occur with positive probability only if there is positive probability that firms stop at some common date $t < \underline{\tau}$. But this cannot be an equilibrium: at a history where both firms have strictly positive probability of stopping, either would do strictly better by stopping later.

Propositions 1 and 2 show that the equilibrium frontier lies in the interval $[\underline{\tau}, \tau^R]$. The next lemma gives a simple characterization of when $\underline{\tau} = \tau^R$ so that the interval collapses to a point.

Lemma 2 (Degenerate-bracket test). Suppose $0 < \tau^R < +\infty$. Then

$$D_i \log \pi(\tau^R, 0) \leq \lambda(\tau^R) = D_i \log \pi(\tau^R, \tau^R),$$

with equality on the left if and only if $\underline{\tau} = \tau^R$; if the inequality is strict, then $\underline{\tau} < \tau^R$.

Proof. Continuity of both indices and of λ gives $D_i \log \pi(\tau^R, \tau^R) = \lambda(\tau^R)$ at the first crossing, and log-supermodularity then gives $D_i \log \pi(\tau^R, 0) \leq \lambda(\tau^R)$, whence $\underline{\tau} \leq \tau^R$. Since $t \mapsto D_i \log \pi(t, 0) - \lambda(t)$ is continuous and strictly decreasing (**Assumption 1**(iv), the rival's technology fixed at zero), its first crossing $\underline{\tau}$ falls strictly below τ^R when its value at τ^R is strictly negative. The final claim combines $T \leq \tau^R$ (**Proposition 1**) with $T \geq \underline{\tau}$ (**Proposition 2**). $\square$

Proposition 3 (Pure priority). There exists a pure-strategy subgame perfect equilibrium.

The idea behind **Proposition 3** is that we can fix one firm (say, 1) as the designated first stopper. Along any history at which both firms are active, let firm 1 pick its favorite stopping date, knowing that firm 2 will respond optimally. This is just one among many equilibria; the following proposition develops a simple classification of all of them.

Proposition 4 (Spreading and bunching). For any $\tau \in \text{SPE}(\pi, \lambda)$, almost surely the realized profile is one of:

(i) **Spread:** $\min_i \tau_i < \max_i \tau_i < \tau^R$; or

(ii) **Bunched:** $\min_i \tau_i = \max_i \tau_i = \tau^R$.

Propositions 1 and **2** already rule out the frontier exceeding τ^R and falling short of $\underline{\tau}$. As in the proof of **Proposition 2**, there cannot be positive probability that both firms plan to stop simultaneously at some $t < \tau^R$. A sole survivor at t optimally continues well past t , all the way to $\chi(t) > t$, which is strictly more profitable than stopping at t . The prospect of this gain makes the simultaneous stop self-defeating. And if a firm planned to stop first at some $t < \tau^R$, with the other continuing to τ^R , the early-stopper would do strictly better by waiting until τ^R itself: This induces its rival to either stop before τ^R , gaining higher own technology, a weaker rival, and less accumulated hazard, or stops simultaneously with its rival at the ceiling, which is also an improvement. Together these two exclusions and the bounds from **Propositions 1** and **2**, leave exactly the taxonomy: either firms stop at different times with the frontier strictly inside the bracket, or they stop together at the ceiling.

Finally, we develop some simple comparative statics for how the ceiling τ^R moves with the primitives. Write $\tau^R(\pi, \lambda)$ for the representative threshold associated with primitives (π, λ) . The key observation is an asymmetry in how π and λ enter the representative firm's stopping rule (**Lemma 1**): the payoff function matters only through the *diagonal index* $t \mapsto D_i \log \pi(t, t)$, whereas the hazard enters only as the threshold this index must fall below.

Proposition 5 (Comparative statics). For (π, λ) and $(\tilde{\pi}, \tilde{\lambda})$ that satisfy **Assumption 1**:

- (i) If $D_i \log \tilde{\pi}(t, t) \geq D_i \log \pi(t, t)$ for every $t \geq 0$, then $\tau^R(\tilde{\pi}, \lambda) \geq \tau^R(\pi, \lambda)$.
If $D_i \log \tilde{\pi}(t, 0) \geq D_i \log \pi(t, 0)$ for every $t \geq 0$ then $\underline{\tau}(\tilde{\pi}, \lambda) \geq \underline{\tau}(\pi, \lambda)$.
- (ii) If $\tilde{\lambda}(t) \geq \lambda(t)$ for every $t \geq 0$, then $\tau^R(\pi, \tilde{\lambda}) \leq \tau^R(\pi, \lambda)$ and $\underline{\tau}(\pi, \tilde{\lambda}) \leq \underline{\tau}(\pi, \lambda)$.

Proposition 5 shows how the frontier's ceiling and floor respond to the primitives. The ceiling τ^R is set by the diagonal index $D_i \log \pi(t, t)$, the neck-and-neck racing incentive: it strengthens with the sensitivity of payoffs to relative position (the contest weight γ , the complementarity B), pushing the ceiling out. The floor $\underline{\tau}$ is set by the origin index $D_i \log \pi(t, 0)$, the marginal value of scaling for a firm whose rival is stuck at the origin, so it moves with the technology's standalone profitability rather than with rivalry. Raising the hazard (λ_0, κ) pulls both bounds in.

Proof of Proposition 5. By **Lemma 1** and the definition of $\underline{\tau}$, both thresholds are first-crossing times of an index with the hazard: $\tau^R(\pi, \lambda) = \inf\{t \geq 0 : \partial_i \log \pi(t, t) \leq \lambda(t)\}$; and $\underline{\tau}(\pi, \lambda) = \inf\{t \geq 0 : \partial_i \log \pi(t, 0) \leq \lambda(t)\}$. All four claims follow from the same set inclusion. If an index is replaced by a pointwise weakly larger one, any t in the new crossing set belongs a fortiori to the old one, so the crossing set shrinks and its infimum weakly rises; this gives both statements in (i), with the diagonal index for τ^R and the origin index for $\underline{\tau}$. If instead the hazard is replaced by a pointwise weakly larger one, any t in the old crossing set belongs to the new one, so the crossing set grows and its infimum weakly falls; this gives both statements in (ii). $\square$

4.1 Illustrative Examples. We now illustrate how the upper- and lower-bounds behave across a number of parametric examples. These examples specialize to an exponential hazard rate, $\lambda(t) = \lambda_0 e^{\kappa t}$.

Example 1 (Score contests). Firm i earns a share $s_i(\mathbf{x}) = e^{\gamma x_i} / (e^{\gamma x_1} + e^{\gamma x_2})$, with $\gamma \geq 0$, of an own-technology prize $A_0 e^{\beta x_i}$, so its flow payoff is $\pi(x_i, x_j) = A_0 e^{\beta x_i} s_i(\mathbf{x})$. At a

symmetric profile, the marginal log-payoff from scaling is $D_i \log \pi(t, t) = \beta + \frac{\gamma}{2}$, while with the rival at the origin it is $D_i \log \pi(t, 0) = \beta + \frac{\gamma}{e^{\gamma t} + 1}$. Suppose first that $\kappa > 0$. Then

$$\tau^R = \left[\frac{1}{\kappa} \log \frac{\beta + \gamma/2}{\lambda_0} \right]^+, \quad \underline{\tau} = \inf \left\{ t \geq 0 : \beta + \frac{\gamma}{e^{\gamma t} + 1} \leq \lambda_0 e^{\kappa t} \right\}.$$

If $\lambda_0 \geq \beta + \gamma/2$, both thresholds equal zero. If $\lambda_0 < \beta + \gamma/2$, the ceiling is interior; the two thresholds coincide when $\gamma = 0$, so that the prize is split evenly regardless of play, and satisfy $\underline{\tau} < \tau^R$ when $\gamma > 0$. If instead $\kappa = 0$, the diagonal index is constant: the ceiling is zero when $\beta + \gamma/2 \leq \lambda_0$ and infinite otherwise, and the $1/\kappa$ formula does not apply.

Figure 1 plots both thresholds as the curvature of the prize and the level of the hazard vary, together with the pure priority equilibrium at selected parameter values: it bunches at the ceiling when the contest is shallow and spreads once the contest is steep. **Figure 2** shows why: the left panel varies the share sensitivity γ . The bracket $[\underline{\tau}, \tau^R]$ collapses at $\gamma = 0$, where the prize is split evenly, and widens with γ : the ceiling rises with rivalry, while the floor first rises (encouragement) and then falls (discouragement: a laggard's share gain vanishes when the contest is steep). The right panel varies the hazard level λ_0 . Both thresholds fall and reach zero together at $\lambda_0 = \beta + \gamma/2$. Dots mark the pure priority equilibrium at selected parameter values: it bunches at the ceiling for moderate curvature, and the two stops separate strictly inside the bracket once the contest is steep ($\gamma \gtrsim 1.8$).




Figure 1: Score-contest thresholds; parameters: $\beta = 0.3, \kappa = 0.4, r = 0.1$; the left panel fixes $\lambda_0 = 0.2$, the right panel fixes $\gamma = 1$




Figure 2: Follower's reaction curve; parameters: $\beta = 0.1$, $\lambda_0 = 0.2$, $\kappa = 0.4$, $r = 0.1$

Figure 2 illustrates both bunching and spreading regimes for the score contest. It shows the follower's reaction curve (red) and the first stopper's isopayoff curve in the plane of stopping dates. The first stopper's payoff rises toward the lower right: delaying its own stop raises its flow profit, while the follower's delay erodes it and prolongs the hazard. The priority equilibrium maximizes the first stopper's payoff along the reaction curve. At $\gamma = 2$ (right panel) the optimum is an interior tangency: the isopayoff curve through the equilibrium touches the reaction curve at the dot and lies below it elsewhere. At $\gamma = 1$ (left panel) the payoff rises along the entire reaction curve, so the optimum is the corner (τ^R, τ^R) where the reaction curve meets the 45° line, and the firms bunch.

Example 2 (Cournot competition). Let inverse demand be $P = a - b(q_1 + q_2)$, and let capability lower marginal cost, $c(x) = c_0 e^{-\delta x}$. Assume $a > 2c_0$; since $c(x) \in (0, c_0]$, every technology profile then has $M_i = a - 2c(x_i) + c(x_j) \geq a - 2c_0 > 0$, so the static Cournot quantities are interior and the log-profit index below is defined globally. In the interior Cournot equilibrium,

$$\pi_i(x) = \frac{M_i^2}{9b}, \quad M_i := a - 2c(x_i) + c(x_j),$$

so

$$D_i \log \pi_i(x) = \frac{4\delta c(x_i)}{a - 2c(x_i) + c(x_j)}.$$

The ceiling and the floor are therefore

$$\tau^R = \inf \left\{ t \geq 0 : \frac{4\delta c_0 e^{-\delta t}}{a - c_0 e^{-\delta t}} \leq \lambda_0 e^{\kappa t} \right\}, \quad \underline{\tau} = \inf \left\{ t \geq 0 : \frac{4\delta c_0 e^{-\delta t}}{a + c_0 - 2c_0 e^{-\delta t}} \leq \lambda_0 e^{\kappa t} \right\}.$$

For $c_0 > 0$, the two thresholds coincide when $4\delta c_0 \leq \lambda_0(a - c_0)$, in which case both equal zero; when $4\delta c_0 > \lambda_0(a - c_0)$, both crossings are positive and $\underline{\tau} < \tau^R$. When $c_0 = 0$, the first-crossing definitions give $\underline{\tau} = \tau^R = 0$, with the logarithmic formulas interpreted by continuity. **Figure 3** illustrates these numerically. The left panel varies the hazard growth κ : at $\kappa = 0$, $\underline{\tau} = 0.511$ and $\tau^R = 0.588$. As κ rises, both crossings move earlier and the bracket tightens. A larger market size a also moves both bounds earlier: with high baseline demand, firms earn substantial margins, so the proportional gain from further cost reduction is smaller. The right panel varies the market size a ; both thresholds reach zero together at $a = c_0(4\delta + \lambda_0)/\lambda_0$. Dots mark the pure priority equilibrium at selected parameter values: at this calibration it bunches at the ceiling, so the two dots coincide on τ^R .




Figure 3: Cournot bracket; parameters: $c_0 = 1$, $\delta = 1$, $\lambda_0 = 0.5$; the left panel fixes $a = 5$, the right panel fixes $\kappa = 0$

Example 3 (Log-separable payoffs). Finally, suppose log-flow payoffs are additively separable with one interaction term. Let $\psi(t) := 1 - e^{-t}$ and take

$$\log \pi(t_i, t_j) = A\psi(t_i) - C\psi(t_j) + B\psi(t_i)\psi(t_j),$$

with $A > 0$, $C \geq B \geq 0$, and $C > 0$. The parameter B controls the complementarity between own and rival technology; C governs how much the rival's progress erodes a

firm's profits. The origin index and its diagonal counterpart are $D_i \log \pi(t, 0) = A\psi'(t)$, and $D_i \log \pi(t, t) = \psi'(t)(A + B\psi(t))$, with gap $B\psi'(t)\psi(t)$, so the two indices coincide when $B = 0$. The hazard-free parts of **Assumption 1** are immediate: $D_i \log \pi = \psi'(t_i)[A + B\psi(t_j)] > 0$; $D_j \log \pi = \psi'(t_j)[B\psi(t_i) - C] < 0$, since $\psi < 1$ and $C \geq B$; $D_{ij} \log \pi = B\psi'(t_i)\psi'(t_j) \geq 0$; and since $\pi(t, t)$ is bounded above the integrability condition (v) holds. The slope of the follower's response to a first stop is governed by $D_{ij} \log \pi$, so under strong complementarity a marginal delay of the first stop drags the follower's stop outward by a comparable amount, which can pull the first stopper's optimum strictly below the ceiling.




Figure 4: Separable bracket; parameters: $A = 2$, $B = \frac{1}{2}$, $r = 0.1$, $C = 1$

Propositions 1 and **2** bracket the equilibrium frontier $T := \max_i \tau_i$ within $[\underline{\tau}, \tau^R]$ in every equilibrium, so competition carries the frontier all the way to the representative threshold precisely when the bracket is degenerate: $\underline{\tau} = \tau^R$. By **Lemma 2**, $\underline{\tau} < \tau^R$ when the origin index falls strictly below the hazard at τ^R ; here the origin index differs from the diagonal index by the gap, so for $0 < \tau^R < +\infty$,

$$\underline{\tau} < \tau^R \iff B\psi'(\tau^R)\psi(\tau^R) > 0 \iff B > 0,$$

using $\psi' > 0$ and $\psi(\tau^R) > 0$: the bracket is degenerate iff $B = 0$, and the magnitude of the complementarity is irrelevant. With an exponential hazard both thresholds are explicit. We have $\partial_i \log \pi(t, t) = e^{-t}[A + B(1 - e^{-t})]$ and $\partial_i \log \pi(t, 0) = Ae^{-t}$, so for $\lambda_0 \in (0, A)$ the floor solves $Ae^{-t} = \lambda_0 e^{\kappa t}$, giving $\underline{\tau} = \log(A/\lambda_0)/(1 + \kappa)$, while the ceiling solves $e^{-t}[A + B(1 - e^{-t})] = \lambda_0 e^{\kappa t}$; for a constant hazard ($\kappa = 0$) the latter is a quadratic in

e^{-t} , and for $\kappa > 0$ it is computed numerically. [Figure 4](#) plots both thresholds. The left panel keeps the hazard constant and varies the risk level λ_0 . Both thresholds fall as the technology becomes more dangerous, with $\underline{\tau} = \log(A/\lambda_0)$, and the bracket narrows from width $\log(1 + B/A)$ as $\lambda_0 \downarrow 0$ to zero as $\lambda_0 \uparrow A$. The right panel considers exponential hazard with $\lambda_0 = 1$, varying the growth rate κ ; both thresholds fall as risk accelerates. Dots mark the pure priority equilibrium at selected parameter values.

5 IMPERFECT TRANSPARENCY

We now relax perfect observability. If firm i stops at τ_i , its rival receives a conclusive signal after an independent delay $E_i \sim \text{Exp}(\omega)$. Thus Y_t^i jumps from zero to one at $\tau_i + E_i$. A signal reveals that the firm has stopped but not when; silence is therefore ambiguous. The news rate ω measures transparency, with mean detection lag $1/\omega$; the perfectly observed benchmark is the limit $\omega \rightarrow \infty$.

Firm i 's filtration is, as before, generated by its private randomization device $\mathbf{Z}_i^{*6}$ and the signals it receives: $\mathcal{F}_t^i = \sigma(\mathbf{Z}_i^*, (Y_s^j)_{s \leq t, j \neq i})$. As before, a strategy specifies, at each of the firm's histories, a distribution over future stopping dates, and an outcome is a profile τ in which each τ_i is $\mathcal{F}^i$ -adapted. The formal construction ([Appendix B.2](#)) mirrors that of the perfectly observed game with detections in place of stops: a strategy is a family of *plans* indexed by the firm's *signal record*—which rivals have been detected, and when—each plan a measurable map from the device to a stopping date. With two firms, a strategy therefore reduces to a pair: a *concession plan*, the (possibly mixed) date at which the firm stops on its no-signal history, and a *response plan*, the date at which it stops following a detection.

We study perfect Bayesian equilibria and write $\text{PBE}(\omega)$ for the set of equilibrium outcomes. Beliefs are largely pinned down by Bayes' rule: on no-signal histories, the posterior over whether and when the rival has stopped is determined by the rival's strategy and the absence of a signal; on detection histories, the signal is conclusive, so firm i assigns probability one to firm j having stopped in the past. The only residual freedom is the posterior over the exact stopping date of a rival detected off path. Before τ^R , these beliefs may affect the optimal response because they determine the rival's frozen technology, but after it [Assumption 2](#) implies that immediate stopping is optimal for every feasible stopping-date belief.

We henceforth maintain the following condition that guarantees that after τ^R , a firm

⁶This consists of two iid uniform random variables. The first is drawn at time 0, and the second is drawn upon *detection* of the rival's stop.

that discovers that its rival has already stopped also prefers to stop.

Assumption 2 (Diagonal single-crossing). The map $t \mapsto D_i \log \pi(t, t) - \lambda(t)$ crosses zero at most once, from above; equivalently, if $D_i \log \pi(t, t) \leq \lambda(t)$ then $D_i \log \pi(s, s) \leq \lambda(s)$ for every $s \geq t$.

Proposition 6 (Existence under imperfect transparency). Suppose that $\tau^R < +\infty$. Then for every news rate $\omega > 0$, $\text{PBE}(\omega) \neq \emptyset$.

Let $\alpha(s) := e^{-rs - \int_0^s \lambda(u) du}$ be the discount-survival kernel, so $\alpha'(s)/\alpha(s) = -(r + \lambda(s))$ and α is continuously differentiable, strictly positive, strictly decreasing, with $\alpha(s) \downarrow 0$.

Definition 1 (Race and stop values). Fix a symmetric no-news history at which both firms are active at technology t . The *race value* is the payoff from both racing forever,

$$V_R(t) := \frac{1}{\alpha(t)} \int_t^\infty \alpha(s) \pi(s, s) ds = \int_t^\infty e^{-\int_t^s (r + \lambda(u)) du} \pi(s, s) ds.$$

The *stop value* $V_S(t, \omega)$ is the payoff from stopping at t while the rival continues until it detects the stop and then chooses its optimal stopping date against a rival frozen at t . With detection date $a := t + \tilde{D}$ and the rival's frozen-rival optimum $\sigma^*(a) := \inf\{u \geq a : D_i \log \pi(u, t) \leq \lambda(u)\}$,

$$V_S(t, \omega) := \mathbb{E}_{\tilde{D}} \left[\int_t^{\sigma^*(a)} \frac{\alpha(s)}{\alpha(t)} \pi(t, s) ds + \frac{\alpha(\sigma^*(a))}{\alpha(t)} \frac{\pi(t, \sigma^*(a))}{r} \right].$$

Write $\Delta(t, \omega) := V_S(t, \omega) - V_R(t)$ for the difference between these values.

The race and stop values capture stopping incentives. Unlike the case with perfect observability, a firm that contemplates stopping knows that it will only be observed at a mean lag $1/\omega$. In the interim, its technology is frozen while its rival keeps scaling and catastrophic risks continue to accrue.




Figure 5: Stopping first under delayed detection.

For a stop at $t \geq \tau^R$, **Assumption 2** implies that the rival stops immediately upon detection ($\sigma^*(a) = a$); **Figure 5** illustrates this post-ceiling case: firm i stops at t and its technology freezes; firm j , seeing no signal, races on. During the detection lag (middle band) firm i is doubly exposed: its flow $\pi(t, s)$ is eroded by the rival's still-rising technology, and the disaster hazard remains alive because the frontier is still moving. Upon detection at $a = t + X$ the rival stops, the hazard dies, and firm i collects the safe perpetuity $\pi(t, a)/r$. The race value $V_R(t)$ is the counterfactual in which both firms ride the diagonal forever (dotted extension). Transparency shortens the middle band in expectation: as $\omega \rightarrow \infty$ it collapses and $V_S(t, \omega) \rightarrow \pi(t, t)/r$ (**Lemma 5**).

Lemma 3 (Single-agent optimum). Fix a rival frozen at c and a detection date $a \geq c$. The value of stopping at $\sigma \geq a$,

$$\int_a^\sigma \frac{\alpha(s)}{\alpha(a)} \pi(s, c) \, ds + \frac{\alpha(\sigma)}{\alpha(a)} \frac{\pi(\sigma, c)}{r},$$

has σ -derivative $\frac{\alpha(\sigma)}{\alpha(a)} \frac{\pi(\sigma, c)}{r} [D_i \log \pi(\sigma, c) - \lambda(\sigma)]$; by single-crossing (**Assumption 1**(iv)) on $[c, \infty)$ it is single-peaked, with unique maximizer the first crossing $\inf\{\sigma \geq a : D_i \log \pi(\sigma, c) \leq \lambda(\sigma)\}$.

Proof. Differentiate: the flow term contributes $\frac{\alpha(\sigma)}{\alpha(a)} \pi(\sigma, c)$ and the terminal term $\frac{\alpha(\sigma)}{\alpha(a)} [-(r + \lambda(\sigma)) \frac{\pi(\sigma, c)}{r} + \frac{D_i \pi(\sigma, c)}{r}]$; collecting gives the derivative. Its sign is that of $D_i \log \pi(\sigma, c) - \lambda(\sigma)$, strictly decreasing on $[c, \infty)$ by (iv), so the value is single-peaked. The argument lives entirely on $\{x_1 \geq x_2 = c\}$, the native domain of (iii). $\square$

Lemma 4 (Closed form past the ceiling). For $t \geq \tau^R$ the rival stops immediately on detection ($\sigma^*(a) = a$), so

$$V_S(t, \omega) = \int_0^\infty \omega e^{-\omega x} \left[\int_t^{t+x} \frac{\alpha(s)}{\alpha(t)} \pi(t, s) \, ds + \frac{\alpha(t+x)}{\alpha(t)} \frac{\pi(t, t+x)}{r} \right] dx.$$

Proof. At detection $a = t + x \geq \tau^R$, log-supermodularity ($t \leq a$) and **Assumption 2** give $D_i \log \pi(a, t) \leq D_i \log \pi(a, a) \leq \lambda(a)$, so $\sigma^*(a) = a$ by **Lemma 3**. $\square$

Lemma 5 (Monotonicity and endpoints). For every $t \geq 0$:

- (a) $\omega \mapsto V_S(t, \omega)$ is nondecreasing, and strictly increasing when $\tau^R < \infty$;
- (b) $V_S(t, 0) := \lim_{\omega \downarrow 0} V_S(t, \omega) = \int_t^\infty \frac{\alpha(s)}{\alpha(t)} \pi(t, s) \, ds < V_R(t)$;
- (c) for $t \geq \tau^R$, $\lim_{\omega \rightarrow \infty} V_S(t, \omega) = \pi(t, t)/r \geq V_R(t)$ with strict inequality for $t > \tau^R$.

Cooperation threshold. $\Delta(t, \omega)$ captures the value of stopping now versus racing forever, conditional on no news at t . Define

$$\bar{\omega} := \sup \left\{ \omega \geq 0 : \Delta(t, \omega) \leq 0 \text{ for all } t \geq 0 \right\}.$$

Because $\Delta(t, \omega)$ is nondecreasing in ω for every t (**Lemma 5(a)**), the set inside the supremum is downward closed. Hence $\omega < \bar{\omega}$ implies $\Delta(t, \omega) \leq 0$ for every t , while every $\omega > \bar{\omega}$ admits some date t with $\Delta(t, \omega) > 0$. If $\tau^R < \infty$, then **Lemma 5(c)**, applied at any $t > \tau^R$, implies $\bar{\omega} < \infty$.

Transparency level $\bar{\omega}$ is the boundary below which no date offers a profitable unilateral stop against a rival that races until detection. Hence, both firms prefer to race on; this is the sense in which observation lags hinder coordination. The following lemma gives an expression for the payoff from continuing until $\sigma \geq t$ then stopping.

Lemma 6 (Recursive identity). Against a rival who continues until news arrives, the payoff to “continue to $\sigma \geq t$, then stop” is

$$\begin{aligned} W_t(\sigma) &= \int_t^\sigma \frac{\alpha(s)}{\alpha(t)} \pi(s, s) \, ds + \frac{\alpha(\sigma)}{\alpha(t)} V_S(\sigma, \omega) \\ &= V_R(t) + \frac{\alpha(\sigma)}{\alpha(t)} \Delta(\sigma, \omega) \end{aligned}$$

and $W_t(\infty) = V_R(t)$.

Proof. On $[t, \sigma]$ both race (flow $\pi(s, s)$); after σ the continuation is $V_S(\sigma, \omega)$ by definition. Subtract the splitting identity $V_R(t) = \int_t^\sigma \frac{\alpha(s)}{\alpha(t)} \pi(s, s) ds + \frac{\alpha(\sigma)}{\alpha(t)} V_R(\sigma)$. As $\sigma \rightarrow \infty$, $\int_t^\sigma \frac{\alpha(s)}{\alpha(t)} \pi(s, s) ds \rightarrow V_R(t)$. It remains to show the tail term vanishes: $\frac{\alpha(\sigma)}{\alpha(t)} V_S(\sigma, \omega) \rightarrow 0$. By Lemma 5, $0 \leq V_S(\sigma, \omega) \leq \pi(\sigma, \sigma)/r$ for $\sigma \geq \tau^R$, so it suffices that $\alpha(\sigma) \pi(\sigma, \sigma) \rightarrow 0$. For $\sigma \geq \tau^R$,

$$\frac{d}{d\sigma} \log \pi(\sigma, \sigma) = D_i \log \pi(\sigma, \sigma) + D_j \log \pi(\sigma, \sigma) \leq \lambda(\sigma),$$

using $D_i \log \pi(\sigma, \sigma) \leq \lambda(\sigma)$ (Assumption 2) and $D_j \log \pi < 0$. Integrating from τ^R gives $\pi(\sigma, \sigma) \leq \pi(\tau^R, \tau^R) \exp\left(\int_{\tau^R}^\sigma \lambda(u) du\right)$; since $\alpha(\sigma) = e^{-r\sigma - \int_0^\sigma \lambda(u) du}$,

$$\alpha(\sigma) \pi(\sigma, \sigma) \leq \pi(\tau^R, \tau^R) e^{-\int_0^{\tau^R} \lambda(u) du} e^{-r\sigma} \rightarrow 0.$$

Hence $W_t(\infty) = V_R(t)$. □

Proposition 7 (Catastrophe under noisy monitoring). Suppose $\omega < \bar{\omega}$. Then there exists $\tau \in \text{PBE}(\omega)$ such that $\tau_1 = \tau_2 = +\infty$ almost surely.

Proposition 7 states that imperfect observation can generate equilibria in which both firms scale forever even when no such equilibrium exists under perfect transparency.⁷

For $\omega > \bar{\omega}$, some date offers a strictly profitable unilateral stop against a rival that races until detection. Even so, each firm would rather that their rival be the one to first stop, so the game after τ^R resembles a war of attrition. It remains to quantify how much equilibrium delay pushes the frontier far beyond τ^R . The force that limits this delay is simple. Once both firms are near τ^R , the private gain from being the frontier firm is small when detection is fast, while the joint loss from letting both firms keep racing is first order.

To measure this joint loss, fix a no-signal history at which both firms have continued until t , and ask what payoff a firm would get if the two firms could commit to a common cap $\sigma \geq t$: both race until σ and then both stop. This benchmark is

$$g_t(\sigma) := \int_t^\sigma \frac{\alpha(s)}{\alpha(t)} \pi(s, s) ds + \frac{\alpha(\sigma)}{\alpha(t)} \frac{\pi(\sigma, \sigma)}{r} \quad (\text{JOINT})$$

$g_t(\sigma)$ is the value of a coordinated stop at σ when both firms are active at time t . Differentiating gives

$$-g'_t(\sigma) = \frac{\alpha(\sigma)}{\alpha(t)} \frac{\pi(\sigma, \sigma)}{r} \left[\lambda(\sigma) - D_i \log \pi(\sigma, \sigma) - D_j \log \pi(\sigma, \sigma) \right].$$

⁷Note, however, that this is a statement about the “worst” equilibrium; others in which both firms stop in finite time still survive.

At $\sigma = \tau^R$, the private stopping condition gives $\lambda(\tau^R) = D_i \log \pi(\tau^R, \tau^R)$ whenever $0 < \tau^R < +\infty$, so the term in the bracket above reduces to $-D_j \log \pi(\tau^R, \tau^R) > 0$.⁸ The representative firm is locally indifferent at τ^R , but when both firms move together, each firm imposes the rival externality on the other. Hence coordinated delay past τ^R burns payoff at a rate bounded away from zero.

Fast detection makes the prize from holding out small. If one firm concedes, the other learns this after an exponential delay with mean $1/\omega$, so the winner's extra frontier advantage is only $O(1/\omega)$; equivalently, a firm can secure $\pi(t, t)/r - O(1/\omega)$ by stopping at once and it cannot be optimal to hold out for an $O(1/\omega)$ prize while the common frontier loses value at a first-order rate. We will show that this implies that the frontier can overshoot τ^R by at most order $1/\omega$. Before we show this, ruling out an infinite frontier requires the profitable stop to survive into the tail. This is guaranteed by the following condition on primitives.

Definition 2 (Tail condition). Suppose $\tau^R < \infty$, and define

$$\kappa_0 := \inf_{t \geq \tau^R} [-D_j \log \pi(t, t)], \quad \Lambda_\infty := \sup_{t \geq \tau^R} \sup_{u \geq t} \{\lambda(u) + |D_j \log \pi(t, u)|\}.$$

If $\kappa_0 > 0$ and $\Lambda_\infty < \infty$ we say that the *tail condition* is fulfilled.

The lower bound $\kappa_0 > 0$ says that the competitive loss from a racing rival does not disappear in the tail. The upper bound $\Lambda_\infty < \infty$ limits both the hazard and the erosion suffered while a stopper waits to be detected. Both conditions are satisfied by the exponential example below. Under the tail condition, define

$$\omega^{\text{cap}} := \frac{e\Lambda_\infty(r + \Lambda_\infty)}{\kappa_0} \quad \text{noting that} \quad \omega^{\text{cap}} \geq \bar{\omega}.$$

Proposition 8 (Fast monitoring uniformly caps the frontier). Suppose the tail condition holds. If $\omega > \omega^{\text{cap}}$, then

$$\sup_{\tau \in \text{PBE}(\omega)} \mathbb{E} \left[(\max_i \tau_i - \tau^R)^+ \right] \leq \frac{1}{\omega} + \frac{e\Lambda_\infty}{\kappa_0(\omega - \omega^{\text{cap}})}.$$

In particular, every equilibrium stops in finite time almost surely, and $(\max_i \tau_i - \tau^R)^+ \rightarrow 0$ in probability as $\omega \rightarrow \infty$.

⁸When $\tau^R = 0$, the condition gives $\lambda(0) \geq D_i \log \pi(0, 0)$ so the term in the bracket is at least $-D_j \log \pi(\tau^R, \tau^R) > 0$ may strictly prefer to stop immediately and coordinated delay once again burns payoff.

Proof of Proposition 8. Fix an equilibrium and define

$$P(t) := \frac{\pi(t, t)}{r} \quad \text{and} \quad \phi(t) := \lambda(t) - D_i \log \pi(t, t) - D_j \log \pi(t, t).$$

Here $P(t)$ is each firm's payoff if both firms stop at t , while $\phi(t)$ is the net instantaneous cost of delaying both firms: the hazard rate less growth rate of flow profits evaluated when firms are neck-and-neck.

We first obtain a uniform bound on delay over a short interval and then iterate it. By **Assumption 2** and the tail condition, $\kappa_0 \leq \phi(t) \leq \Lambda_\infty$ ($t \geq \tau^R$), where the lower bound uses $\lambda - D_i \log \pi \geq 0$ and $-D_j \log \pi \geq \kappa_0$ and the upper bound uses $D_i \log \pi > 0$ and $\lambda + |D_j \log \pi| \leq \Lambda_\infty$. For $t \geq \tau^R$, define the loss from a detection lag of duration x by

$$\psi_t(x) := P(t) - \left[\int_t^{t+x} \frac{\alpha(s)}{\alpha(t)} \pi(t, s) \, ds + \frac{\alpha(t+x)}{\alpha(t)} \frac{\pi(t, t+x)}{r} \right].$$

By **Lemma 4**, $P(t) - V_S(t, \omega) = \int_0^\infty \omega e^{-\omega x} \psi_t(x) \, dx$. Moreover, $\psi_t(0) = 0$ and

$$0 \leq \psi'_t(x) = \frac{\alpha(t+x)}{\alpha(t)} \frac{\pi(t, t+x)}{r} [\lambda(t+x) - D_j \log \pi(t, t+x)] \leq P(t) \Lambda_\infty.$$

Hence $\psi_t(x) \leq P(t) \Lambda_\infty x$, and therefore

$$V_S(t, \omega) \geq P(t) \left(1 - \frac{\Lambda_\infty}{\omega} \right). \quad (\text{STOP})$$

Using the coordinated-stop value g_t from **(JOINT)**, for $0 \leq x \leq 1/(r + \Lambda_\infty)$ we have

$$-\frac{g'_t(t+x)}{P(t)} = e^{-rx - \int_t^{t+x} \phi(v) \, dv} \phi(t+x) \geq \kappa_0 e^{-(r+\Lambda_\infty)x} \geq \frac{\kappa_0}{e}.$$

It follows that, for every $s \geq t \geq \tau^R$,

$$g_t(s) \leq P(t) \left[1 - \frac{\kappa_0}{e} ((s-t) \wedge h) \right]. \quad (\text{DELAY})$$

Taking $s \rightarrow \infty$ in (DELAY) and using $g_t(\infty) = V_R(t)$ gives

$$V_R(t) \leq P(t) \left(1 - \frac{\kappa_0}{e(r + \Lambda_\infty)} \right).$$

Together with (STOP), this implies $V_S(t, \omega) > V_R(t)$ for every $t \geq \tau^R$ whenever $\omega > \omega^{\text{cap}}$. Hence $\bar{\omega} \leq \omega^{\text{cap}}$.

We next use equilibrium optimality. Fix $t \geq \tau^R$ with $\mathbb{P}(\min_i \tau_i \geq t) > 0$. At an active no-news history, a firm may be uncertain whether its rival is still active or has stopped without being detected. Under the latter event, stopping immediately is weakly optimal by [Assumption 2](#), log-supermodularity, and [Lemma 3](#). This thus cancels in the comparison between equilibrium continuation and the deviation to stop immediately. On the former event, stopping immediately yields at least $V_S(t, \omega)$. This is because the definition of V_S assumes that the rival continues until it detects the stop which is the worst case scenario—the rival might stopper sooner.

We next bound equilibrium continuation conditional on the rival still being active. Both firms earn the diagonal flow until the first stop. At that date, the continuation payoff of either firm is at most $P(\min_i \tau_i)$.⁹ where we recall that $g_t(X)$ defined above is the ‘coordinated stop value’ from coordinating a stop at x . Thus equilibrium continuation is at most

$$\mathbb{E} \left[g_t(\min_i \tau_i) \mid \min_i \tau_i \geq t \right].$$

Combining this bound with (STOP) and (DELAY) we have:

$$\begin{aligned} P(t) \left(1 - \frac{\Lambda_\infty}{\omega} \right) &\leq V_S(t, \omega) \\ &\leq \mathbb{E} \left[g_t(\min_i \tau_i) \mid \min_i \tau_i \geq t \right] \\ &\leq P(t) \left\{ 1 - \frac{\kappa_0}{e} \mathbb{E} \left[((\min_i \tau_i - t)^+) \wedge \frac{1}{r + \Lambda_\infty} \mid \min_i \tau_i \geq t \right] \right\}. \end{aligned}$$

Canceling $P(t)$ and rearranging yields

$$\mathbb{E} \left[((\min_i \tau_i - t)^+) \wedge \frac{1}{r + \Lambda_\infty} \mid \min_i \tau_i \geq t \right] \leq \frac{e\Lambda_\infty}{\kappa_0\omega} = \frac{\omega^{\text{cap}}}{\omega(r + \Lambda_\infty)}. \quad (\text{BLOCK})$$

where the equality is just from the definition of ω^{cap} . Because the random variable in (BLOCK) equals $1/(r + \Lambda_\infty)$ on $\{\min_i \tau_i \geq t + 1/(r + \Lambda_\infty)\}$,

$$\mathbb{P} \left(\min_i \tau_i \geq t + \frac{1}{r + \Lambda_\infty} \mid \min_i \tau_i \geq t \right) \leq \frac{\omega^{\text{cap}}}{\omega} < 1.$$

Applying this inequality successively at $t = \tau^R, \tau^R + 1/(r + \Lambda_\infty), \dots, \tau^R + (k - 1)/(r +$

⁹For the leader that stopped, her best case is for the follower to also stop which delivers $P(\min_i \tau_i)$; for the follower who knew that the leader has stopped at $\min_i \tau_i$, her best response is to stop immediately which also yields $P(\min_i \tau_i)$.

Λ_∞) yields

$$\mathbb{P} \left(\min_i \tau_i \geq \tau^R + \frac{k}{r + \Lambda_\infty} \mid \min_i \tau_i \geq \tau^R \right) \leq \left(\frac{\omega^{\text{cap}}}{\omega} \right)^k.$$

Applying (BLOCK) at $t = \tau^R + k/(r + \Lambda_\infty)$ and using the preceding bound gives

$$\mathbb{E} \left[\left(\min_i \tau_i - \tau^R - \frac{k}{r + \Lambda_\infty} \right)^+ \wedge \frac{1}{r + \Lambda_\infty} \right] \leq \frac{1}{r + \Lambda_\infty} \left(\frac{\omega^{\text{cap}}}{\omega} \right)^{k+1}.$$

Observe that we have

$$\left(\min_i \tau_i - \tau^R \right)^+ = \sum_{k=0}^{\infty} \left[\left(\min_i \tau_i - \tau^R - \frac{k}{r + \Lambda_\infty} \right)^+ \wedge \frac{1}{r + \Lambda_\infty} \right].$$

This partitions the positive overshoot into consecutive intervals of length $1/(r + \Lambda_\infty)$. If the overshoot lies between $k/(r + \Lambda_\infty)$ and $(k + 1)/(r + \Lambda_\infty)$, the first k terms each equal $1/(r + \Lambda_\infty)$, the next term equals the remaining overshoot, and every later term is zero. Monotone convergence then gives

$$\mathbb{E} \left[\left(\min_i \tau_i - \tau^R \right)^+ \right] \leq \frac{\omega^{\text{cap}}/\omega}{(r + \Lambda_\infty)(1 - \omega^{\text{cap}}/\omega)} = \frac{e\Lambda_\infty}{\kappa_0(\omega - \omega^{\text{cap}})}.$$

Finally, after the first stop at $\min_i \tau_i$, the other firm detects it after $\tilde{D} \sim \text{Exp}(\omega)$. By **Assumption 2**, log-supermodularity, and **Lemma 3**, it stops no later than $\max\{\tau^R, \min_i \tau_i + \tilde{D}\}$. Thus, pathwise,

$$\left(\max_i \tau_i - \tau^R \right)^+ \leq \left(\min_i \tau_i - \tau^R \right)^+ + \tilde{D}.$$

Taking expectations and using the preceding bound proves the result. The bound is finite and uniform across equilibria. Its right-hand side converges to zero as $\omega \rightarrow \infty$, which gives the uniform L^1 convergence and convergence in probability follows from Markov's inequality. $\square$

Proposition 8 shows how cooperation partially succeeds when news arrives sufficiently quickly ($\omega > \omega^{\text{cap}}$). If so, then for the equilibrium in which the technology is pushed past τ^R , outlasting one's rival is worth only the extra flow earned during the rival's detection lag (of order $1/\omega$) while joint delay burns value at a rate bounded away from zero. Since it cannot be rational to hold out for a long time to gain a vanishing prize, the expected delay until concession is order $1/\omega$, so the frontier overshoots τ^R by

$O(1/\omega)$ in expectation where this bound is uniform across all equilibria.¹⁰ As $\omega \rightarrow \infty$, the overshoot vanishes in probability so the perfect-observability ceiling is recovered in the limit.

Proposition 8 is stated for the monitoring regime $\omega > \omega^{\text{cap}} \geq \bar{\omega}$ in order to discipline stopping incentives in the tail. However, for stationary environments such as those analyzed in the exponential below as well as in [Section 6](#), the threshold is tight: when $\omega < \bar{\omega}$ there exists an equilibrium in which firms race forever; when $\omega > \bar{\omega}$ in every equilibria $\max_i \tau < +\infty$ almost surely; we show this in [Online Appendix III](#).

Example 4 (Exponential race with imperfect transparency). Let flow payoffs be exponential, $\pi(t_i, t_j) = e^{\alpha t_i - \beta t_j}$, $\alpha, \beta > 0$. Both indices of the previous section are then constant and equal: $D_i \log \pi(t, t) = D_i \log \pi(t, 0) = \alpha$ and $D_{ij} \log \pi = 0$, so with an increasing hazard $\lambda(t) = \lambda_0 e^{\kappa t}$ ($\kappa > 0$, $\lambda_0 < \alpha$) the bracket is degenerate: $\underline{\tau} = \tau^R = \frac{1}{\kappa} \log \frac{\alpha}{\lambda_0}$, so under perfect transparency, firms stop at τ^R in every equilibrium. This makes the exponential race a natural setting to isolate the effect of imperfect transparency: any departure from τ^R is attributable to monitoring frictions alone. We take the hazard constant at a level exceeding the index: $\lambda(t) \equiv \lambda > \alpha$, with $r > \alpha - \beta$, so that $\tau^R = 0$.¹¹ The environment is then *stationary*: at a symmetric no-signal history at technology t , every continuation payoff scales with the current flow $\pi(t, t) = e^{(\alpha - \beta)t}$, so values per unit of current flow are date-independent, and $\Delta(t, \omega)$ has the same sign at every date. Direct computation gives the normalized values

$$V_R = \frac{1}{r + \lambda - \alpha + \beta}, \quad V_S(0) = \frac{1}{r + \lambda + \beta}, \quad V_S(\omega) = V_S(0) + \frac{\omega[1/r - V_S(0)]}{\omega + r + \lambda + \beta} :$$

racing forever discounts at $r + \lambda$ while the diagonal flow grows at $\alpha - \beta$; stopping first freezes one's own flow, which then decays at the erosion rate β under a still-live hazard until the detection date, whose expected discount factor $\omega/(\omega + r + \lambda + \beta)$ weights the upgrade to the joint-stop perpetuity $1/r$. Setting $V_S(\omega) = V_R$ and solving, $\bar{\omega} = \frac{r\alpha}{\lambda - \alpha + \beta}$. Every parameter moves $\bar{\omega}$ in the intuitive direction. The transparency needed for cooperation rises with the scaling gain α , which pays only while racing, and with impatience r , which discounts the reciprocation upgrade. It falls with the hazard λ and the erosion β : both depress the race value faster than the stop value (the race lasts forever; the exposure

¹⁰The proof gives a geometric block-tail bound and an $O(1/\omega)$ bound on the expected frontier overshoot which does not require equilibrium stopping distributions to have densities. In the smooth symmetric equilibrium considered below, the short delay is implemented by a concession hazard proportional to ω .

¹¹This is the boundary case of [Assumption 1](#)(iv): the index α never crosses the hazard from above because it starts below it, and [Assumption 2](#) holds trivially. Interpret (π, λ) as the continuation environment once the race has reached the threshold; the next section adopts the same specification.

window ends at detection), so more dangerous technologies and stronger competitive erosion make cooperation easier to sustain.

By stationarity, $\Delta(\cdot, \omega)$ has the same sign at every date. Hence, when $\omega > \bar{\omega}$ so that $\Delta > 0$ everywhere, a unilateral stop is profitable at every date and [Proposition 12](#) in [Online Appendix III](#) shows that every equilibrium stops in finite time. [Figure 6](#) displays the threshold at which V_S crosses V_R , and the vanishing prize. Panel (a) illustrates how the stop value $V_S(\omega)$ rises from $V_S(0) = 1/(r + \lambda + \beta)$ toward $1/r$ as news speeds up, crossing the race value $V_R = 1/(r + \lambda - \alpha + \beta)$ at $\bar{\omega} = r\alpha/(\lambda - \alpha + \beta) = 1/15$: below the threshold racing forever is self-enforcing ([Proposition 7](#)), and above it all equilibria are finite, conceding quickly in expectation when transparency is high ([Proposition 8](#)).




Figure 6: Exponential under imperfect monitoring ($r = 0.1$, $\lambda = 1$, $\alpha = 0.5$, $\beta = 0.25$, values per unit of current flow)

To see why concession speeds up as monitoring becomes more precise, note that the (normalized¹²) value upon detection that one's rival has stopped instead of first is: $V_W(\omega) - V_S(\omega) \simeq \frac{\alpha + \beta}{r\omega}$, where $V_W(\omega)$ is the expected payoff from racing against an opponent who has stopped and stopping only upon detection. This is illustrated by panel (b). By contrast, joint delay burns value at the (normalized) constant rate $(\lambda - \alpha + \beta)/r$ per unit of time. Hence, equilibrium indifference equires concession at a hazard of order ω , leaving expected delay of order $1/\omega$.

6 TRUST

We are interested in firms' ability to coordinate when they fear their opponent might irrationally never stop. Continue with the stationary exponential environment of [Example 4](#),

¹²We normalize this by the flow profits $\pi(t, t)$.

imposing in addition $0 < \beta < \alpha$ and $\omega \geq \beta$. Thus $\pi(t_i, t_j) = e^{\alpha t_i - \beta t_j}$, $\lambda > \alpha$, and $r > \alpha - \beta$, and the representative threshold is $\tau^R = 0$. We now suppose that each firm is independently either a *rational* type with these payoffs or a *crazy* type that never stops. Each firm is rational with probability p , and we write $\ell := p/(1-p)$ for the prior odds of rationality.

As in Example 4, we normalize continuation payoffs by current flow and write

$$V_R = \frac{1}{r + \lambda - \alpha + \beta}, \quad V_S(0) = \frac{1}{r + \lambda + \beta} \text{ and } V_S(\omega) = V_S(0) + \frac{\omega[1/r - V_S(0)]}{\omega + r + \lambda + \beta}.$$

Recall that $V_S(\omega) > V_R$ if and only if $\omega > \bar{\omega} := \frac{r\alpha}{\lambda - \alpha + \beta}$.

We additionally define the *verification value*

$$V_W(\omega) := \frac{r + \omega}{r(\omega + r + \lambda - \alpha)}$$

as the normalized payoff from racing beside a rival that has already stopped at the same technology, and exiting upon detection; this is the value of waiting for confirmation that a firm's rival is not crazy. Note $V_R < V_W(\omega) < 1/r$: verifying beats racing forever but falls short of joining, by the *lag cost* $1/r - V_W(\omega)$.

Two kinds of coordination. We now identify two distinct coordination modes and the conditions under which each is incentive-compatible. For an agent to stop first i.e., without knowing if their rival has stopped, she gambles on both their rival's type and on news arriving quickly: if their rival is rational, it stops upon receiving the news of their stop, and never stops otherwise:

$$\frac{\ell}{1+\ell} V_S(\omega) + \frac{1}{1+\ell} V_S(0) \geq V_R \iff \ell \geq \bar{\ell} := \begin{cases} \frac{V_R - V_S(0)}{V_S(\omega) - V_R} & \text{if } \omega > \bar{\omega} \\ +\infty & \text{otherwise} \end{cases}$$

Alternatively, rational firms might stop at a common date. In this case, we have a new deviation that was not present when firms were commonly known to be rational—firms are tempted to *wait and verify*: hold out for news that their rival has, in fact, stopped. If the rival is rational, this delivers $V_W(\omega)$ instead of $1/r$. If the rival is crazy, then the deviator obtains V_R (rival never stops) instead of $V_S(0)$. Hence, stopping at a known date is optimal if and only if

$$\ell [1/r - V_W(\omega)] \geq V_R - V_S(0) \iff \ell \geq \underline{\ell} := \frac{V_R - V_S(0)}{1/r - V_W(\omega)}.$$

We assume that $\underline{\ell} < \bar{\ell}$ which implies that simultaneous coordination is less demanding than sequential coordination.¹³ Because $\omega \geq \beta$, waiting indefinitely for verification that their rival has stopped is the binding deviation. The two thresholds have the same numerator $V_R - V_S(0)$, the cost of stopping unreciprocated, and differ in the denominator: the *prize of reciprocation* $V_S(\omega) - V_R$ for sequential stopping and the *lag cost* $1/r - V_W(\omega)$ for simultaneous stopping. Figure 7(a) shows $V_S(\omega)$ and $V_W(\omega)$ rising toward $1/r$: the prize of reciprocation grows while the lag cost shrinks, so faster news makes both stopping first and waiting to verify more attractive. Panel (b) plots the trust-weighted values of the four candidate strategies at $\omega = 0.5$. The lighter curves repeat the exercise at $\omega' = 0.8$. Racing and stopping together are unaffected by ω ; only stopping first and waiting to verify shift. Faster news pulls $\bar{\ell}$ down and pushes $\underline{\ell}$ up, squeezing the medium zone until it closes at $\omega^\dagger$.




Figure 7: Transparency and trust; stationary environment: $(r, \lambda, \alpha, \beta) = (0.1, 1, 0.5, 0.25)$.

Transparency thus moves the two thresholds in opposite directions. Faster monitoring speeds reciprocation, so $\bar{\ell}$ falls—from $+\infty$ at $\bar{\omega}$ (no trust substitutes for a minimum of monitoring) toward the strictly positive limit $\frac{V_R - V_S(0)}{1/r - V_R}$ —while it also cheapens verification, inviting free-riding on the signal, so $\underline{\ell}$ rises without bound. This partitions the trust line into three zones (Figure 8): *low* ($\ell < \underline{\ell}$), *medium* ($\underline{\ell} \leq \ell < \bar{\ell}$), and *high* ($\ell \geq \bar{\ell}$).

¹³Each of V_R , $V_S(0)$, $V_S(\omega)$, $V_W(\omega)$, and $1/r$ is the payoff of a fixed course of play and does not depend on equilibrium except for the fact that a sole survivor who learns their rival has stopped also stops immediately, which is implied by dominance ($\alpha < \lambda$). If $\underline{\ell} > \bar{\ell}$, the middle zone disappears. Never stopping remains equilibrium for $\ell \leq \bar{\ell}$ and is ruled out for $\ell > \bar{\ell}$ when stopping first is strictly profitable.




Figure 8: Trust and transparency. Parameters: $(r, \lambda, \alpha, \beta) = (0.1, 1, 0.5, 0.25)$.

Figure 8 illustrates the thresholds $\underline{\ell}$ and $\bar{\ell}$. The threshold $\underline{\ell}$ rises with ω (cheaper verification invites free-riding on the signal; below $\omega = \beta$ it is constant at its $\omega = \beta$ value—the vertical segment), while the threshold $\bar{\ell}$ falls (faster reciprocation), with horizontal asymptote $\bar{\omega} = 1/15$ and vertical asymptote $\bar{\ell}(\infty) \approx 0.049$: the high-trust zone lies entirely above the main text’s cooperation threshold, and no transparency pushes its boundary below $\bar{\ell}(\infty)$. The curves cross once, at $(\ell^\dagger, \omega^\dagger) \approx (0.132, 0.914)$. The dotted vertical, at trust $\ell = 0.09$, shows the double-edged effect of transparency: raising ω first destroys simultaneous coordination ($M \rightarrow L$ at $\omega \approx 0.43$) and only later makes sequential coordination viable ($L \rightarrow H$ at $\omega \approx 1.79$).

Why do trust and transparency interact so differently across the two modes of coordination? Sequential coordination asks a firm to stop first, gambling that a rational rival will reciprocate once the news lands. Hence, faster news raises the prize of reciprocation $V_S(\omega) - V_R$, so transparency helps. Conversely, simultaneous coordination requires that a firm not be tempted to keep racing, and stop only after seeing that the rival really did stop. But faster news shrinks the lag cost $1/r - V_W(\omega)$ of that deviation, so transparency hurts. This double-edged effect of transparency is visible in Figure 8: at intermediate levels of trust (e.g., $\ell = 0.10$), raising ω can destroy the less demanding simultaneous mode ($M \rightarrow L$) long before it makes the sequential mode viable ($L \rightarrow H$).

Extending the notation of Section 5, write $\text{PBE}(\omega, \ell)$ for the set of perfect Bayesian equilibrium *outcomes*—profiles of stopping times $\tau = (\tau_1, \tau_2)$ —when the news rate is ω and the prior odds of rationality are ℓ .

Proposition 9. The following characterizes equilibrium outcomes across trust zones:

- (i) *Low trust.* If $\ell < \underline{\ell}$, then every $\tau \in \text{PBE}(\omega, \ell)$ has $\tau_1 = \tau_2 = +\infty$ almost surely: no firm ever stops, and the disaster arrives with probability one.
- (ii) *Medium trust.* If $\underline{\ell} \leq \ell \leq \bar{\ell}$, then $\text{PBE}(\omega, \ell)$ contains both (a) the outcome in which rational firms stop immediately, so that the risk survives only through crazy types, and (b) the never-stopping outcome.
- (iii) *High trust.* If $\ell > \bar{\ell}$, then every $\tau \in \text{PBE}(\omega, \ell)$ satisfies

$$\mathbb{P}(\tau_1 = \tau_2 = +\infty \mid \text{both firms rational}) \leq (\bar{\ell}/\ell)^2,$$

and stopping at zero remains $\text{PBE}(\omega, \ell)$.

Low trust makes any unilateral stop too likely to go unreciprocated, so every equilibrium races forever. At intermediate trust, simultaneous stopping becomes sustainable, but neither firm is willing to stop first, leaving both stopping and racing equilibria. At high trust, stopping first becomes profitable whenever the rival's holdout probability is too large, yielding the quadratic bound.

Although [Proposition 9](#) was stated for the case in which $\underline{\ell} < \bar{\ell}$, similar results hold in the opposite case. There, Part (i) survives verbatim with $\bar{\ell}$ (now the smaller threshold) as the boundary of the low zone. There is no middle region. Part (iii) survives for $\ell > \bar{\ell}$ and the $(\bar{\ell}/\ell)^2$ bound holds unchanged. [Table 1](#) summarizes equilibrium behavior across various combinations of trust and transparency.

		perfect observability ($\omega = \infty$)	imperfect observability ($\omega < \infty$)
no crazy types ($p = 1$)		Coordination is easy: in every equilibrium, both firms stop at once.	Single threshold $\bar{\omega}$: below it, never-stopping is an equilibrium alongside the joint stop; above it, every equilibrium coordinates.
crazy types ($p < 1$)		Single threshold $\bar{\ell}(\infty)$: below it, every equilibrium races to ruin; above it, ruin $\leq (\bar{\ell}(\infty)/\ell)^2$, with zero attained only by leader-follower outcomes	Two thresholds $\underline{\ell}$ rising and $\bar{\ell}$ falling in ω , and three zones (Proposition 9).

Table 1: Equilibria across transparency and trust regimes

7 CONCLUSION

We have shown that competition and coordination play distinct roles in determining how far a dangerous technology advances. Under perfect transparency, competition places the frontier between monopoly and representative-firm thresholds. Delayed observation can sustain racing forever or generate a short war of attrition beyond the perfect-transparency ceiling. Imperfect trust is more damaging: when firms doubt their rivals' rationality, racing to ruin may occur in every equilibrium. Transparency and trust are therefore complements in some regions but substitutes in others.

Profile effects We might alternatively suppose that risk also depends on the entire profile of technology states, and that when the second firm's technology is near the frontier it matters almost as much as the leader.¹⁴ We conjecture that these profile effects can reverse a force that facilitated coordination in the baseline perfect-transparency model: A firm that stops now exerts a positive externality on the racers it leaves behind: its exit thins the frontier and lowers the hazard they face, so stopping first might embolden one's competitors to race for longer.

Accumulative risks Suppose that once the technology crosses a threshold, the industry lives with a permanent background chance of collapse: the hazard depends on the peak scale reached so far and persists after every firm has stopped.¹⁵ Stopping then ends the gains from scaling but not the exposure to risk, so the value of settling at technology T is the flow payoff discounted at $r + \lambda(T)$ rather than at r . This changes the representative firm's trade-off: the hazard level is borne whether or not the firm stops, so the cost of continuing is only the permanent increment to background risk and the stopping rule becomes¹⁶

$$D_i \log \pi(t, t) \leq \frac{\lambda'(t)}{r + \lambda(t)}.$$

Hence, what disciplines the race is the growth of risk, not its level: a hazard that is severe but flat cannot stop the race. Moreover, comparative statics in the level of risk invert: a uniform upward shift of the hazard lowers the right-hand side and so prolongs the race. Firms that already live with severe background risk have little continuation value left to

¹⁴We might further suppose that risks are persistent, so that firms at frontier T generate risk even after they stop scaling; the next discussion takes up persistence.

¹⁵In the language of Trammell and Aschenbrenner (2024) this is *state* risk rather than *transition* risk.

¹⁶Differentiating the analog of $R_i(\tau)$ with terminal value $\pi(\tau, t)/(r + \lambda(\tau))$ gives a marginal value of delay proportional to $D_i \log \pi(\tau, t, \dots, t) - \lambda'(\tau)/(r + \lambda(\tau))$. A formal statement and the requisite assumptions are collected in Online Appendix II.

protect, so marginal risk is cheap to them—a “nothing left to lose” effect that is absent from the baseline model, where stopping extinguishes risk and uniformly higher hazards always shorten the race (**Proposition 5(ii)**).

Reversibility Finally, if stopping is reversible, so that a firm that halts scaling may later resume, the ceiling τ^R no longer applies since the backward-induction argument leans on the permanence of stopping: Now races beyond τ^R can be sustained by rivals’ threats to resume scaling. Moreover, those threats can support efficient early stopping strictly below τ^R , sustained by the threat of racing. We are exploring these issues in ongoing work.

REFERENCES

- AGHION, P., C. HARRIS, P. HOWITT, AND J. VICKERS (2001): “Competition, Imitation and Growth with Step-by-Step Innovation,” *The Review of Economic Studies*, 68, 467–492.
- ALIPRANTIS, C. D. AND K. C. BORDER (2006): *Infinite Dimensional Analysis: A Hitchhiker’s Guide*, Berlin: Springer, 3rd ed.
- ARMSTRONG, S., N. BOSTROM, AND C. SHULMAN (2016): “Racing to the Precipice: A Model of Artificial Intelligence Development,” *AI & Society*, 31, 201–206.
- BAYE, M. R., D. KOVENOCK, AND C. G. DE VRIES (1996): “The All-Pay Auction with Complete Information,” *Economic Theory*, 8, 291–305.
- BENGIO, Y., G. HINTON, A. YAO, D. SONG, P. ABBEEL, T. DARRELL, Y. N. HARARI, Y.-Q. ZHANG, L. XUE, S. SHALEV-SHWARTZ, ET AL. (2024): “Managing extreme AI risks amid rapid progress,” *Science*, 384, 842–845.
- BOSTROM, N. (2014): “Superintelligence: Paths, dangers, strategies.”
- BUENO DE MESQUITA, E., W. DZIUDA, AND M. POLBORN (2026): “The AGI Race and Existential Risk,” Working paper, University of Chicago.
- BULOW, J. AND P. KLEMPERER (1999): “The Generalized War of Attrition,” *The American Economic Review*, 89, 175–189.
- DALEY, B. AND B. GREEN (2012): “Waiting for News in the Market for Lemons,” *Econometrica*, 80, 1433–1504.
- DASGUPTA, P. AND J. E. STIGLITZ (1980): “Uncertainty, Industrial Structure, and the Speed of R&D,” *The Bell Journal of Economics*, 11, 1–28.
- FUDENBERG, D., R. GILBERT, J. STIGLITZ, AND J. TIROLE (1983): “Preemption, Leapfrogging and Competition in Patent Races,” *European Economic Review*, 22, 3–31.

- FUDENBERG, D., Y. ISHII, AND S. D. KOMINERS (2014): “Delayed-response strategies in repeated games with observation lags,” *Journal of Economic Theory*, 150, 487–514.
- FUDENBERG, D. AND J. TIROLE (1985): “Preemption and Rent Equalization in the Adoption of New Technology,” *The Review of Economic Studies*, 52, 383–401.
- (1986): “A theory of exit in duopoly,” *Econometrica: Journal of the Econometric Society*, 943–960.
- GALE, D. (1995): “Dynamic coordination games,” *Economic theory*, 5, 1–18.
- HALL, R. E. AND C. I. JONES (2007): “The value of life and the rise in health spending,” *The Quarterly Journal of Economics*, 122, 39–72.
- HARRIS, C. AND J. VICKERS (1985): “Perfect Equilibrium in a Model of a Race,” *The Review of Economic Studies*, 52, 193–209.
- (1987): “Racing with Uncertainty,” *The Review of Economic Studies*, 54, 1–21.
- HASHIM, S. AND C. FORD (2026): “AI CEOs Want to Slow Down. The World’s Too Busy to Help,” *Transformer*, <https://www.transformernews.ai/p/ai-ceos-want-to-slow-down-the-worlds-davos-demis-hassabis-dario-amodei>, january 23, 2026. Remarks by Demis Hassabis and Dario Amodei at the World Economic Forum Annual Meeting, Davos.
- HENDRICKS, K., A. WEISS, AND C. WILSON (1988): “The War of Attrition in Continuous Time with Complete Information,” *International Economic Review*, 29, 663–680.
- HOPENHAYN, H. A. AND F. SQUINTANI (2011): “Preemption Games with Private Information,” *The Review of Economic Studies*, 78, 667–692.
- JONES, C. I. (2024): “The AI dilemma: Growth versus existential risk,” *American Economic Review: Insights*, 6, 575–590.
- KOH, A. AND S. SANGUANMOO (2024): “Robust Technology Regulation,” *arXiv preprint arXiv:2408.17398*.
- KONRAD, K. A. (2009): *Strategy and Dynamics in Contests*, Oxford: Oxford University Press.
- KREPS, D. M. AND R. WILSON (1982): “Reputation and imperfect information,” *Journal of economic theory*, 27, 253–279.
- KRISHNA, V. AND J. MORGAN (1997): “An Analysis of the War of Attrition and the All-Pay Auction,” *Journal of Economic Theory*, 72, 343–362.
- LARAKI, R., E. SOLAN, AND N. VIEILLE (2005): “Continuous-Time Games of Timing,” *Journal of Economic Theory*, 120, 206–238.
- LEE, T. AND L. L. WILDE (1980): “Market Structure and Innovation: A Reformulation,” *The Quarterly Journal of Economics*, 94, 429–436.
- LOURY, G. C. (1979): “Market Structure and Innovation,” *The Quarterly Journal of Eco-*

- nomics*, 93, 395–410.
- MAYNARD SMITH, J. (1974): “The Theory of Games and the Evolution of Animal Conflicts,” *Journal of Theoretical Biology*, 47, 209–221.
- MOLDOVANU, B. AND A. SELA (2001): “The Optimal Allocation of Prizes in Contests,” *The American Economic Review*, 91, 542–558.
- MURPHY, K. M. AND R. H. TOPEL (2006): “The value of health and longevity,” *Journal of Political Economy*, 114, 871–904.
- ORD, T. (2020): *The precipice: Existential risk and the future of humanity*, Hachette UK.
- REINGANUM, J. F. (1982): “A Dynamic Game of R and D: Patent Protection and Competitive Behavior,” *Econometrica*, 50, 671–688.
- RIEDEL, F. AND J.-H. STEG (2017): “Subgame-Perfect Equilibria in Stochastic Timing Games,” *Journal of Mathematical Economics*, 72, 36–50.
- ROSEN, S. (1988): “The value of changes in life expectancy,” *Journal of Risk and Uncertainty*, 1, 285–304.
- SIEGEL, R. (2009): “All-Pay Contests,” *Econometrica*, 77, 71–92.
- SKAPERDAS, S. (1996): “Contest Success Functions,” *Economic Theory*, 7, 283–290.
- TOUZI, N. AND N. VIEILLE (2002): “Continuous-Time Dynkin Games with Mixed Strategies,” *SIAM Journal on Control and Optimization*, 41, 1073–1088.
- TRAMMELL, P. AND L. ASCHENBRENNER (2024): “Existential risk and growth,” Tech. rep., GPI Working Paper 13-2024, Global Priorities Institute, University of Oxford.
- TULLOCK, G. (1980): “Efficient Rent Seeking,” in *Toward a Theory of the Rent-Seeking Society*, ed. by J. M. Buchanan, R. D. Tollison, and G. Tullock, College Station, TX: Texas A&M University Press, 97–112.
- WEEDS, H. (2002): “Strategic Delay in a Real Options Model of R&D Competition,” *The Review of Economic Studies*, 69, 729–747.

APPENDIX

A OMITTED PROOFS

Technical lemmas We first record a few useful lemmas we will reference throughout the proofs. For stopping dates $(\tau_i, \tau_j) \in \overline{\mathbb{R}}_+^2$, define

$$U_i(\tau_i, \tau_j) := \int_0^\infty e^{-rs - \int_0^s \lambda(u) \mathbf{1}_{\{u \leq \tau_i \vee \tau_j\}} du} \pi(s \wedge \tau_i, s \wedge \tau_j) ds.$$

For a date $s \geq 0$, interpret a stopping date below s as a technology frozen in the past and a stopping date in $[s, +\infty]$ as the plan of a firm that remains active. The *date- s continuation*

payoff is

$$U_i^s(\tau_i, \tau_j) := \int_s^\infty \exp \left\{ -r(u-s) - \int_s^u \lambda(v) \mathbf{1}\{v \leq \tau_i \vee \tau_j\} dv \right\} \pi(u \wedge \tau_i, u \wedge \tau_j) du,$$

so that $U_i^0 = U_i$.

Lemma 7 (Payoff continuity). Write $\overline{\mathbb{R}}_+ = [0, \infty]$ and measure the distance between two dates by $d(x, y) = |e^{-x} - e^{-y}|$. Under **Assumption 1**, for every firm i and date $s \geq 0$, the payoff $U_i^s(\tau_i, \tau_j)$ is finite, continuous in the two stopping dates, and bounded uniformly over those dates. In particular, U_i is finite and continuous, and for every rival stopping date τ_j and every finite $t \geq 0$,

$$\lim_{\varepsilon \downarrow 0} U_i(t + \varepsilon, \tau_j) = U_i(t, \tau_j).$$

Thus the payoff from stopping immediately after t converges to the payoff from stopping at t , whether or not the rival stops at t .

Proof. Let $(\tau_i^k, \tau_j^k) \rightarrow (\tau_i, \tau_j)$ in $\overline{\mathbb{R}}_+^2$ and write $T_k := \tau_i^k \vee \tau_j^k$ and $T := \tau_i \vee \tau_j$. Fix $u \geq s$. Both capped stopping dates converge, and

$$\int_s^u \lambda(v) \mathbf{1}\{v \leq T_k\} dv \longrightarrow \int_s^u \lambda(v) \mathbf{1}\{v \leq T\} dv,$$

because the integrands converge at every $v \neq T$ and λ is integrable on $[s, u]$. Since π is continuous, the integrand of U_i^s converges pointwise in u .

Monotonicity of π gives $\pi(u \wedge \tau_i^k, u \wedge \tau_j^k) \leq \pi(u \wedge T_k, 0)$, while, since $\lambda \geq 0$ and $u \wedge T_k \leq u$,

$$\exp \left\{ - \int_s^u \lambda(v) \mathbf{1}\{v \leq T_k\} dv \right\} \leq \exp \left\{ \int_0^s \lambda(v) dv - \int_0^{u \wedge T_k} \lambda(v) dv \right\}.$$

Hence the integrand is bounded by

$$e^{rs + \int_0^s \lambda(z) dz} e^{-ru} \sup_{v \leq u} \left\{ e^{-\int_0^v \lambda(z) dz} \pi(v, 0) \right\},$$

which is integrable on $[s, \infty)$ by **Assumption 1(v)**, uniformly over stopping pairs. Dominated convergence delivers finiteness, the uniform bound, and continuity. $\square$

Lemma 8 (Earliest optimal plans). Let $f(t, s)$ be continuous on

$$\{(t, s) \in \mathbb{R}_+ \times \overline{\mathbb{R}}_+ : t \leq s\}.$$

Then $a(t) := \min \operatorname{argmax}_{s \in [t, +\infty]} f(t, s)$ is well defined, lower semicontinuous, and hence Borel measurable. If the maximizer is unique for every t , then a is continuous.

Proof. The feasible tails $[t, +\infty]$ are compact and vary continuously with t . Berge's maximum theorem therefore gives a nonempty, compact-valued, upper-hemicontinuous maximizer correspondence. Given $t_m \rightarrow t$, choose a subsequence along which $a(t_m)$ converges to its limit inferior, say q . Upper hemicontinuity implies that q is a maximizer at t , so $q \geq a(t)$. Thus a is lower semicontinuous and therefore Borel measurable. If every maximizer is unique, every convergent subsequence of $a(t_m)$ has limit $a(t)$, which gives continuity. $\square$

Proof of Lemma 1. For $\tau \geq t$, split the date-zero payoff when the rival stops at t into the flow earned before t and the continuation payoff from t onward:

$$U_i(\tau, t) = \int_0^t e^{-ru - \int_0^u \lambda(z) dz} \pi(u, u) du + e^{-rt - \int_0^t \lambda(z) dz} R_t(\tau).$$

Equivalently,

$$R_t(\tau) = e^{rt + \int_0^t \lambda(z) dz} \left[U_i(\tau, t) - \int_0^t e^{-ru - \int_0^u \lambda(z) dz} \pi(u, u) du \right].$$

The same identity holds at $\tau = +\infty$ by taking limits. Hence Lemma 7 makes $(t, \tau) \mapsto R_t(\tau)$ continuous on the feasible set with compactified stopping-date coordinate. Applying Lemma 8 to $f(t, \tau) = R_t(\tau)$ shows that a maximizer exists for every t .

For $\tau \geq t$, differentiating the frozen-rival payoff gives

$$\frac{\partial R_t(\tau)}{\partial \tau} = \frac{e^{-\int_t^\tau (r + \lambda(z)) dz} \pi(\tau, t)}{r} \left[D_i \log \pi(\tau, t) - \lambda(\tau) \right].$$

The prefactor is strictly positive. By Assumption 1(iv), the bracket is strictly decreasing in τ when the rival is fixed at t . Hence R_t is single-peaked on $[t, +\infty]$, with a unique maximizer. The unique-maximizer part of Lemma 8 therefore makes the maximizer $\chi(t)$ continuous in t . Since $\chi(t) \geq t$, setting $\chi(+\infty) := +\infty$ extends the continuity of χ to $[0, +\infty]$.

Stopping immediately is optimal exactly when the right derivative at t is nonpositive, which gives the displayed equivalence. That equivalence and continuity of the two rates make the set defining τ^R closed, so finite τ^R belongs to it. If $t < \tau^R$, then t does not belong to that set, so the displayed equivalence gives the strict inequality. $\square$

Proof of Proposition 2. If $\tau = 0$, the claim is immediate. For a date $t \geq 0$ at which a firm

stopped, define the sole survivor's crossing date $\chi(t) := \inf \{s \geq t : D_i \log \pi(s, t) \leq \lambda(s)\}$. By [Lemma 1](#), at every date in $[t, \chi(t)]$ a firm whose rival stopped at t has unique optimal stopping time $\chi(t)$. Moreover, $\chi(t) \geq \underline{\tau}$ for every $t \geq 0$: for every $s < \underline{\tau}$, log-supermodularity and the definition of $\underline{\tau}$ give $D_i \log \pi(s, t) \geq D_i \log \pi(s, 0) > \lambda(s)$, so no crossing occurs before $\underline{\tau}$. Finally, χ is continuous because the remaining firm's payoff is jointly continuous in its stopping date and in the rival's stopping date t on the compactified stopping-time space ([Lemma 7](#)), and its maximizer is unique; continuity follows by Berge's maximum theorem: every subsequential limit of $\chi(t_m)$ is optimal against the limit t and hence equals $\chi(t)$.

Before either firm has stopped, let X_i denote firm i 's candidate stopping date; X_1 and X_2 are independent because the firms' randomization devices are. By the preceding paragraph, the realized profile is $(X_1, \chi(X_1))$ if $X_1 < X_2$, $(\chi(X_2), X_2)$ if $X_2 < X_1$, and (X_1, X_1) if $X_1 = X_2$. On $\{X_1 \neq X_2\}$ the frontier equals $\chi(\min_i X_i) \geq \underline{\tau}$. Hence, up to null events, $\{\max_j \tau_j < \underline{\tau}\} \subseteq \{X_1 = X_2 < \underline{\tau}\}$, and by independence a positive-probability tie requires a common atom: some $t < \underline{\tau}$ with $\mathbb{P}(X_1 = t) > 0$ and $\mathbb{P}(X_2 = t) > 0$.

Suppose such an atom exists. On the device event $\{X_1 = t\}$, consider the deviation in which firm 1 replaces its candidate t by $t + \epsilon$, where $0 < \epsilon < \underline{\tau} - t$, and responds at $\chi(t')$ if firm 2 stops first at some date t' ; leave the strategy unchanged on $\{X_1 \neq t\}$. Because $\{X_1 = t\}$ is measurable with respect to firm 1's private randomization device and the response depends only on subsequent public records, this deviation is admissible. Conditional on $X_1 = t$: if $X_2 < t$, the outcome is unchanged. If $X_2 = t$, the profile changes from (t, t) to $(\chi(t), t)$, a strict gain $U_1(\chi(t), t) - U_1(t, t) > 0$, since $\chi(t) \geq \underline{\tau} > t$ and the frozen-rival payoff strictly increases up to its unique maximizer $\chi(t)$. The event $\{t < X_2 \leq t + \epsilon\}$ has probability converging to zero as $\epsilon \downarrow 0$, and payoffs are uniformly bounded ([Lemma 7](#)). If $X_2 > t + \epsilon$, the profile changes from $(t, \chi(t))$ to $(t + \epsilon, \chi(t + \epsilon))$, whose payoff effect vanishes as $\epsilon \downarrow 0$ by continuity of χ and of U_1 . The conditional expected gain from the deviation therefore converges to $\mathbb{P}(X_2 = t)[U_1(\chi(t), t) - U_1(t, t)] > 0$, so the deviation is strictly profitable for all sufficiently small ϵ , contradicting subgame perfection. Hence $\mathbb{P}(\max_j \tau_j < \underline{\tau}) = 0$. $\square$

Proof of [Proposition 3](#). We construct an equilibrium in which firm 1 is the designated first stopper. By [Lemma 1](#), once one firm stops at s , its rival's unique sequentially rational response is $\chi(s)$. Fix a no-stop history at date t . If firm 1 stops first at $s \geq t$, write

$$L_t(s) := U_1^t(s, \chi(s)), \quad s \in [t, +\infty],$$

for firm 1's continuation payoff from that choice. For $t \leq t' \leq s$, splitting the continuation

payoff from racing to s as evaluated at t at the intermediate time t' gives

$$L_t(s) = \int_t^{t'} e^{-r(u-t) - \int_t^u \lambda(z) dz} \pi(u, u) du + e^{-r(t'-t) - \int_t^{t'} \lambda(z) dz} L_{t'}(s).$$

The identity also holds at $s = +\infty$ by taking limits. Take the first date to be 0 and solve for $L_{t'}(s)$. The function $L_0(s) = U_1(s, \chi(s))$ is continuous by [Lemma 7](#) and continuity of χ , while the prefix and the positive coefficient vary continuously with t' . Hence $(t', s) \mapsto L_{t'}(s)$ is jointly continuous on the feasible set with compactified stopping-date coordinate. Applying [Lemma 8](#) to $f(t, s) = L_t(s)$ shows that the earliest maximizer $a(t) := \min \arg \max_{s \in [t, +\infty]} L_t(s)$ is well defined and Borel measurable. The same identity verifies that the selected plan does not change as time passes. Its first term and the positive coefficient on $L_{t'}(s)$ do not depend on s .

Prescribe the following profile: at every no-stop history at date t , firm 1 plans to stop at $a(t)$ and firm 2 waits; after a first stop at s , the remaining firm, active at date t' , stops at $\chi(s) \vee t'$. If, off path, both firms remain active beyond firm 1's previously prescribed stopping date, restart the same construction from the current date: the designated firm stops at $a(t)$, and firm 2 continues. These prescriptions define a valid pure strategy profile. Before firm 1 stops, $a(t') = a(t)$, so its planned stopping date remains unchanged. After a stop at s , the survivor continues to target $\chi(s)$, stopping immediately if that date has already passed.

We verify sequential rationality at an arbitrary two-active history at date t . Write $a := a(t)$. Since firm 2 waits until a stop occurs, any deviation by firm 1 reduces to a first stopping date $s \geq t$, with payoff $L_t(s)$; since $a(t)$ maximizes L_t , $L_t(s) \leq L_t(a)$. Thus the designated firm has no profitable deviation.

Along the prescribed path, firm 2 receives $F_t(a) := U_2^t(a, \chi(a))$. Before a the two firms have identical technologies; after a , firm 2's own technology is weakly higher than firm 1's, while its rival's technology is weakly lower, and both firms face the same survival kernel. Because flow profit rises in own technology and falls in rival technology, $F_t(a) \geq L_t(a)$. Consider any deviation by the waiting firm. If it stops first at $s < a$, firm 1 responds at $\chi(s)$; symmetry implies that the deviator's payoff is $L_t(s)$, and therefore $L_t(s) \leq L_t(a) \leq F_t(a)$ by the two preceding displays. If it stops at $s = a$, the firms stop together and its payoff is $U_2^t(a, a)$; once firm 1 stops at a , immediate stopping is feasible for firm 2 whereas $\chi(a)$ is its unique optimal response, so $U_2^t(a, a) \leq F_t(a)$. If its candidate is later than a , firm 1 stops first; for every successor choice $\tau \geq a$, optimality of $\chi(a)$ gives $U_2^t(a, \tau) \leq U_2^t(a, \chi(a)) = F_t(a)$. These cases cover every deviation by firm 2. (If $a = +\infty$, every finite s belongs to the first case and $s = +\infty$ gives the equilibrium tie.)

The same arguments apply at every history at which both firms are active, and the single-peakedness established at the start of the proof gives optimality at every history at which one firm has stopped. $\square$

Proof of Proposition 4. Fix $\tau \in \text{SPE}(\pi, \lambda)$. By Proposition 1, $\max_i \tau_i \leq \tau^R$ almost surely. It thus suffices to rule out two cases: both firms stop at the same date strictly below τ^R , or the firms stop at different dates and the later date is τ^R .

First suppose that one firm stops at time t while its rival remains active. By Lemma 1, the best response $\chi(t)$ is unique and continuous, and $\chi(\tau^R) = \tau^R$.¹⁷ Subgame perfection therefore requires a sole survivor to stop at $\chi(t)$ almost surely after its rival stops at t .

Before either firm has stopped, let X_i denote firm i 's candidate stopping date. The candidates X_1 and X_2 are independent because the firms' randomization devices are independent. The realized profile is

$$(\tau_1, \tau_2) = \begin{cases} (X_1, \chi(X_1)), & X_1 < X_2, \\ (\chi(X_2), X_2), & X_2 < X_1, \\ (X_1, X_1), & X_1 = X_2. \end{cases}$$

We now rule out the two cases.

Case 1: bunching below τ^R . Since $\chi(t) > t$ for every $t < \tau^R$, bunching below τ^R can occur only when $X_1 = X_2 < \tau^R$. Independence then implies that a positive-probability bunching event requires a common atom: there must be some $t < \tau^R$ such that $\mathbb{P}(X_1 = t) > 0$ and $\mathbb{P}(X_2 = t) > 0$. On the event $\{X_1 = t\}$, consider the deviation in which firm 1 replaces its candidate stopping date t with $t + \epsilon$, where $\epsilon > 0$ is small enough that $t + \epsilon < \tau^R$. If firm 2 stops first at some date s , firm 1 responds at $\chi(s)$. Leave the strategy unchanged on $\{X_1 \neq t\}$. Because $\{X_1 = t\}$ is measurable with respect to firm 1's private randomization device and the response depends only on subsequent public records, this deviation is admissible.

Conditional on $X_1 = t$, the deviation leaves the outcome unchanged when $X_2 < t$.

¹⁷Payoffs earned before t do not depend on τ , so $\chi(t)$ also uniquely maximizes $U_i(\tau, t)$ over $\tau \geq t$. To verify continuity directly, take $t_m \rightarrow t$ and a convergent subsequence, relabelled t_m , for which $\chi(t_m) \rightarrow q$. Compactness guarantees such a subsequence. Since $\chi(t_m) \geq t_m$, we have $q \geq t$. For any $\tau \geq t$, optimality gives $U_i(\chi(t_m), t_m) \geq U_i(\max\{\tau, t_m\}, t_m)$. Passing to the limit using Lemma 7 gives $U_i(q, t) \geq U_i(\tau, t)$. Thus q is optimal at t , and uniqueness gives $q = \chi(t)$. Every subsequential limit therefore equals $\chi(t)$, proving continuity. Finally, $\chi(t) \leq \tau^R$ for every $t \leq \tau^R$: for $\tau > \tau^R$, log-supermodularity gives $D_i \log \pi(\tau, t) \leq D_i \log \pi(\tau, \tau)$, and single-crossing gives $D_i \log \pi(\tau, \tau) < \lambda(\tau)$ strictly, so $\partial R_i(\tau)/\partial \tau < 0$ throughout $(\tau^R, +\infty)$ and the maximizer cannot exceed τ^R .

When $X_2 = t$, it changes the profile from (t, t) to $(\chi(t), t)$ and yields the strict gain $U_1(\chi(t), t) - U_1(t, t) > 0$, because t is not optimal against a rival frozen at t . The event $t < X_2 \leq t + \epsilon$ has probability converging to zero as $\epsilon \downarrow 0$, and payoffs are bounded by **Lemma 7**. Finally, when $X_2 > t + \epsilon$, the profile changes from $(t, \chi(t))$ to $(t + \epsilon, \chi(t + \epsilon))$, whose payoff effect converges to zero by continuity of χ and U_1 . The conditional expected gain from the deviation therefore converges to $\mathbb{P}(X_2 = t) [U_1(\chi(t), t) - U_1(t, t)] > 0$. For all sufficiently small ϵ , the deviation is strictly profitable, contradicting subgame perfection. Hence $\mathbb{P}(\tau_1 = \tau_2 < \tau^R) = 0$.

Case 2: spread profile reaching τ^R . Suppose, toward a contradiction, that $\mathbb{P}(\tau_1 < \tau^R = \tau_2) > 0$; where we assume that firm 1 is the laggard without loss since payoffs are symmetric. This implies that this event is contained in $\{X_1 < X_2, X_1 < \tau^R, \chi(X_1) = \tau^R\}$.

On the event $\{X_1 < \tau^R, \chi(X_1) = \tau^R\}$, consider the deviation in which firm 1 replaces its candidate stopping date X_1 with τ^R . If firm 2 stops first at y , firm 1 responds at $\chi(y)$. Leave the strategy unchanged on the complement of this event. Continuity of χ makes the event measurable with respect to firm 1's private randomization device, so this is an admissible deviation.

Fix a realization of firm 1's device to which this deviation applies, and write $x = X_1$. If $X_2 < x$, firm 2 stops first under both strategies and the outcome is unchanged. If $X_2 = x$, the outcome changes from (x, x) to $(\chi(x), x) = (\tau^R, x)$, which is strictly better because by construction $\chi(x)$ is the unique optimum against a rival frozen at x .

Now suppose $x < X_2 = y < \tau^R$. Originally, firm 1 stops at x and firm 2 responds at $\chi(x) = \tau^R$, giving the profile (x, τ^R) . Under the deviation, firm 2 stops first at y and firm 1 responds at $\chi(y)$, giving $(\chi(y), y)$. Since $x < y < \chi(y) \leq \tau^R$, firm 1's technology is everywhere weakly higher, firm 2's technology is everywhere weakly lower, so the disaster risk is weakly lower. At least one comparison is strict on a positive interval, so firm 1's payoff rises strictly.

Finally, suppose $X_2 \geq \tau^R$. If $\tau^R < +\infty$, the deviating profile is (τ^R, τ^R) : either the firms stop together at τ^R , or firm 1 stops there and firm 2 responds immediately because $\chi(\tau^R) = \tau^R$. If $\tau^R = +\infty$, this case is simply $X_2 = +\infty$, and both firms race forever under the deviation. In either case, relative to (x, τ^R) , firm 1's technology rises while its rival's technology and the frontier remain unchanged. The deviation is therefore weakly profitable for every realization and hence strictly profitable on the positive probability event above. This contradicts subgame perfection. Hence $\mathbb{P}(\tau_1 < \tau^R = \tau_2) = \mathbb{P}(\tau_2 < \tau^R = \tau_1) = 0$. Combining the two exclusions with **Proposition 1** yields the result. $\square$

Proof of Lemma 5. Fix $t \geq 0$. For each detection date $a \geq t$, let $\sigma^*(a) := \inf\{u \geq a :$

$D_i \log \pi(u, t) \leq \lambda(u)\}$, and define

$$\Phi_t(a) := \int_t^{\sigma^*(a)} \frac{\alpha(s)}{\alpha(t)} \pi(t, s) \, ds + \frac{\alpha(\sigma^*(a))}{\alpha(t)} \frac{\pi(t, \sigma^*(a))}{r}.$$

Thus, if $E_\omega \sim \text{Exp}(\omega)$, then $V_S(t, \omega) = \mathbb{E}[\Phi_t(t + E_\omega)]$.

Let $q_t := \inf\{u \geq t : D_i \log \pi(u, t) \leq \lambda(u)\}$. By **Assumption 1**(iv), $D_i \log \pi(u, t) - \lambda(u)$ is strictly decreasing in u , so $\sigma^*(a) = \max\{a, q_t\}$. Hence Φ_t is constant for $a \leq q_t$. For $a > q_t$, we have $\sigma^*(a) = a$, and differentiation gives $\Phi'_t(a) = \frac{\alpha(a)}{\alpha(t)} \frac{\pi(t, a)}{r} [D_j \log \pi(t, a) - \lambda(a)] < 0$, because $D_j \pi < 0$ and $\lambda \geq 0$. Therefore Φ_t is weakly decreasing in the detection date.

Let $Z \sim \text{Exp}(1)$ and couple the detection delays by $E_\omega = Z/\omega$. If $\omega_2 > \omega_1$, then $t + Z/\omega_2 < t + Z/\omega_1$ almost surely. Since Φ_t is weakly decreasing, $\Phi_t\left(t + \frac{Z}{\omega_2}\right) \geq \Phi_t\left(t + \frac{Z}{\omega_1}\right)$. Taking expectations shows that $V_S(t, \omega)$ is nondecreasing in ω .

If $\tau^R < \infty$, then $q_t < \infty$. Indeed, for every $u \geq \max\{t, \tau^R\}$, log-supermodularity and **Assumption 2** give $D_i \log \pi(u, t) \leq D_i \log \pi(u, u) \leq \lambda(u)$. Moreover, Φ_t is strictly decreasing above q_t . The coupled inequality is therefore strict on an event of positive probability, so $V_S(t, \omega)$ is strictly increasing in ω . This proves part (a).

As $\omega \downarrow 0$, the coupling gives $E_\omega \rightarrow \infty$ almost surely. Since $\sigma^*(a) \geq a$, it follows that $\sigma^*(a) \rightarrow \infty$ as $a \rightarrow \infty$. Hence $\int_t^{\sigma^*(a)} \frac{\alpha(s)}{\alpha(t)} \pi(t, s) \, ds \rightarrow \int_t^\infty \frac{\alpha(s)}{\alpha(t)} \pi(t, s) \, ds$. Because $\pi(t, \cdot)$ is decreasing, $0 \leq \frac{\alpha(\sigma^*(a))}{\alpha(t)} \frac{\pi(t, \sigma^*(a))}{r} \leq \frac{\alpha(\sigma^*(a))}{\alpha(t)} \frac{\pi(t, t)}{r} \rightarrow 0$. Therefore $\Phi_t(a) \rightarrow \int_t^\infty \frac{\alpha(s)}{\alpha(t)} \pi(t, s) \, ds$. Since Φ_t is nonnegative and weakly decreasing, it is bounded by $\Phi_t(t) < \infty$. Bounded convergence therefore gives $V_S(t, 0) := \lim_{\omega \downarrow 0} V_S(t, \omega) = \int_t^\infty \frac{\alpha(s)}{\alpha(t)} \pi(t, s) \, ds$. For every $s > t$, strict monotonicity in own technology gives $\pi(t, s) < \pi(s, s)$. Therefore

$$V_S(t, 0) = \int_t^\infty \frac{\alpha(s)}{\alpha(t)} \pi(t, s) \, ds < \int_t^\infty \frac{\alpha(s)}{\alpha(t)} \pi(s, s) \, ds = V_R(t),$$

which proves part (b).

Finally, suppose $t \geq \tau^R$. For every $a \geq t$, log-supermodularity and **Assumption 2** imply $D_i \log \pi(a, t) \leq D_i \log \pi(a, a) \leq \lambda(a)$. Thus $\sigma^*(a) = a$, and

$$\Phi_t(a) = \int_t^a \frac{\alpha(s)}{\alpha(t)} \pi(t, s) \, ds + \frac{\alpha(a)}{\alpha(t)} \frac{\pi(t, a)}{r}.$$

As $\omega \rightarrow \infty$, the coupling gives $E_\omega \rightarrow 0$ almost surely. Since $0 \leq \Phi_t(t + E_\omega) \leq \Phi_t(t) = \pi(t, t)/r$, bounded convergence yields $\lim_{\omega \rightarrow \infty} V_S(t, \omega) = \frac{\pi(t, t)}{r}$.

It remains to compare this limit with $V_R(t)$. Since $D_j \pi < 0$, we have $\pi(s, s) < \pi(s, t)$ for every $s > t$, and hence $V_R(t) \leq \int_t^\infty \frac{\alpha(s)}{\alpha(t)} \pi(s, t) \, ds$. For $s \geq t \geq \tau^R$, log-supermodularity

and **Assumption 2** give $D_i \log \pi(s, t) \leq \lambda(s)$. Consequently,

$$-\frac{d}{ds} \left[\frac{\alpha(s)}{\alpha(t)} \frac{\pi(s, t)}{r} \right] = \frac{\alpha(s)}{\alpha(t)} \frac{\pi(s, t)}{r} [r + \lambda(s) - D_i \log \pi(s, t)] \geq \frac{\alpha(s)}{\alpha(t)} \pi(s, t).$$

Moreover, $\frac{\alpha(s)}{\alpha(t)} \pi(s, t) \leq \pi(t, t) e^{-r(s-t)} \rightarrow 0$, where the inequality follows from $D_i \log \pi(s, t) \leq \lambda(s)$. Integrating the preceding derivative inequality from t to infinity therefore gives $\int_t^\infty \frac{\alpha(s)}{\alpha(t)} \pi(s, t) ds \leq \frac{\pi(t, t)}{r}$. It follows that $V_R(t) \leq \frac{\pi(t, t)}{r} = \lim_{\omega \rightarrow \infty} V_S(t, \omega)$. The first inequality is strict for $t > \tau^R$, proving part (c). $\square$

Proof of Proposition 7. We construct a symmetric assessment. At every no-signal history, each firm races forever. After a first signal at date a , the active firm assigns probability one to the rival having stopped at a . At every subsequent date $s \geq a$ at which it remains active, it stops at

$$\inf \{u \geq s : D_i \log \pi(u, a) \leq \lambda(u)\},$$

with the convention $\inf \emptyset = +\infty$. By **Lemma 3**, this is the unique optimal continuation under the assigned belief. In particular, if crossing never occurs, racing forever is optimal.

We verify optimality at a no-signal history at date t . The rival never stops under the proposed strategy, so no signal arrives and racing forever is worth $V_R(t)$. Consider first a pure deviation that plans to stop at some finite $\sigma \geq t$ absent a signal. The rival then receives its signal at a date $a \geq \sigma$. Under the off-path belief above, its prescribed response is $\inf \{u \geq a : D_i \log \pi(u, a) \leq \lambda(u)\}$. If instead the rival knew the deviator's true stopping date σ , its response would be $\inf \{u \geq a : D_i \log \pi(u, \sigma) \leq \lambda(u)\}$. Because $a \geq \sigma$, log-supermodularity implies that the prescribed response is weakly later. A stopped firm is hurt by any delay in its rival's stop (**Lemma 12**), so the deviator's continuation payoff after stopping at σ is at most $V_S(\sigma, \omega)$. By **Lemma 6**, its total deviation payoff is thus at most

$$\begin{aligned} \int_t^\sigma \frac{\alpha(s)}{\alpha(t)} \pi(s, s) ds + \frac{\alpha(\sigma)}{\alpha(t)} V_S(\sigma, \omega) &= V_R(t) + \frac{\alpha(\sigma)}{\alpha(t)} \Delta(\sigma, \omega) \\ &\leq V_R(t). \end{aligned}$$

The last inequality follows because $\omega < \bar{\omega}$ and **Lemma 5(a)** imply $\Delta(\sigma, \omega) \leq 0$ at every date. A deviation that races forever yields $V_R(t)$, and randomized deviations cannot improve on this (by linearity of expectations). Hence racing forever is optimal at every no-signal history.

The specified beliefs obey Bayes' rule on path. Signal histories are off path, and the assigned point belief is admissible. Play is sequentially rational. Thus the assessment is a

perfect Bayesian equilibrium, and both firms race forever almost surely. $\square$

Proof of Proposition 9. Throughout the proof, abbreviate the two effective discount rates $\delta := r + \lambda - \alpha + \beta = 1/V_R$ (along the race path) and $\rho := r + \lambda + \beta = 1/V_S(0)$ (stopped beside a racing rival).

Part (i) Crazy firms never stop, and a signal can arrive only if some firm has already stopped, so the first stop in any equilibrium—if there is one—must occur at a *quiet* history (active, no signal received). It therefore suffices to show that no rational firm ever stops at a quiet history: then no stop ever occurs, no signal is ever sent, and no responsive stop ever follows. Fix an equilibrium and let F_j be the (sub-)distribution of rational firm j 's stopping time along its quiet histories, $S_j := 1 - F_j$. Since an active firm i has sent no signal, its rival's behavior is governed by F_j , and a stop at c remains undetected with probability $e^{-\omega(t-c)}$; firm i 's posterior at a quiet history at date t therefore puts weights proportional to

$$\underbrace{1 - p,}_{\text{crazy}} \quad \underbrace{p S_j(t)}_{\text{rational, active}}, \quad \underbrace{p e^{-\omega(t-c)} dF_j(c)}_{\text{rational, stopped at } c, \text{ undetected}}.$$

Compare stopping now (denoted STOP) with the feasible plan WAIT: continue, and stop immediately when a signal arrives. Cell by cell, per unit of current flow:

Crazy. Stopping yields $V_S(0)$; WAIT races forever, worth V_R : the deviation gains $V_R - V_S(0)$.

Stopped at c . STOP ends the risk at once, worth $e^{\beta(t-c)}/r$ (the factor is the flow advantage over a rival frozen at c); WAIT yields $e^{\beta(t-c)}V_W(\omega)$, so the deviation loses $e^{\beta(t-c)}[1/r - V_W(\omega)]$.

Rational, active. Let τ be the rival's residual quiet-path stopping time and $X \sim \text{Exp}(\omega)$ the detection lag. STOP ends the risk at $X \wedge \tau$, worth $V_S(0) + [1/r - V_S(0)]\mathbb{E}[e^{-\rho(X \wedge \tau)}]$; WAIT races until τ and then exits upon detection, yielding $V_R + [V_W(\omega) - V_R]\mathbb{E}[e^{-\delta\tau}]$ (recall $V_W(\omega) > V_R$). Using $\mathbb{E}_X[e^{-\rho(X \wedge \tau)}] = \frac{\omega}{\omega + \rho} + \frac{\rho}{\omega + \rho}e^{-(\omega + \rho)\tau}$, the gap (STOP minus WAIT) on this cell is $\mathbb{E}[\varphi(\tau)]$ where

$$\varphi(\tau) = V_S(0) - V_R + [1/r - V_S(0)] \left[\frac{\omega}{\omega + \rho} + \frac{\rho}{\omega + \rho} e^{-(\omega + \rho)\tau} \right] - [V_W(\omega) - V_R]e^{-\delta\tau}.$$

Evaluating φ' shows that is quasi-convex, so for every (random) τ ,

$$\mathbb{E}[\varphi(\tau)] \leq \max\{\varphi(0), \varphi(\infty)\} = \max\{1/r - V_W(\omega), V_S(\omega) - V_R\}.$$

Finally, on the stopped cell the weight carries the damping $e^{-\omega(t-c)} \leq e^{-\beta(t-c)}$ (here

$\omega \geq \beta$), so its weighted gap is at most $[1/r - V_W(\omega)] p \int dF_j$. Summing the cells and using $S_j(t) + \int_0^t dF_j \leq 1$ and $p = \ell(1-p)$,

$$\begin{aligned} (\text{stop}) - (\text{WAIT}) &\propto -(1-p)[V_R - V_S(0)] + p \cdot \max\{1/r - V_W(\omega), V_S(\omega) - V_R\} \\ &= (1-p) \left[\frac{\ell}{\underline{\ell} \wedge \bar{\ell}} - 1 \right] [V_R - V_S(0)] < 0, \end{aligned}$$

since $\min\{\underline{\ell}, \bar{\ell}\} = [V_R - V_S(0)] / \max\{1/r - V_W(\omega), V_S(\omega) - V_R\}$ and $\ell < \underline{\ell} < \bar{\ell}$. So at every quiet history stopping is strictly dominated by WAIT: rational firms never stop at quiet histories, under any pure or mixed rule, and by the reduction above no stop ever occurs.

Part (ii) For (b), consider the profile in which no firm ever stops at a quiet history, and any active firm stops immediately upon a signal. On path no signal arrives, so quiet histories carry the prior odds ℓ and no undetected-stop weight. A deviation is a plan “stop at σ unless a signal arrives first”; no signal ever comes, so its value is

$$(1 - e^{-\delta\sigma}) V_R + e^{-\delta\sigma} \frac{\ell V_S(\omega) + V_S(0)}{1 + \ell},$$

since a stop at σ is reciprocated upon detection by a rational rival (worth $V_S(\omega)$; the response is forced by $\alpha < \lambda$) and never by a crazy one (worth $V_S(0)$). For $\ell < \bar{\ell}$ the gamble is strictly below V_R , so every $\sigma < \infty$ is strictly worse than racing forever, and mixtures are averages of such plans; at $\ell = \bar{\ell}$ every finite plan is weakly worse—payoff-equivalent at the optimum—so the profile remains an equilibrium.

For (a), consider the bunched profile: rational firms stop at $t = 0$, crazy firms never, and any active firm stops immediately upon a signal. A rational firm’s deviations are again plans $\sigma \in [0, \infty]$ (continue until σ or the first signal, whichever is first). With $X \sim \text{Exp}(\omega)$ the date at which the deviator would detect a rational rival’s stop at 0, the plan’s value is

$$W(\sigma) = p \underbrace{\mathbb{E} \left[\frac{1 - e^{-(r+\lambda-\alpha)(X \wedge \sigma)}}{r + \lambda - \alpha} + \frac{e^{-(r+\lambda-\alpha)(X \wedge \sigma)}}{r} \right]}_{\text{rational rival: race alone, stop at } X \wedge \sigma} + (1-p) \underbrace{\left[V_R - (V_R - V_S(0)) e^{-\delta\sigma} \right]}_{\text{crazy rival: race, stop at } \sigma}.$$

Then $W(0) = p/r + (1-p) V_S(0)$ is the bunching payoff, $W(\infty) = p V_W(\omega) + (1-p) V_R$, and

$$W'(\sigma) = -p \frac{\lambda - \alpha}{r} e^{-(\omega + r + \lambda - \alpha)\sigma} + (1-p) \frac{V_R - V_S(0)}{V_R} e^{-\delta\sigma}.$$

For $\omega \geq \beta$ the negative term decays faster, so W' changes sign at most once, from $-$ to $+$: W is maximized at an endpoint. And $W(0) \geq W(\infty)$ iff $p[1/r - V_W(\omega)] \geq (1-p)[V_R -$

$V_S(0)]$, i.e. $\ell \geq \underline{\ell}$. So no plan improves on stopping at once.¹⁸

Part (iii) Since $(\bar{\ell}/\ell)^2 < 1$, the bound already excludes the never-stopping outcome. To see this, conjecture that neither firm ever stops. After every quiet history, each firm continues to believe that their rival is rational with odds ℓ since the conjectured profile assigns zero probability to its rival stopping. Stopping immediately yields

$$\frac{\ell V_S(\omega) + V_S(0)}{1 + \ell} > V_R,$$

and is therefore a profitable deviation. Note here that our assumption that $\alpha < \lambda$ ensures that a rational rival finds it dominant to stop after receiving news. For the bound, fix an equilibrium and, with F_j, S_j as in part (i), let $s_j := \lim_{t \rightarrow \infty} S_j(t)$. Conditional on both firms rational, the risk never dies only if neither firm ever stops spontaneously (responsive stops require a prior stop), an event of probability $s_1 s_2$. If $s_j \leq \bar{\ell}/\ell$ for both firms, then $s_1 s_2 \leq (\bar{\ell}/\ell)^2$ and we are done. Otherwise, say $s_1 > \bar{\ell}/\ell$; we claim $s_2 = 0$, so ruin has probability zero. On firm 2's quiet path, the odds on "rational and active" are $\ell S_1(t) \geq \ell s_1 > \bar{\ell}$, so the immediate-stop gamble

$$u(t) := \frac{\ell S_1(t) V_S(\omega) + V_S(0)}{1 + \ell S_1(t)} \geq V_R + \eta \quad \text{for some } \eta > 0$$

(the gamble increases in the odds because $V_S(\omega) > V_R$, which holds whenever $\bar{\ell} < \infty$; the undetected-stop cell only raises the stop value further). Consider any continuation plan of firm 2 at a late quiet history t . Relative to the deterministic benchmark "race to σ , then stop," worth

$$(1 - e^{-\delta\sigma}) V_R + e^{-\delta\sigma} u(t + \sigma) \leq u(t) - (1 - e^{-\delta\sigma}) \eta$$

(using that u decreases along the quiet path), the plan can gain only through two channels that vanish as $t \rightarrow \infty$: the rival's residual spontaneous concession of conditional probability $(S_1(t) - s_1)/S_1(t) \rightarrow 0$, since $S_1 \downarrow s_1 > 0$, and the undetected-stop weight, which converges to zero; each contributes at most $1/r$ per unit of weight. Hence for t large, any plan that delays stopping beyond $t + w$ with positive probability is strictly dominated by stopping at once, and sequential rationality forces firm 2 to stop by $t + w$ almost surely on its quiet path—contradicting $s_2 > 0$.

Finally, zero is attained: since $\ell > \bar{\ell} > \underline{\ell}$, the bunched profile of part (ii)(a) remains an equilibrium here, and conditional on both firms rational it ends the risk immediately. $\square$

¹⁸Off the bunched path, a deviator continues optimally against its updated posterior, which affects no on-path incentive. The same computation at any common date T shows delayed bunching is also an equilibrium: the environment is stationary and nothing is learned before T .

B TECHNICAL DETAILS ON STRATEGIES

This appendix defines strategies and equilibria. When we discuss firm i , let j denote its rival. [Appendix B.1](#) defines the extensive form as well as mixing under perfect observability. It then shows that a firm's stopping date uses only information the firm has observed. [Appendix B.2](#) gives the corresponding definition when firms observe delayed signals and defines perfect Bayesian equilibrium for [Section 5](#) and [Online Appendix I](#). We omit “almost surely” when it is clear from context.

B.1 The extensive form under perfect observability. We now describe play after each observed stop. At any public history, each active firm uses its current private draw to choose a candidate stopping date. The firm cannot wait to see its rival's action at the same instant before choosing its own. When a rival stops, the firm observes a new public history, makes a fresh draw, and uses its response plan. This construction follows [Riedel and Steg \(2017\)](#) and [Laraki et al. \(2005\)](#). Here public information changes only when a firm stops, so at most two such events occur.

Records, plans, and strategies. A *record* is either empty, $p = \emptyset$, or a pair $p = (j, t_j)$ stating that firm j stopped at t_j . Set $t(\emptyset) := 0$ and $t(j, t_j) := t_j$. Under perfect observability the record is public. A *plan* for an active firm i is a Borel-measurable map $\sigma_i(p, z) \in [t(p), +\infty]$, where $z \in [0, 1]$. When record p begins, the firm's candidate date is $\sigma_i(p, Z_i^{(p)})$, unless the record changes first. The draw $Z_i^{(p)}$ is the relevant entry of $\mathbf{Z}_i^*(t) := (Z_{i,1}, \mathbf{1}\{\tau_j \leq t\}Z_{i,2})$: It equals $Z_{i,1}$ at the empty record and $Z_{i,2}$ after the rival stops.

At a history (t, p) , date t has arrived, record p has not changed, and the firm remains active. A strategy also specifies a distribution $G_i(\cdot \mid t, p)$ for its candidate date at every such history. This distribution assigns probability one to $[t, +\infty]$, and its probability of every Borel set varies measurably with (t, p) . When record p begins, the draw $Z_i^{(p)}$ generates distribution $G_i(B \mid t(p), p) = \int_0^1 \mathbf{1}\{\sigma_i(p, z) \in B\} dz$. We impose the following dynamic-consistency condition: for every $t \leq s$ and every Borel set B , whenever $G_i([s, +\infty] \mid t, p) > 0$,

$$G_i(B \mid s, p) = \frac{G_i(B \cap [s, +\infty] \mid t, p)}{G_i([s, +\infty] \mid t, p)}. \quad (\text{DC})$$

Equation [\(DC\)](#) says that the firm does not redraw its candidate merely because time has passed. If it reaches s , its candidate is the old candidate conditional on being at least s . If the denominator is zero, the earlier plan would never reach s , but the strategy must still say what the firm does there. A fresh private uniform draw then selects a new candidate

from $G_i(\cdot \mid s, p)$; **Lemma 9** shows that one draw can generate this distribution. The firm observes the draw, so its later choices may depend on it. After selecting the new candidate, the firm again keeps it as time passes according to (DC). These extra draws are used only at histories that the firm's earlier plan would never reach. Play from the initial history uses only $Z_{i,1}$ and, after the rival stops, $Z_{i,2}$.

Within a fixed record, the firm's draw fixes its candidate date. The firm therefore cannot wait to see whether its rival also stops at date t before choosing its own date- t action. If the rival stops at t , a new record begins and activates a new plan, which may also prescribe stopping at t . By **Lemma 7**, this response has the same payoff as stopping at $t + \varepsilon$ in the limit as $\varepsilon \downarrow 0$, holding terminal stopping dates fixed. A strategy specifies a plan even for records that equilibrium play never reaches, so subgame perfection can test behavior after deviations.

Mixing. A mixed plan at record p is a distribution over candidate stopping dates in $[t(p), +\infty]$. When the record begins, a fresh uniform draw selects one date from this distribution. Draws used by different firms and at different records are independent. The firm observes its selected date, but its rival observes only its eventual stopping action. The next lemma shows that one uniform draw can implement any such distribution.

Lemma 9 (Implementing a stopping-date distribution). Let $\mathcal{R}$ be the set of records, and let the earliest feasible date $b(\rho)$ vary measurably with the record ρ . Suppose that $G(\rho, \cdot)$ is a distribution on $[b(\rho), +\infty]$ and that, for every Borel set B , its probability $G(\rho, B)$ varies measurably with ρ . Then there is a Borel-measurable map $Q : \mathcal{R} \times [0, 1] \rightarrow \overline{\mathbb{R}}_+$ such that $Q(\rho, Z)$ has distribution $G(\rho, \cdot)$ whenever Z is uniform and independent of the information available when record ρ begins. Conversely, every such map defines a family of distributions with these measurability properties.

Proof. This is the usual inverse-distribution construction. We first map all possible dates, including $+\infty$, into the unit interval. We then apply the inverse distribution function there and map the result back. Map $\overline{\mathbb{R}}_+$ to $[0, 1]$ by $\psi(x) = x/(1+x)$ and $\psi(+\infty) = 1$. Let F_ρ be the distribution function of the image of $G(\rho, \cdot)$ under ψ , and define

$$Q(\rho, z) := \psi^{-1}(\inf\{y \in [\psi(b(\rho)), 1] : F_\rho(y) \geq z\}), \quad z \in [0, 1].$$

For every finite x , $\{(\rho, z) : Q(\rho, z) \leq x\} = \{(\rho, z) : b(\rho) \leq x, z \leq G(\rho, [0, x])\}$, so Q is Borel-measurable. It lies in $[b(\rho), +\infty]$ and has the required distribution. Conversely, $G(\rho, B) = \int_0^1 \mathbf{1}\{Q(\rho, z) \in B\} dz$ has the required measurability for every Borel-measurable Q . $\square$

Lemma 10 (Outcomes). Every strategy profile σ determines a unique pair of stopping dates $\tau(\sigma) = (\tau_1, \tau_2)$. Each firm's stopping date uses only information that the firm has observed: formally, $\tau_i(\sigma)$ is an $(\mathcal{F}_t^i)_t$ -stopping time.

Proof. Construction. Call $\sigma_i(p, Z_i^{(p)})$ firm i 's candidate date at record p . At the empty record, compare the two candidates. If both are $+\infty$, both firms race forever. If the candidates tie at a finite date, both firms stop there. Otherwise, the firm with the earlier candidate stops first. Its rival observes that stop, activates the record (j, t_j) , and uses its second draw to select a new candidate. The rival stops at that date or races forever if the candidate is $+\infty$. These cases give a unique stopping pair.

The stopping date uses only observed information. Fix i and t . Firm i first computes its candidate at the empty record. If its rival does not stop strictly before that date, the candidate is its stopping date, with a tie producing a simultaneous stop. Otherwise, when the rival stops first at τ_j , firm i activates record (j, τ_j) , uses $Z_{i,2}$ to compute its second candidate, and stops at that date. This reproduces the construction above. On $\{\tau_j > t\}$ the computation uses only $Z_{i,1}$, so $Z_{i,2}$ enters only after the rival's stop, as the filtration requires. Hence $\{\tau_i \leq t\} \in \sigma(\mathbf{Z}_i^*(t), (\mathbf{1}\{\tau_j \leq s\})_{s \leq t}) = \mathcal{F}_t^i$. $\square$

Subgames and equilibrium. For a record p and a date $t \geq t(p)$, the *subgame after* (t, p) is the continuation game in which no new stop has occurred since record p began. At the empty record both firms remain active. At a record (j, t_j) , firm j is frozen at t_j and only firm i remains active. An active firm's strategy includes its current plan and, at the empty record, its plan after a later rival stop. Given terminal stopping dates (τ_i, τ_j) , where a firm that has already stopped keeps its recorded date, firm i 's continuation payoff is $U_i^t(\tau_i, \tau_j)$.

An active rival's candidate has distribution $G_j(\cdot \mid t, p)$. At a history reached with positive probability, this is its original candidate conditional on being at least t . A draw reserved for a later rival stop has not yet been used. At a history the earlier plan would never reach, or after a firm has continued past its candidate, the strategy uses the fresh draw specified above. Condition (DC) then makes the firm keep that new candidate as time passes.

Definition 3 (Subgame-perfect equilibrium). A strategy profile σ is a *subgame-perfect equilibrium* if, after every history (t, p) and for every active firm i , no replacement of firm i 's plans from that history onward raises its expected continuation payoff

$$\mathbb{E}[U_i^t(\tau_i, \tau_j) \mid \text{firm } i\text{'s information at } (t, p)].$$

The replacement must leave firm i 's earlier play unchanged, but may depend on every

private draw the firm has already observed and may specify its response to every later observed stop. The expectation is over the rival's unobserved candidate and all future draws. We write $\text{SPE}(\pi, \lambda)$ for the set of outcomes $\tau(\sigma)$ generated by subgame-perfect strategy profiles.

B.2 Strategies under imperfect observability. Under imperfect observability, a firm sees delayed signals rather than its rival's stopping date. The first private draw governs play before any signal, and a fresh second draw is made after the first detection. This subsection defines these plans and the perfect Bayesian equilibrium used in [Section 5](#) and [Online Appendix I](#).

Primitives. The game uses independent private $\text{Uniform}[0, 1]$ draws to mix and independent detection clocks $E_1, E_2 \sim \text{Exp}(\omega)$. The draw $Z_{i,1}$ governs firm i 's initial plan, and $Z_{i,2}$ is drawn after its first detection. At any history that its earlier plan would never reach, the firm makes another fresh uniform draw. All private draws and detection clocks are mutually independent. The clock E_j is attached to firm j 's stop: if firm j stops at $\tau_j < \infty$, the signal process $Y_t^j := \mathbf{1}\{\tau_j + E_j \leq t\}$ jumps at the *detection date* $\tau_j + E_j$, at which point its rival learns, conclusively, that firm j has stopped; if $\tau_j = +\infty$, then $Y^j \equiv 0$. A jump of Y^j reveals the event that firm j has stopped but not the date τ_j at which it did so.

Signal records, plans, and strategies. A *signal record* for firm i is either empty, $q = \emptyset$, or a pair $q = (j, a)$ stating that firm i detected firm j at date a . Set $a(\emptyset) := 0$ and $a(j, a) := a$. Unlike the public records under perfect observability, a signal record is private and lists the detection date, not the rival's stopping date. A *plan* is a Borel-measurable map $(q, z) \mapsto \sigma_i(q, z) \in [a(q), +\infty]$. When record q begins, $\sigma_i(q, Z_i^{(q)})$ is firm i 's candidate date; the firm stops there unless a signal changes its record first. The private draws available by date t are $Z_i^*(t) := (Z_{i,1}, Y_t^j Z_{i,2})$. The draw $Z_i^{(q)}$ used at record q equals $Z_{i,1}$ before detection and $Z_{i,2}$ after detection. A strategy specifies the initial plan and a response plan for every possible signal date. A plan cannot condition on an event at the same instant: a detection at a activates the response plan, which may itself prescribe stopping at a . By [Lemma 7](#), the payoff from stopping immediately after a converges to the payoff from this response.

At a history (t, q) , a strategy also specifies a distribution $G_i(\cdot \mid t, q)$ for the firm's current candidate date. This distribution assigns probability one to $[t, +\infty]$, and its probability of every Borel set varies measurably with (t, q) . When record q begins, the draw $Z_i^{(q)}$ generates distribution $G_i(B \mid a(q), q) = \int_0^1 \mathbf{1}\{\sigma_i(q, z) \in B\} dz$. If the history is reached

with positive probability, this is simply the original candidate date conditioned on that candidate date being at least t , so it satisfies (DC) with q in place of p . The firm holds onto that candidate for as long as it lies in the future. If the firm deviates by continuing past that candidate, or if its earlier plan would never reach the history, a fresh private uniform draw selects a new candidate from the distribution specified there, as in Lemma 9. The firm observes this draw, and equilibrium requires that the new plan be optimal. After selecting the new candidate, the firm again keeps it as time passes according to (DC).

Lemma 11 (Outcomes under imperfect observability). Together with the detection clocks, every strategy profile determines a unique stopping pair up to a probability-zero event. Each firm's stopping date uses only information that the firm has observed. Formally, $\tau_i(\sigma)$ is an $(\mathcal{F}_t^i)_t$ -stopping time for $\mathcal{F}_t^i = \sigma(\mathbf{Z}_i^*(t), (Y_s^j)_{s \leq t})$.

Proof. Starting from date zero, list each scheduled stop and each pending detection. Execute the earliest event. A stop starts that firm's detection clock; a detection activates the rival's second plan if the rival is still active. Continue until both firms have stopped or no finite event remains, in which case every active firm races forever. At most two stops and two detections can occur. Tied candidate dates produce a simultaneous stop. A pending detection almost surely does not tie a stop that was already scheduled because exponential delays have continuous distributions. If the response plan chooses the detection date itself, execute that stop immediately after the detection. Thus the stopping outcome is unique except on a probability-zero event.

Each firm i can compute its stopping date from its two private draws and the path Y^j of its rival's signal. It first evaluates the plan at the empty record. If Y^j does not jump strictly before the resulting candidate, that candidate is τ_i . If the signal arrives first, the firm evaluates its second plan at the signal record and follows that candidate. This reproduces the construction above, so $\{\tau_i \leq t\} \in \mathcal{F}_t^i$. $\square$

Histories, beliefs, and equilibrium. An *observable history* for an active firm i at date t is the date together with its signal record: $h = (t, q)$. The firm also knows its own draws and current candidate date, and optimality is tested using all this information. Beliefs can be written as a function of h alone because the firm's draws are independent of whether its rival has stopped.

The rival may have stopped at some date $y \leq t$, leaving its technology frozen while its signal remains pending, or it may still be active. A belief μ_i assigns probabilities to these possibilities at every observable history. To evaluate a continuation plan, firm i averages over them and over any pending detection delay. If the rival is active, its signal

record is empty because firm i has not stopped, so its current candidate has distribution $G_j(\cdot \mid t, \emptyset)$.

Definition 4 (Perfect Bayesian equilibrium). A *perfect Bayesian equilibrium* is a strategy profile σ together with beliefs (μ_1, μ_2) such that:

- (i) *Sequential rationality*: after every observable history and every realization of the firm's private information, the firm cannot gain by changing its plans from that point onward while leaving its earlier play unchanged;
- (ii) *Consistent beliefs*: at every no-signal history reached with positive probability, beliefs follow Bayes' rule. After a detection at date a , if the Bayes denominator implied by the rival's strategy is positive, the posterior is pinned down by Bayes' rule. If the denominator is zero such that detection at a is impossible under the rival's strategy, then we only impose the restriction that the belief must assign probability one to stopping dates $y \leq a$.

Before any signal, the rival's strategy and the absence of news determine the firm's belief. After a signal, the firm uses the Bayes posterior at every detection date for which that posterior has a positive denominator. The existence proof computes this posterior in [Lemma 14](#). At feasible dates not covered pointwise by the fixed-point condition, it selects a measurable best response to this posterior. Only when the rival's strategy makes detection at date a impossible—that is, when the Bayes denominator is zero—may the firm choose another belief; the construction in [Online Appendix I](#) then assigns δ_a , meaning that the rival stopped at the detection date a .

Concession and response plans. A strategy has two parts. The *concession plan* is the distribution of the stopping date before any signal. The *response plan* assigns a distribution of stopping dates in $[a, +\infty]$ to each first-signal date a . [Lemma 9](#) shows how the private draws generate both plans. Together with the instructions at histories that the earlier plan would never reach, these plans define a complete strategy. At any history at which firm i is active it has sent no signal, so an active rival's signal record is empty. The belief μ_i therefore reduces to a distribution over whether, and when, the rival has stopped, as computed in [Lemma 14](#).

Thus mixing has a simple meaning: after each signal record, the strategy assigns a distribution over candidate stopping dates, and a fresh draw selects one candidate if that record is reached. By the checking-deviations argument in the main text, we can verify optimality by checking continuation plans that leave earlier play unchanged and make no further random choices.

ONLINE APPENDIX TO ‘RACING TO RUIN’

DREW FUDENBERG* ANDREW KOH[†]

FOR ONLINE PUBLICATION ONLY

I EQUILIBRIUM EXISTENCE UNDER IMPERFECT TRANSPARENCY

This appendix proves **Proposition 6** for the two-firm game. Set $\alpha(s) = e^{-rs - \int_0^s \lambda(u) du}$. For a realized stopping pair $(\tau_1, \tau_2) \in \overline{\mathbb{R}}_+^2$, write firm 1’s realized payoff as

$$U(\tau_1, \tau_2) := U_1(\tau_1, \tau_2) = \int_0^\infty e^{-rs - \int_0^s \lambda(u) \mathbf{1}_{\{u \leq \tau_1 \vee \tau_2\}} du} \pi(s \wedge \tau_1, s \wedge \tau_2) ds,$$

firm 2’s realized payoff being $U(\tau_2, \tau_1)$ by symmetry. Write $\bar{U} := \sup |U|$, and let η_U denote a modulus of continuity for U .

Overview of proof. **Figure 9** gives a roadmap. If $\omega < \bar{\omega}$, the never-stopping equilibrium from **Proposition 7** proves existence. The argument below covers all remaining news rates, including $\omega = \bar{\omega}$. It proceeds in two steps. First, we solve a smaller game in which each firm chooses a distribution of initial concession dates and a response rule after a signal (**Appendices I.1 to I.3**). Second, we complete that strategy at histories that occur with probability zero and verify optimality after every possible history (**Appendices I.4 to I.6**).

*Department of Economics, MIT, drewf@mit.edu

[†]Department of Economics, Columbia University, andrew.koh@columbia.edu




Figure 9: Logical structure of the existence proof. Each arrow points from an input to the result that uses it.

Step 1: solve the reduced game. A two-firm strategy has two economically relevant parts: an initial *concession date* if no signal arrives, and a *response date* after detecting the rival's stop. **Appendix I.1** shows that, after a signal at a , every optimal response lies between a and $\chi(a)$. The rule K assigns a distribution over that interval for each a , while G is the distribution of the initial concession date. These pairs (G, K) form a compact convex set. Payoffs vary continuously because stops are not observed immediately and detection dates have continuous distributions. A fixed-point theorem therefore gives a symmetric reduced equilibrium (G^*, K^*) : G^* assigns probability only to optimal concession dates, and K^* chooses an optimal response at almost every detection date that can occur.

Step 2: complete the dynamic strategy. The reduced equilibrium specifies behavior only at histories that matter with positive probability. Perfect Bayesian equilibrium also requires a plan and a belief at probability-zero histories. At such detection dates, we prescribe the latest potentially optimal response $\chi(a)$ and the belief that the rival stopped at a . This change does not disturb the optimality of G^* (Lemma 18). We then show that a concession date that is optimal initially remains optimal after any history with no signal (Lemma 19). After the latest date at which the rival might concede, the firm's belief no longer changes and its stopping problem is time-consistent (Lemma 20). Appendix I.6 combines these pieces and checks every class of history.

I.1 The response problem. Note that

$$\alpha(y) \pi(y, 0) \leq e^{-ry} \sup_{v \leq y} \left\{ e^{-\int_0^v \lambda(z) dz} \pi(v, 0) \right\} \xrightarrow{y \rightarrow \infty} 0, \quad (1)$$

where the limit holds because the running supremum is nondecreasing in y . If the right side were at least $\varepsilon > 0$ along $y_m \uparrow \infty$, we could pass to a subsequence with $y_{m+1} > y_m + 1$. Monotonicity of the supremum would then give

$$\int_{y_m}^{y_{m+1}} e^{-rs} \sup_{v \leq s} \left\{ e^{-\int_0^v \lambda(z) dz} \pi(v, 0) \right\} ds \geq \varepsilon e^{-r}$$

on each of these disjoint intervals, contradicting the integrability in Assumption 1(v). Because flow profits decrease in the rival's technology, $\pi(a, b) \leq \pi(a, 0)$ for every $b \geq 0$. Thus (1), which bounds discounted profits against a rival fixed at 0, also bounds discounted profits along every stopping history considered below.

Lemma 12 (Rival's delay hurts). Fix $x \in \mathbb{R}_+$. The map $y \mapsto U(x, y)$ is continuous on $[x, +\infty]$ and strictly decreasing: for every $x \leq y < y' \leq +\infty$, $U(x, y') < U(x, y)$.

Proof. Continuity is Lemma 7. For $y \in (x, \infty)$ we have $\tau_1 \vee \tau_2 = y$, so $\Lambda_s = \int_0^{s \wedge y} \lambda$ and

$$U(x, y) = \int_0^y \alpha(s) \pi(s \wedge x, s) ds + \frac{\alpha(y) \pi(x, y)}{r}.$$

Differentiating and using $\alpha'(y) = -(r + \lambda(y))\alpha(y)$,

$$\begin{aligned} \frac{\partial}{\partial y} U(x, y) &= \alpha(y) \pi(x, y) + \frac{\alpha'(y) \pi(x, y) + \alpha(y) D_j \pi(x, y)}{r} \\ &= \frac{\alpha(y)}{r} \left[D_j \pi(x, y) - \lambda(y) \pi(x, y) \right] < 0, \end{aligned}$$

since $D_i \pi < 0$, $\lambda \geq 0$, and $\pi > 0$. Strict monotonicity on (x, ∞) follows; continuity extends the strict comparison to the endpoints $y = x$ and $y' = +\infty$. $\square$

After a detection the game is a single-agent stopping problem against a frozen rival whose technology may be unknown. For $c \geq 0$ define the *crossing date*

$$\chi(c) := \inf \{y \geq c : D_i \log \pi(y, c) \leq \lambda(y)\}.$$

If a signal arrives at a and reveals a rival frozen at t , the optimal response is $a \vee \chi(t)$. For a detection date $a > 0$ and a belief $\nu \in \Delta([0, a])$ over the rival's frozen technology, the value of stopping at $y \in [a, +\infty]$ is

$$w(y; a, \nu) := \int_{[0, a]} \left[\int_a^y \frac{\alpha(s)}{\alpha(a)} \pi(s, c) \, ds + \frac{\alpha(y)}{\alpha(a)} \frac{\pi(y, c)}{r} \right] \nu(dc),$$

with the terminal term absent at $y = +\infty$, and the *response window* is $\mathcal{T}(a) := [a, \chi(a)]$. From the definition of w , for every $a \leq a' \leq y$ (note $\Delta([0, a]) \subseteq \Delta([0, a'])$),

$$w(y; a, \nu) = \int_a^{a'} \frac{\alpha(s)}{\alpha(a)} \left(\int \pi(s, c) \nu(dc) \right) ds + \frac{\alpha(a')}{\alpha(a)} w(y; a', \nu); \quad (2)$$

the first term does not depend on y , so the two continuation values differ only by an additive constant and a positive scale factor. They therefore rank every stopping date in $[a', +\infty]$ in the same way. The next lemma collects the properties of χ , w , and $\mathcal{T}$ that we use.

Lemma 13 (The response problem). Fix $a > 0$ and $\nu \in \Delta([0, a])$. Then:

(i) χ is nondecreasing and continuous on $\mathbb{R}_+$, with $c \leq \chi(c) \leq \tau^R$ and $\chi(c) > c$ for $c < \tau^R$, and $\chi(c) = c$ for $c \geq \tau^R$; consequently $\mathcal{T}(a) = \{a\}$ for $a \geq \tau^R$ and $\mathcal{T}(a) \subseteq [a, \tau^R]$ for $a \leq \tau^R$.

(ii) For $a \leq \tau^R$ and every $y \in [a, +\infty]$,

$$0 \leq w(y; a, \nu) \leq \bar{w} := \frac{1}{\alpha(\tau^R)} \left[\int_0^\infty \alpha(s) \pi(s, 0) \, ds + \frac{1}{r} \sup_{v \geq 0} \alpha(v) \pi(v, 0) \right] < \infty. \quad (3)$$

(iii) $w(\cdot; a, \nu)$ is continuous on $[a, +\infty]$ and strictly decreasing on $[\chi(a), +\infty]$; hence

$$\operatorname{argmax}_{y \in [a, +\infty]} w(y; a, \nu)$$

is a nonempty compact subset of $\mathcal{T}(a)$, and any response outside $\mathcal{T}(a)$ is strictly suboptimal under every belief in $\Delta([0, a])$.

- (iv) For the point belief δ_c with $c \leq a$, the unique maximizer is $a \vee \chi(c)$. In particular $\chi(a)$ is the unique optimal response under δ_a , and it is the *latest* response that is optimal under any belief in $\Delta([0, a])$.
- (v) w is jointly continuous in (y, a, ν) , where beliefs with supports in varying intervals are compared as elements of $\Delta(\overline{\mathbb{R}}_+)$ with the weak topology, on the domain $\{(y, a, \nu) : a > 0, y \in [a, +\infty], \nu([0, a]) = 1\}$; consequently, by Berge's maximum theorem, $(a, \nu) \mapsto \max_y w(y; a, \nu)$ is continuous and $(a, \nu) \mapsto \operatorname{argmax}_y w(y; a, \nu)$ is upper hemicontinuous with compact values.

Proof. Part (i). For $c \leq \tau^R$, log-supermodularity (**Assumption 1**(iii)) and the definition of τ^R give $D_i \log \pi(\tau^R, c) \leq D_i \log \pi(\tau^R, \tau^R) \leq \lambda(\tau^R)$, so the defining set is nonempty and $\chi(c) \leq \tau^R$. For $c \geq \tau^R$, **Assumption 2** gives $D_i \log \pi(c, c) \leq \lambda(c)$, so $\chi(c) = c$; for $c < \tau^R$, $D_i \log \pi(c, c) > \lambda(c)$ and continuity give $\chi(c) > c$. Monotonicity: for $c \leq c'$, continuity at the crossing and log-supermodularity give $D_i \log \pi(\chi(c'), c) \leq D_i \log \pi(\chi(c'), c') \leq \lambda(\chi(c'))$, and $\chi(c') \geq c' \geq c$, so $\chi(c) \leq \chi(c')$. Continuity: write $g(y, c) := D_i \log \pi(y, c) - \lambda(y)$, which is continuous, and strictly decreasing in y on $[c, \infty)$ by **Assumption 1**(iv). If $\chi(c) > c$, fix $0 < \varepsilon < (\chi(c) - c)/2$. Then $g(\chi(c) - \varepsilon, c) > 0 > g(\chi(c) + \varepsilon, c)$. For c' sufficiently close to c , we have $c' < \chi(c) - \varepsilon$ and both strict inequalities persist, so $\chi(c') \in (\chi(c) - \varepsilon, \chi(c) + \varepsilon)$. If $\chi(c) = c$ then $c \geq \tau^R$; on $[\tau^R, \infty)$ the map χ coincides with the identity, and continuity at the junction $c = \tau^R$ follows from the squeeze $c' \leq \chi(c') \leq \tau^R$ for $c' < \tau^R$, which forces $\chi(c') \rightarrow \tau^R = \chi(\tau^R)$ as $c' \uparrow \tau^R$. The consequences for $\mathcal{T}$ are immediate.

Part (ii). Use $\pi(\cdot, c) \leq \pi(\cdot, 0)$ and $1/\alpha(a) \leq 1/\alpha(\tau^R)$ for $a \leq \tau^R$: the integral is finite by **Assumption 1**(v), whose integrand dominates $\alpha(s) \pi(s, 0)$, and the supremum by (1).

Part (iii). For a fixed rival technology $c \leq a$, the contribution to the derivative of w from that technology is

$$\frac{\alpha(y)}{r \alpha(a)} \pi(y, c) \left[D_i \log \pi(y, c) - \lambda(y) \right].$$

For $y \geq \chi(a)$, log-supermodularity and the definition of $\chi(a)$ give

$$D_i \log \pi(y, c) \leq D_i \log \pi(y, a) \leq \lambda(y),$$

with the second inequality strict for $y > \chi(a)$. Integrating against ν , the derivative of w is strictly negative on $(\chi(a), \infty)$, which yields the strict decrease and the inclusion $\operatorname{argmax} \subseteq$

$\mathcal{T}(a)$. Continuity of w up to $y = +\infty$ follows from dominated convergence: the flow is dominated using (1) and the terminal term satisfies $\alpha(y)\pi(y, c) \leq \alpha(y)\pi(y, 0) \rightarrow 0$. Nonemptiness and compactness of the argmax follow from continuity on the compact $[a, +\infty]$ and closedness of the argmax.

Part (iv). Under the point belief δ_c , the derivative displayed in Part (iii) has the sign of $D_i \log \pi(y, c) - \lambda(y)$, which is strictly decreasing in y by [Assumption 1\(iv\)](#). The point-belief payoff is therefore single-peaked with unique maximizer $a \vee \chi(c)$. Part (iii) shows that no response later than $\chi(a)$ is optimal under any belief, while $\chi(a)$ is uniquely optimal under δ_a . Hence $\chi(a)$ is the latest response that can be optimal.

Part (v). Let $b(y, a, c)$ denote the payoff term inside the integral defining $w(y; a, \nu)$:

$$b(y, a, c) := \int_a^y \frac{\alpha(s)}{\alpha(a)} \pi(s, c) \, ds + \frac{\alpha(y)}{\alpha(a)} \frac{\pi(y, c)}{r},$$

with the terminal term omitted when $y = +\infty$. The function b is jointly continuous in (y, a, c) and dominated by an integrable envelope via (1), so w is jointly continuous in (y, a) uniformly over c in compact sets and affine-continuous in ν . Joint continuity follows: let $(y_m, a_m, \nu_m) \rightarrow (y, a, \nu)$ with $\nu_m([0, a_m]) = 1$ and set $A := \sup_m a_m < \infty$. Extend b beyond the relevant supports by replacing it with $b(y, a, c \wedge A)$, a bounded function that is continuous in c on $\overline{\mathbb{R}}_+$. Then

$$\begin{aligned} |w(y_m; a_m, \nu_m) - w(y; a, \nu)| &\leq \sup_{c \leq A} |b(y_m, a_m, c) - b(y, a, c)| \\ &\quad + \left| \int b(y, a, c \wedge A) \, d(\nu_m - \nu)(c) \right|, \end{aligned}$$

where the first term vanishes by the uniform convergence just noted and the second by weak convergence. The domain correspondence $a \mapsto [a, +\infty]$ is continuous: upper hemicontinuity is closedness of $\{(a, y) : y \geq a\}$ in the compact $\overline{\mathbb{R}}_+^2$, and lower hemicontinuity holds because any $y \in [a, +\infty]$ is the limit of $y \vee a_m \in [a_m, +\infty]$ along $a_m \rightarrow a$. Berge's maximum theorem ([Aliprantis and Border, 2006](#), Theorem 17.31) therefore applies. $\square$

Response rules. We describe every date in the response interval $[a, \chi(a)]$ by a number $\theta \in [0, 1]$: $\chi_\theta(a) := (1 - \theta)a + \theta\chi(a)$. Thus $\theta = 0$ means stop when the signal arrives, while $\theta = 1$ means wait until $\chi(a)$, the latest potentially optimal response. For $a \geq \tau^R$, both dates equal a . A *response kernel* K is simply a measurable rule that assigns a probability distribution over θ to each detection date a . In the terminology of [Aliprantis and Border \(2006, Theorem 19.7 and Definition 19.8\)](#), it is a measurable transition from detection dates to probability distributions over responses. Upon detection, the firm draws θ

from $K(a, \cdot)$ and stops at $\chi_\theta(a)$. By [Lemma 13\(iii\)](#), no optimal response lies outside this interval.

To define convergence of response rules, we associate each kernel K with a joint distribution over detection dates and responses. If $\tau^R = 0$, the response parameter is immaterial at every date, so we take the response-rule space to be a singleton. Suppose $\tau^R > 0$, let $A := [0, \tau^R]$, and associate with each kernel K the probability measure $Q_K(da, d\theta) := \frac{da K(a, d\theta)}{\tau^R}$ on $A \times [0, 1]$. The first marginal of Q_K is the uniform distribution on A . Conversely, disintegration recovers from every joint law with this marginal a measurable kernel $K(a, d\theta)$, unique for almost every a . We therefore treat two kernels as the same when they agree almost everywhere and let $\mathcal{K}$ denote the resulting set. We say that $K_m \rightarrow K$ when $Q_{K_m} \rightarrow Q_K$ weakly. The set $\mathcal{K}$ is nonempty and convex. It is also compact and metrizable. Fix a bounded Borel function f on $A \times [0, 1]$ that is continuous except on a null set $D \subseteq A$. Note that since Q_K has the uniform distribution on A , it assigns zero probability to $D \times [0, 1]$. Hence if Q_{K_m} converges weakly to Q_K then the Portmanteau lemma gives:

$$\int_A \int_0^1 f(a, \theta) K_m(a, d\theta) da = \tau^R \int f dQ_{K_m} \longrightarrow \tau^R \int f dQ_K = \int_A \int_0^1 f(a, \theta) K(a, d\theta) da \quad (J)$$

I.2 The reduced game. Let $\mathcal{G} := \Delta(\overline{\mathbb{R}}_+)$ be the set of probability distributions over concession dates, including the possibility $+\infty$. We use weak convergence on this set, which makes $\mathcal{G}$ compact, convex, and metrizable. A *reduced strategy* is a pair (G, K) : draw the no-signal concession date from G , and use response rule K after the first detection. By [Lemma 9](#) and [Appendix B.2](#), fresh independent uniform draws implement every such pair in the dynamic game.

We first record the beliefs that Bayes' rule assigns against a rival playing a reduced strategy. For $G \in \mathcal{G}$ set

$$N_G(a) := \int_{[0, a]} e^{\omega y} G(dy), \quad \nu_a^G(dy) := \frac{e^{\omega y} \mathbf{1}\{y \leq a\} G(dy)}{N_G(a)} \quad \text{whenever } N_G(a) > 0.$$

Lemma 14 (Beliefs). Fix $\{i, j\} = \{1, 2\}$ and a profile in which firm j plays the reduced strategy (G, K) , and let $X_j \sim G$ denote its no-signal concession date. Consider firm i at a history at which it is active and has received no signal by t . Then:

- (i) any stop by firm j so far occurred at X_j ; the conditional law of X_j given the history assigns mass proportional to $e^{-\omega(t-y)} G(dy)$ on $[0, t]$ and to $G(dy)$ on $(t, +\infty]$;

- (ii) the finite part of the date a of firm i 's first signal has the possibly defective density $a \mapsto \omega e^{-\omega a} N_G(a)$ on $(0, \infty)$, and the signal date has an atom of size $G(\{+\infty\})$ at $+\infty$ (this law does not depend on firm i 's own concession plan, so long as firm i is active); given a finite first signal at a with $N_G(a) > 0$, the conditional law of X_j is ν_a^G ;
- (iii) both conditional laws are unaffected by conditioning on survival of the disaster, and they do not depend on firm i 's own (past or planned) concession behavior.

Proof. Part (i). Because firm i has not stopped, firm j has received no signal and therefore stops, if at all, at its no-signal concession date X_j . Given $X_j = y$, the probability that firm i receives no signal by t is $e^{-\omega(t-y)}$ when $y \leq t$ and is one when $y > t$. While firm i is active, its technology at each date s equals s , so a rival stop at $y \leq s$ leaves the frontier at $\max(s, y) = s$. The survival probability through t is therefore the common factor $e^{-\int_0^t \lambda(s) ds}$ for every realization of X_j . Bayes' rule gives the conditional law in (i).

Part (ii). Given $X_j = y < +\infty$, the finite first-signal date has density $a \mapsto \omega e^{-\omega(a-y)} \mathbf{1}\{a > y\}$. Integrating this density over $y \in [0, a]$ gives $\omega e^{-\omega a} N_G(a)$. If $X_j = +\infty$, no signal arrives, which gives the atom of size $G(\{+\infty\})$ at $+\infty$. At a finite signal date a , Bayes' rule assigns weight proportional to $e^{\omega y} G(dy)$ on $[0, a]$, yielding ν_a^G .

Part (iii). The common survival factor calculated in Part (i) proves that conditioning on survival does not change either posterior. Moreover, X_j and its detection clock are independent of firm i 's private draws and concession plan. Hence firm i 's own past or planned behavior does not change either conditional law. $\square$

The fixed-point proof must choose an optimal response as a measurable function of the detection date. The next lemma proves that this is possible. For a with $N_G(a) > 0$ define

$$\Theta^*(a, G) := \left\{ \theta \in [0, 1] : \chi_\theta(a) \in \operatorname{argmax}_{y \in [a, +\infty]} w(y; a, \nu_a^G) \right\}.$$

Lemma 15 (Measurability in the detection date). Fix $G \in \mathcal{G}$. Then:

- (i) N_G is nondecreasing and right-continuous, and on $\{N_G > 0\}$ the map $a \mapsto \nu_a^G$ is Borel for the weak topology on $\Delta(\overline{\mathbb{R}}_+)$;
- (ii) on $(\{N_G > 0\} \cap (0, \tau^R)) \times [0, 1]$, the map $(a, \theta) \mapsto w(\chi_\theta(a); a, \nu_a^G)$ is Borel in a and continuous in θ (a Carathéodory integrand), and $a \mapsto \max_{y \in [a, +\infty]} w(y; a, \nu_a^G) = \max_{\theta \in [0, 1]} w(\chi_\theta(a); a, \nu_a^G)$ is Borel;
- (iii) $\Theta^*(\cdot, G)$ is nonempty- and compact-valued, and one can choose an optimal $\theta^*(a)$ as a Borel function of a .

Proof. Part (i). Monotonicity is clear, and right-continuity is dominated convergence: as $a_m \downarrow a$, $e^{\omega y} \mathbf{1}\{y \leq a_m\} \downarrow e^{\omega y} \mathbf{1}\{y \leq a\}$ pointwise, dominated by $e^{\omega \sup_m a_m}$. For bounded continuous f , $a \mapsto \int f \, d\nu_a^G = N_G(a)^{-1} \int_{[0,a]} f(y) e^{\omega y} G(dy)$ is then a ratio of Borel functions of a , hence Borel; and Borel measurability against every bounded continuous f is Borel measurability for the weak topology.

Part (ii). The map $a \mapsto (a, \nu_a^G)$ is Borel by (i), and $(a, \theta) \mapsto \chi_\theta(a)$ is jointly continuous (Lemma 13(i)); composing with the jointly continuous w (Lemma 13(v)) gives the Carathéodory property. The displayed identity holds because $\operatorname{argmax}_y w \subseteq \mathcal{T}(a) = \{\chi_\theta(a) : \theta \in [0, 1]\}$ (Lemma 13(iii)), and the maximum is then Borel as the value function of a maximization of a Carathéodory function over the fixed compact $[0, 1]$ (Aliprantis and Border, 2006, Measurable Maximum Theorem 18.19).

Part (iii). By the displayed identity, $\Theta^*(a, G) = \operatorname{argmax}_{\theta \in [0, 1]} w(\chi_\theta(a); a, \nu_a^G)$. The same theorem shows that this set is nonempty, compact, and measurable in a , and that an optimal $\theta^*(a)$ can be chosen as a Borel function. $\square$

Outcomes and payoffs. Fix concession dates x for firm 1 and y for firm 2. Use the detection clocks E_1, E_2 from Appendix B.2, where E_i is attached to firm i 's stop, and let θ_1, θ_2 be the response draws. For a detection at $a \geq \tau^R$, the response parameter is immaterial because $\chi_\theta(a) = a$ for every θ . The following three cases compare which firm plans to concede first and whether the resulting signal arrives before the other firm's planned concession. The realized stopping pair is:

- if $x < y$: $\tau_1 = x$; with $a_2 := x + E_1$, $\tau_2 = y$ if $y < a_2$, and $\tau_2 = \chi_{\theta_2}(a_2)$ if $a_2 \leq y$;
- if $y < x$: $\tau_2 = y$; with $a_1 := y + E_2$, $\tau_1 = x$ if $x < a_1$, and $\tau_1 = \chi_{\theta_1}(a_1)$ if $a_1 \leq x$;
- if $x = y$: $\tau_1 = \tau_2 = x$.

The first stopper concedes without receiving a signal—its rival stops strictly later, so no signal precedes its concession—and the later firm either concedes on schedule or, if the signal arrives first, responds optimally. Note that the law Q_K determines $K(a, d\theta)$ for almost every date a , and we can choose its values arbitrarily at remaining dates. Firm 1's reduced payoff against a rival playing (G, K) , when firm 1 concedes at x and responds per the same kernel K , is

$$U_K(x, y) := \mathbb{E}[U(\tau_1, \tau_2)], \quad U(x; G, K) := \int_{\mathbb{R}_+} U_K(x, y) G(dy),$$

the expectation running over $(E_1, E_2, \theta_1, \theta_2)$ with $\theta_i \sim K(a_i, \cdot)$. Note that θ_1 enters only on $\{y < x\}$ and θ_2 only on $\{x < y\}$, so $U_K(x, y)$ is *affine* in K : averaging two response rules

averages their payoffs. Since all concession draws, response draws, and detection clocks are independent, $U(x; G, K)$ equals firm 1's expected payoff in the dynamic game under the corresponding plans (Lemmas 7 and 9).

The reduced game treats concession and response choices separately. A response rule that is optimal at almost every detection date remains optimal after any concession deviation: detection dates have atomless laws, and their posteriors do not depend on the firm's concession plan (Lemma 14(ii)–(iii)). Thus $U(x; G, K)$ can hold the response rule K fixed while varying the concession date x . Correlating the response draw with the concession draw also cannot help: after a detection, the abandoned concession draw no longer affects payoffs, and a fresh response draw attains the pointwise optimum.

Lemma 16 (Continuity of reduced payoffs). (i) The family $\{U_K : K \in \mathcal{K}\}$ is uniformly bounded by $\bar{U}$ and uniformly equicontinuous on $\bar{\mathbb{R}}_+^2$.

(ii) For each fixed (x, y) , the map $K \mapsto U_K(x, y)$ is continuous on $\mathcal{K}$.

(iii) U is jointly continuous on $\bar{\mathbb{R}}_+ \times \mathcal{G} \times \mathcal{K}$.

Proof. Part (i). Boundedness is immediate. For equicontinuity, fix K and consider the branch $y < x$ (the branch $x < y$ is symmetric; the diagonal is treated at the end). Substituting $a = y + E_2$,

$$U_K(x, y)|_{y < x} = \underbrace{e^{-\omega(x-y)} U(x, y)}_{\text{clock rings after } x} + \int_y^x \left[\int_0^1 U(\chi_{\theta_1}(a), y) K(a, d\theta_1) \right] \omega e^{-\omega(a-y)} da,$$

using $\chi_\theta(a) = a$ on $[\tau^R, \infty)$. Perturbing x to $x' > x$ changes the first term by at most $\eta_U(d(x, x')) + \bar{U} (e^{-\omega(x-y)} - e^{-\omega(x'-y)})$ and adds to the integral a piece over (x, x') of mass $e^{-\omega(x-y)} - e^{-\omega(x'-y)}$; since

$$e^{-\omega(x-y)} - e^{-\omega(x'-y)} \leq 1 - e^{-\omega(x'-x)} \quad \text{and} \quad e^{-\omega(x-y)} - e^{-\omega(x'-y)} \leq e^{-\omega(x-y)},$$

the change is bounded by $\eta_U(d(x, x')) + 2\bar{U} \min\{1 - e^{-\omega(x'-x)}, e^{-\omega(x-y)}\}$. The minimum is small uniformly: if $x \leq L$ then $x' - x \leq e^{L+1}d(x, x')$ for $d(x, x')$ small, making the first entry small, while if $x > L$ and $y \leq x - \sqrt{L}$ the second entry is at most $e^{-\omega\sqrt{L}}$; and if $x > L$, $y > x - \sqrt{L}$, then all realized coordinates under either concession date exceed $L - \sqrt{L}$, so the two payoffs lie within $2\eta_U(2e^{-(L-\sqrt{L})})$ of one another. Choosing L large and then $d(x, x')$ small yields a modulus independent of (y, K) .

Now we compare the two integrals at the same detection date. For $z < x$, write

$p_z(a) := \omega e^{-\omega(a-z)}$. The inner expectation is bounded by $\bar{U}$ and, uniformly in (a, K) ,

$$\left| \int_0^1 U(\chi_\theta(a), y') K(a, d\theta) - \int_0^1 U(\chi_\theta(a), y) K(a, d\theta) \right| \leq \eta_U(d(y, y')).$$

Let $y' = y + \delta < x$. On the common interval $[y', x]$ we have $p_{y'}(a) = e^{\omega\delta} p_y(a)$. The probability removed from $[y, y']$ is $1 - e^{-\omega\delta}$, and $\int_{y'}^x (p_{y'}(a) - p_y(a)) da \leq 1 - e^{-\omega\delta}$. It follows that the detection terms differ by at most $\eta_U(d(y, y')) + 2\bar{U}(1 - e^{-\omega\delta})$. The no-detection terms differ by at most $\eta_U(d(y, y')) + \bar{U}(1 - e^{-\omega\delta})$. Hence, whenever $y < y' < x$,

$$|U_K(x, y') - U_K(x, y)| \leq 2\eta_U(d(y, y')) + 3\bar{U}(1 - e^{-\omega\delta}),$$

and the case $y' < y < x$ follows by interchanging y and y' .

Uniformity in $d(y, y')$ follows by considering two cases: if $y \leq L$, then $\delta \leq e^{L+1}d(y, y')$ for $d(y, y')$ small, making every δ -dependent term above small; if $y > L$, then every realized coordinate under either concession date exceeds L , so the two payoffs lie within $2\eta_U(2e^{-L})$ of one another. Choosing L large and then $d(y, y')$ small yields a modulus independent of (x, K) . Perturbations across the diagonal $x = y$ compose the two branch moduli with the observation that, as $x \rightarrow y$ from either side, both branch formulas converge to $U(y, y)$ at rate $\eta_U(d(x, y)) + 2\bar{U}(1 - e^{-\omega|x-y|})$, again uniformized by the two-case split: from below, the integral term has vanishing mass and the first term tends to $U(y, y)$; from above, symmetrically.

Part (ii). In the displayed formula the kernel appears only through

$$\int_0^{\tau^R} \int_0^1 f(a, \theta) K(a, d\theta) da \quad \text{and} \quad f(a, \theta) = U(\chi_\theta(a), y) \omega e^{-\omega(a-y)} \mathbf{1}\{y < a < x\},$$

a bounded measurable function continuous in θ ; recall $(a, \theta) \mapsto \chi_\theta(a)$ and U are continuous. If $x < y$, the kernel-dependent term instead has integrand

$$U(x, \chi_\theta(a)) \omega e^{-\omega(a-x)} \mathbf{1}\{x < a < y\}$$

which is also bounded and continuous except possibly at the boundary dates $a = x$ and $a = y$. Since these dates form a Lebesgue-null set, (J) implies that the kernel-dependent integral varies continuously with K . The remaining no-detection term does not depend on K . If $x = y$, the firms stop simultaneously and the payoff is independent of K . Thus $K \mapsto U_K(x, y)$ is continuous for every fixed (x, y) .

Part (iii). Let $(x_m, G_m, K_m) \rightarrow (x, G, K)$ and estimate

$$\begin{aligned} |U(x_m; G_m, K_m) - U(x; G, K)| &\leq \underbrace{\sup_y |U_{K_m}(x_m, y) - U_{K_m}(x, y)|}_{\rightarrow 0 \text{ by (i)}} \\ &\quad + \left| \int U_{K_m}(x, y) \, d(G_m - G) \right| + \underbrace{\left| \int [U_{K_m} - U_K](x, y) \, G(dy) \right|}_{\rightarrow 0 \text{ by (ii) and domination}}. \end{aligned}$$

For the middle term, suppose along a subsequence it stays above $\varepsilon > 0$; by (i) and Arzelà–Ascoli we may pass to a further subsequence with $U_{K_m}(x, \cdot) \rightarrow h$ uniformly on $\overline{\mathbb{R}}_+$. Then

$$\left| \int U_{K_m}(x, y) \, d(G_m - G)(y) \right| \leq 2 \|U_{K_m}(x, \cdot) - h\|_\infty + \left| \int h \, d(G_m - G) \right| \rightarrow 0,$$

a contradiction. $\square$

I.3 A symmetric equilibrium of the reduced game.

Lemma 17 (Reduced equilibrium). There exists $(G^*, K^*) \in \mathcal{G} \times \mathcal{K}$ such that

- (i) $G^*(\operatorname{argmax}_{x \in \overline{\mathbb{R}}_+} U(x; G^*, K^*)) = 1$; equivalently, the support of G^* is contained in $\operatorname{argmax}_x U(x; G^*, K^*)$;
- (ii) for Lebesgue-almost every $a \in (0, \tau^R)$ with $N_{G^*}(a) > 0$,

$$K^*\left(a, \left\{ \theta \in [0, 1] : \chi_\theta(a) \in \operatorname{argmax}_{y \in [a, +\infty]} w(y; a, v_a^{G^*}) \right\}\right) = 1.$$

Proof. Define, for $(G, K) \in \mathcal{G} \times \mathcal{K}$,

$$\begin{aligned} \operatorname{BR}_G(G, K) &:= \left\{ G' \in \mathcal{G} : G'(\operatorname{argmax}_x U(x; G, K)) = 1 \right\}, \\ \operatorname{BR}_K(G) &:= \left\{ K' \in \mathcal{K} : K'(a, \Theta^*(a, G)) = 1 \text{ for a.e. } a \in (0, \tau^R) \text{ with } N_G(a) > 0 \right\}, \end{aligned}$$

with $\Theta^*(a, G)$ as defined before [Lemma 15](#). We verify the hypotheses of the Kakutani–Fan–Glicksberg theorem ([Aliprantis and Border, 2006](#), Corollary 17.55) on the nonempty, compact, convex set $\mathcal{G} \times \mathcal{K}$ (a subset of the product of two locally convex spaces).

The fixed-point theorem requires four properties: best replies must exist; each set of best replies must be convex and compact; and best replies must vary continuously in the closed-graph sense. We check these properties in turn.

Best replies exist. $\operatorname{argmax}_x U(\cdot; G, K)$ is nonempty and closed by [Lemma 16\(iii\)](#) and compactness of $\overline{\mathbb{R}}_+$, so $\operatorname{BR}_G \neq \emptyset$. For BR_K , [Lemma 15\(iii\)](#) lets us choose an optimal response $\theta^*(a)$ as a Borel function of the detection date. The rule $K'(a, \cdot) = \delta_{\theta^*(a)}$ (extended arbitrarily where $N_G(a) = 0$) therefore lies in $\operatorname{BR}_K(G)$.

Best-reply sets are convex and compact. Both sets require a distribution to put full probability on a fixed set of optimal choices. This condition is linear in the distribution, so each best-reply set is convex. The surrounding strategy spaces are compact. The closed-graph arguments below imply that each best-reply set is closed and hence compact.

Concession best replies vary continuously. $G' \in \operatorname{BR}_G(G, K)$ iff $\int U(x; G, K) G'(dx) \geq \max_x U(x; G, K)$. Both sides are continuous in (G', G, K) : the left side because $(G', G, K) \mapsto \int U dG'$ inherits continuity from [Lemma 16\(iii\)](#) (uniform equicontinuity in x makes the integrand vary continuously in the sup norm, so the same Arzelà–Ascoli argument as in the proof of [Lemma 16\(iii\)](#) applies), the right side by Berge. The inequality therefore defines a closed set.

Response best replies vary continuously. If $\tau^R = 0$ then $\mathcal{K}$ is a singleton and there is nothing to prove. Suppose $\tau^R > 0$. Define the *optimality gap*—the average payoff loss from using K' instead of an optimal response—by

$$\Gamma(G, K') := \int_0^{\tau^R} (N_G(a) \wedge 1) \left[\max_{y \in [a, +\infty]} w(y; a, \nu_a^G) - \int_0^1 w(\chi_\theta(a); a, \nu_a^G) K'(a, d\theta) \right] e^{-a} da.$$

Where $N_G(a) = 0$, set the payoff-loss term inside the integral—the best-response payoff minus the payoff under K' —equal to 0. This payoff loss is nonnegative, so $\Gamma \geq 0$, and $K' \in \operatorname{BR}_K(G)$ iff $\Gamma(G, K') = 0$. Let $(G_m, K_m) \rightarrow (G, K)$ with $K_m \in \operatorname{BR}_K(G_m)$; we show $\Gamma(G, K) = 0$. First, for all but countably many a (the atoms of G), $N_{G_m}(a) \rightarrow N_G(a)$, and where $N_G(a) > 0$ also $\nu_a^{G_m} \rightarrow \nu_a^G$ weakly (the integrand $y \mapsto e^{\omega y} \mathbf{1}\{y \leq a\}$ is bounded and G -a.s. continuous at non-atoms a). Hence, writing $I(a, \theta; G') := (N_{G'}(a) \wedge 1) w(\chi_\theta(a); a, \nu_a^{G'})$ and $W(a; G') := (N_{G'}(a) \wedge 1) \max_y w(y; a, \nu_a^{G'})$ (both 0 where $N_{G'}(a) = 0$), [Lemma 13\(v\)](#) gives, for a.e. a , $W(a; G_m) \rightarrow W(a; G)$ and $\sup_{\theta \in [0, 1]} |I(a, \theta; G_m) - I(a, \theta; G)| \rightarrow 0$ (where $N_G(a) = 0$, the factor $N_{G_m}(a) \wedge 1 \rightarrow 0$ forces both convergences because the underlying response payoffs are bounded by $\bar{w}$ from (3)). Since the θ -suprema bound the kernel integrals uniformly over K_m , all terms lying in $[0, \bar{w}]$, dominated convergence yields

$$\begin{aligned} \Gamma(G_m, K_m) &= \tilde{\Gamma}_m + o(1), \\ \tilde{\Gamma}_m &:= \int_0^{\tau^R} \left[W(a; G) - \int_0^1 I(a, \theta; G) K_m(a, d\theta) \right] e^{-a} da. \end{aligned}$$

To apply (J), we verify that the integrand is continuous outside atoms of G . At any non-atom a with $N_G(a) > 0$, both $N_G(a)$ and the posterior ν_a^G vary continuously with a , so $I(a, \theta; G)e^{-a}$ is continuous in (a, θ) . If $N_G(a) = 0$, then $N_G(a') \wedge 1 \rightarrow 0$ as $a' \rightarrow a$; because response payoffs are uniformly bounded, $I(a', \theta; G)e^{-a'} \rightarrow 0$ uniformly in θ . Thus the only possible discontinuities occur at atoms of G . These atoms are countable so the discontinuity set has Q_K -measure zero. Since $Q_{K_m} \rightarrow Q_K$ weakly, (J) gives $\tilde{\Gamma}_m \rightarrow \Gamma(G, K)$ so $\Gamma(G, K) = \lim_m \Gamma(G_m, K_m) = 0$. Kakutani–Fan–Glicksberg now delivers a fixed point $(G^*, K^*) \in BR_G(G^*, K^*) \times BR_K(G^*)$, and since $\arg\max_x U$ is closed, (i) is equivalent to the support inclusion. $\square$

I.4 Completing behavior at probability-zero histories. Fix (G^*, K^*) from Lemma 17. When evaluating firm 1's payoff, let $X_2 \sim G^*$ denote firm 2's concession date. Let $t_0 := \inf \text{supp} G^* \in [0, +\infty]$, so that $N_{G^*}(a) > 0$ for every $a > t_0$ and $N_{G^*}(a) = 0$ for every $a < t_0$. Let $Z \subseteq (0, \tau^R)$ denote the (Lebesgue-null, measurable) set of dates a with $N_{G^*}(a) > 0$ at which the optimality property of Lemma 17(ii) fails. By Lemma 15(iii), we can choose a Borel function $\theta^*(a) \in \Theta^*(a, G^*)$ at every date with $N_{G^*}(a) > 0$. Define the modified kernel

$$\hat{K}(a, \cdot) := \begin{cases} K^*(a, \cdot) & \text{if } N_{G^*}(a) > 0 \text{ and } a \notin Z, \\ \delta_{\theta^*(a)} & \text{if } N_{G^*}(a) > 0 \text{ and } a \in Z, \\ \delta_1 & \text{if } N_{G^*}(a) = 0. \end{cases}$$

At every feasible date in Z , the new response is optimal under the Bayes posterior $\nu_a^{G^*}$. When $N_{G^*}(a) = 0$, the detection history is impossible; there the response is $\chi_1(a) = \chi(a)$, which is uniquely optimal under the belief δ_a that the rival stopped at a (Lemma 13(iv)).

Lemma 18 (Responses at probability-zero dates). (i) $U(x; G^*, \hat{K}) = U(x; G^*, K^*)$ for every $x \in \overline{\mathbb{R}}_+$ with $x \geq t_0$;

(ii) $U(x; G^*, \hat{K}) \leq U(x; G^*, K^*)$ for every $x < t_0$.

Consequently

$$\max_x U(x; G^*, \hat{K}) = \max_x U(x; G^*, K^*) \quad \text{and} \quad \text{supp} G^* \subseteq \arg\max_x U(x; G^*, \hat{K}).$$

Proof. The kernels $\hat{K}$ and K^* differ only on the exceptional set $Z \cup \{a : N_{G^*}(a) = 0\}$, and $\{a : N_{G^*}(a) = 0\} \subseteq (0, t_0]$.

Part (i). Fix $x \geq t_0$. If $x = X_2$, both firms stop at x and no response kernel is used; this is the only branch when $t_0 = +\infty$. In the branch $x < X_2$ the rival's response draw is

taken at $a_2 = x + E_1$, whose conditional law is absolutely continuous and supported on $(x, \infty) \subseteq (t_0, \infty)$, on which $N_{G^*} > 0$; in the branch $X_2 < x$ firm 1's own response draw is taken at $a_1 = X_2 + E_2$, with $X_2 \geq t_0$ almost surely, so $a_1 > t_0$ almost surely. In both strict branches the detection date avoids the Lebesgue-null set Z and the point t_0 almost surely, so the realized responses under $\hat{K}$ and K^* coincide almost surely, and the payoffs are equal.

Part (ii). Fix $x < t_0$. Since $X_2 \geq t_0 > x$ almost surely, only the branch $x < X_2$ occurs: firm 1 concedes at $\tau_1 = x$, and firm 1's own responses are never triggered. The rival stops at $\tau_2 = X_2 \geq t_0 > x$ or at $\tau_2 = \chi_{\theta_2}(a_2) \geq a_2 > x$. The detection date $a_2 = x + E_1$ has an absolutely continuous distribution, so it belongs to the null set Z with probability zero. A modification on Z therefore has no effect on the expected payoff. At any remaining date where the kernels differ, $N_{G^*}(a_2) = 0$ and the response under $\hat{K}$ is $\chi_1(a_2) = \chi(a_2) \geq \chi_\theta(a_2)$ for every θ . Holding all other randomness fixed, the rival therefore never stops earlier under $\hat{K}$ than under K^* on this event. By [Lemma 12](#), a later rival stop lowers the payoff of a firm that has already stopped, so taking expectations gives $U(x; G^*, \hat{K}) \leq U(x; G^*, K^*)$. By (i), $U(\cdot; G^*, \hat{K})$ and $U(\cdot; G^*, K^*)$ agree on $[t_0, +\infty]$, a set that contains $\text{supp} G^*$ and on which the latter function attains its global maximum ([Lemma 17\(i\)](#), with $\text{supp} G^* \neq \emptyset$); by (ii) the modification does not raise the value anywhere below t_0 . Hence the maximum is unchanged and is still attained on $\text{supp} G^*$. $\square$

I.5 Dynamic consistency. The next lemma asks whether an initially optimal concession date remains optimal after the firm reaches date t without a signal. It does. The reason is that the initial payoff equals a term already determined before t plus a positive multiple of the continuation payoff at t . Thus the two problems rank every future concession date in the same way. Fix a rival reduced strategy (G, K) , let firm 1 respond per the same kernel K , and for $\sigma \in [t, +\infty]$ let $U(\sigma; G, K)$ be as above. Let N_t denote the event that firm 1 has received no signal by t (under any concession plan $\sigma \geq t$, this event does not depend on the plan), and let $\tilde{U}_t(\sigma)$ denote firm 1's expected continuation payoff, discounted to t , at the no-signal history at t , under the plan "concede at σ absent a signal, respond per K ."

Lemma 19 (Conditional optimality at no-signal histories). For every $t \geq 0$: $m(t) := \mathbb{P}(N_t) > 0$, and there is a constant A_t , independent of the plan, such that

$$U(\sigma; G, K) = A_t + \alpha(t) m(t) \tilde{U}_t(\sigma) \quad \text{for every } \sigma \in [t, +\infty].$$

Consequently $\arg\max_{\sigma \in [t, +\infty]} \tilde{U}_t(\sigma) = \arg\max_{\sigma \in [t, +\infty]} U(\sigma; G, K)$, and the same identity extends to plans that mix over $\sigma \geq t$, by linearity.

Proof. Since firm 1 has not stopped before t under any plan $\sigma \geq t$, firm 2 is unsignalled and stops, if at all, at its concession date $X_2 \sim G$; the first signal to firm 1 arrives at $X_2 + E_2$. Thus $N_t = \{X_2 > t\} \cup \{X_2 \leq t, X_2 + E_2 > t\}$, with

$$m(t) = G((t, +\infty]) + \int_{[0,t]} e^{-\omega(t-y)} G(dy) > 0,$$

the two terms not both null. Decompose $U(\sigma; G, K) = \mathbb{E}[U \mathbf{1}_{N_t^c}] + \mathbb{E}[U \mathbf{1}_{N_t}]$. On N_t^c the first signal arrives at some $a \leq t$ and firm 1 stops at $\chi_{\theta_1}(a)$ while firm 2 stopped at X_2 : the realized pair, hence the payoff, does not depend on σ . On N_t , firm 1 is active at t , so $\tau_1 \geq t$ and the frontier equals s for all $s \leq t$; splitting the payoff integral at t ,

$$\begin{aligned} U(\tau_1, \tau_2) \mathbf{1}_{N_t} &= \underbrace{\mathbf{1}_{N_t} \int_0^t \alpha(s) \pi(s, s \wedge X_2) ds}_{\text{independent of } \sigma} \\ &\quad + \mathbf{1}_{N_t} \alpha(t) \times (\text{continuation flows discounted to } t). \end{aligned}$$

Collecting the σ -independent pieces into A_t and conditioning the last term on N_t gives the display: the conditional law of $(X_2, E_1, E_2, \theta_1, \theta_2)$ given N_t is the distribution of the primitives at the no-signal history at t , so the conditional expectation of the continuation flows is $\tilde{U}_t(\sigma)$ by definition. $\square$

The final lemma treats no-signal histories after the latest date at which the rival might concede. Such histories can occur only after the firm has itself deviated from its plan.

Lemma 20 (After the latest possible concession date). Let $G \in \mathcal{G}$ with $T_G := \sup \text{supp} G < \infty$, and fix a history at which firm 1 is active at some $t > T_G$ against a rival playing (G, K) , and at which firm 1 has either received no signal, or received its first signal at a date $a > T_G$. The conclusions below also hold at $t = T_G$ whenever $G(\{T_G\}) = 0$. Under this no-atom condition, they also hold when the first signal arrives at T_G , because the rival stopped strictly before T_G with probability one. Then:

- (i) the rival has stopped, and the Bayes posterior over its (frozen) stopping date is

$$\hat{\mu}(dy) \propto e^{\omega y} G(dy) \quad \text{on } [0, T_G],$$

at every such history, independently of t and of the signal date, if any;

- (ii) in particular, signal arrivals after T_G carry no information: $\nu_a^G = \hat{\mu}$ for every $a > T_G$; the same equality holds at $a = T_G$ under the endpoint condition $G(\{T_G\}) = 0$;

(iii) the continuation problem is the deterministic stopping problem with objective $w(\cdot; t, \hat{\mu})$.

This objective is well defined because $\hat{\mu}$ is supported on $[0, T_G] \subseteq [0, t]$. Its maximum is attained on $\mathcal{T}(t) \subseteq [t, t \vee \tau^R]$. Let $y^*(t)$ be the earliest maximizing date. This rule is time-consistent: if the firm plans at t to stop at $y^*(t)$ and reaches an intermediate date $t' \leq y^*(t)$, it still chooses the same stopping date. Ignoring later signals and following this rule is optimal at every such history. When $t \geq \tau^R$, $y^*(t) = t$, so the firm stops immediately.

Proof. Parts (i)–(ii). Let $X_2 \sim G$ denote firm 2's concession date. Since $T_G < \infty$, we have $X_2 \leq T_G < t$ almost surely, so the rival has stopped. By Lemma 14(i), the no-signal posterior on $[0, t]$ is proportional to $e^{-\omega(t-y)}G(dy) \propto e^{\omega y}G(dy)$, and by Lemma 14(ii) the posterior upon a first signal at $a > T_G$ is $\nu_a^G \propto e^{\omega y} \mathbf{1}\{y \leq a\}G(dy) = e^{\omega y}G(dy)$: both equal $\hat{\mu}$ after normalization, for every $t, a > T_G$. (As recorded in Lemma 14(iii), firm 1's own past behavior does not enter this inference.)

Part (iii). Since the rival is frozen at $c \sim \hat{\mu}$ with $c \leq T_G < t$ and no future observation changes $\hat{\mu}$, the continuation payoff of any (possibly signal-contingent) plan is that of some distribution over stopping dates $y \geq t$ evaluated through $w(\cdot; t, \hat{\mu})$; it therefore suffices to maximize the latter. By Lemma 13(i),(iii), the maximum is attained on $\mathcal{T}(t) = [t, \chi(t)] \subseteq [t, t \vee \tau^R]$, and for $t \geq \tau^R$ we have $\mathcal{T}(t) = \{t\}$, so $y^*(t) = t$. Time-consistency is the splitting identity (2): $w(\cdot; t, \hat{\mu})$ and $w(\cdot; t', \hat{\mu})$ share their maximizers on $[t', +\infty]$, so if $y^*(t) \geq t'$ then $y^*(t') = y^*(t)$. Optimality at every such history follows since the prescription attains the maximum there.

Endpoint. If $t = T_G$ and $G(\{T_G\}) = 0$, the rival stopped strictly before t almost surely. The no-signal belief on $[0, T_G]$ is again $\hat{\mu}(dy) \propto e^{\omega y}G(dy)$, and the same formula holds after any signal at $a \geq T_G$. When $a = T_G$, every possible rival stopping date already lies below a , so the signal formula removes no possible date. The continuation objective is therefore $w(\cdot; t, \hat{\mu})$, and parts (i)–(iii) follow by the arguments just given. $\square$

I.6 Proof of Proposition 6.

Proof of Proposition 6. Fix (G^*, K^*) from Lemma 17, let $\hat{K}, t_0, Z$ be as in Appendix I.4, and write $T^* := \sup \text{supp } G^* \in [0, +\infty]$. Define the following symmetric profile for both firms:

- (a) *No-signal plan.* Draw a concession date X_i from G^* (Lemma 9). If the firm reaches date t without a signal and X_i is still in the future, it keeps that candidate. Conditional on reaching t without a signal while following this plan, its candidate has distribution G^* conditional on $[t, +\infty]$. If the firm deviated by continuing past X_i , it uses a fresh draw from this same conditional distribution whenever $G^*([t, +\infty]) >$

0. If that probability is zero, all possible rival concession dates have passed, and the firm stops at $y^*(t)$ from [Lemma 20\(iii\)](#). Keeping a surviving candidate, or drawing once and then keeping the new candidate, gives condition (DC).

- (b) *Response plan.* Upon a first signal at a , stop immediately if $a \geq \tau^R$. If $0 < a < \tau^R$, draw θ from $\hat{K}(a, \cdot)$ and stop at $\chi_\theta(a)$. At the impossible history $a = 0$, stop at $\chi(0)$. Let H_a denote the resulting distribution of response dates and let $\mu_a \in \Delta([0, a])$ be the belief held after the signal. The firm keeps its realized response date while that date remains in the future. If it deviates by continuing past that date to some s , it uses a fresh draw from H_a conditional on $[s, +\infty]$ whenever this event has positive probability. If no candidate in H_a lies at or after s , it stops at $\min \arg\max_{y \in [s, +\infty]} w(y; s, \mu_a)$. [Lemma 13\(iii\)](#) guarantees that this set of optimal dates is nonempty and compact. Equation (2) shows that a surviving candidate remains optimal and that the earliest optimal fallback remains unchanged as time passes. Thus the response plan also satisfies (DC).

All these rules vary measurably with the observed history. The conditional distributions are Borel, and the measurable maximum theorem used in [Lemma 15](#) implies that the earliest optimal dates above can also be chosen as Borel functions. [Lemma 9](#) then implements each required distribution with a uniform draw.

Beliefs. At every no-signal history at date t , use the posterior in [Lemma 14\(i\)](#), treating $X_j \leq t$ as a past stop and $X_j > t$ as an active rival. By [Lemma 14\(iii\)](#), the same posterior applies after histories created by the firm's own deviations.

At every detection date a with $N_{G^*}(a) > 0$, use the Bayes posterior $\nu_a^{G^*}$, including dates in the null set Z . The fixed-point condition already makes K^* optimal outside Z , while the measurable selection $\theta^*(a)$ makes $\hat{K}$ optimal on Z . If $N_{G^*}(a) = 0$, a signal at a is impossible under G^* , so Bayes' rule is undefined there. We then assign δ_a , under which the prescribed response $\chi(a)$ is optimal. At the impossible history $a = 0$, we likewise assign δ_0 and prescribe $\chi(0)$.

We now verify that no firm can gain after any history. As usual, it is enough to consider deviations that use no new random draw. By [Lemma 13\(iii\)](#), we need consider only response dates in $\mathcal{T}(\cdot)$, because every date outside that interval is strictly worse at the history where it is chosen.

Signal histories. At $a \geq \tau^R$: by [Lemma 13\(iii\)](#), $\mathcal{T}(a) = \{a\}$ and $w(\cdot; a, \nu)$ is strictly decreasing beyond a for every belief $\nu \in \Delta([0, a])$, so immediate stopping is optimal regardless of the belief held. At $a < \tau^R$ with $N_{G^*}(a) > 0$ and $a \notin Z$: the prescribed response mixes over $\arg\max_y w(y; a, \nu_a^{G^*})$ by [Lemma 17\(ii\)](#) and the definition of Z , hence is opti-

mal under the (Bayes) belief. At $a < \tau^R$ with $N_{G^*}(a) > 0$ and $a \in Z$, the prescribed response $\chi_{\theta^*(a)}(a)$ is optimal under the same Bayes belief by the definition of $\theta^*(a)$. If $N_{G^*}(a) = 0$, the assigned belief is δ_a and the prescribed response $\chi(a)$ is uniquely optimal (Lemma 13(iv)). Finally, at histories strictly between a signal at a and the prescribed stop y : no further signals arrive (the rival has already stopped) and survival is uninformative (Lemma 14(iii)), so the belief is unchanged. By (2), moving from a to an intermediate date s only adds a constant and multiplies later payoffs by a positive number. It does not change which remaining stopping date is best, so continuing to y stays optimal. At histories at some s strictly past the prescribed stop—reached only by the firm's own deviations—the belief is still unchanged: no signal has intervened (the rival is stopped), survival is uninformative, and the firm's own play is uninformative about its rival (Lemma 14(iii)). If H_a assigns positive probability to dates at or after s , draw a new response date from that part of H_a . Every date selected this way was optimal at a and remains optimal at s by (2). If H_a assigns no probability to such dates, choose the earliest date that maximizes $w(\cdot; s, \mu_a)$. This choice is optimal, and the same equation shows that both rules remain optimal as time passes.

No-signal histories with $G^*([t, +\infty]) > 0$. A continuation deviation consists of a (possibly mixed) concession date $\sigma \geq t$ on the no-signal history, together with responses at any future signal histories that may arise.

We treat responses first. Whenever a finite first signal arrives, its date satisfies $a_1 = X_2 + E_2 > t_0$ and its law is absolutely continuous; if $X_2 = +\infty$, no signal arrives (Lemma 14(ii)). Every such finite signal date satisfies $N_{G^*}(a_1) > 0$, so the posterior at a_1 is $\nu_{a_1}^{G^*}$ by Lemma 14(iii), irrespective of the deviator's own behavior. The prescribed response is optimal at every such date: K^* is optimal outside Z , and $\theta^*(a_1)$ is optimal on Z . Response deviations therefore cannot raise the continuation value.

It remains to compare the payoffs of the concession dates. By Lemma 19, applied with the rival's strategy $(G^*, \hat{K})$ and own responses $\hat{K}$, we have

$$\tilde{U}_t(\sigma) = \frac{U(\sigma; G^*, \hat{K}) - A_t}{\alpha(t) m(t)} \quad \text{for all } \sigma \in [t, +\infty],$$

so maximizing the continuation value over $\sigma \geq t$ is equivalent to maximizing $U(\cdot; G^*, \hat{K})$ over $[t, +\infty]$. The prescribed rule continues according to $G^*(\cdot \mid [t, +\infty])$, whose support is contained in

$$\text{supp } G^* \cap [t, +\infty] \subseteq \arg\max_{x \in \overline{\mathbb{R}}_+} U(x; G^*, \hat{K}) \cap [t, +\infty],$$

by **Lemma 18**. Every date selected by the prescription therefore attains the global maximum of $U(\cdot; G^*, \hat{K})$, and hence also the maximum restricted to dates no earlier than t , so the specified dates are optimal

No-signal histories with $G^([t, +\infty]) = 0$.* Here $T^* < \infty$ and either $t > T^*$, or $t = T^*$ with $G^*({T^*}) = 0$. By **Lemma 20**, the rival has already stopped, the belief is $\hat{\mu}$ both before and after any later signal, and stopping at $y^*(t)$ is optimal. At every later feasible detection date for which Bayes' rule is defined, the posterior equals $\hat{\mu}$. The response plan is optimal under this belief: the fixed-point kernel supplies an optimal response outside Z , and the measurable selection supplies one on Z . The time-consistency result in **Lemma 20**(iii) then completes the strategy. No deviation improves on it.

Histories at which the firm has already stopped involve no further decisions. Thus this profile, together with the specified beliefs, is a perfect Bayesian equilibrium of the game with news rate ω , and $\text{PBE}(\omega) \neq \emptyset$. $\square$

II EXTENSIONS

This appendix formalizes the two risk extensions discussed in the conclusion. We retain perfect transparency and common knowledge of rationality. The results bound every subgame-perfect equilibrium, including mixed equilibria. The statements and proofs are written for an arbitrary finite set of firms and specialize to the two-firm model of the main text.

II.1 Additive Risk . Keep the baseline technology law, but let risk depend on the number of firms that remain active. For a stopping profile τ , let

$$k_s(\tau) := \#\{j : \tau_j > s\}.$$

When $k_s(\tau) \geq 1$, every active firm is at the frontier s , and the date- s hazard is $\lambda(s, k_s(\tau))$; when $k_s(\tau) = 0$, the hazard is zero. Whether a firm that stops at s enters the count is immaterial because a single date has measure zero. Payoffs are otherwise as in the baseline model, using this hazard process.

Assumption 3 (additive risk regularity). For each $k \in \{1, \dots, n\}$, the hazard $\lambda(\cdot, k)$ is continuous and nondecreasing in technology. It is also nondecreasing in the number of active firms, so in particular

$$\lambda(t, k) \geq \lambda(t, 1) \quad \text{for every } t \geq 0 \text{ and } k \geq 1.$$

Flow payoffs satisfy **Assumption 1**(ii)–(iii). We will also impose the analog of **Assumption 1** (iv): for every firm i and fixed rival technology vector $\mathbf{c}$, the map

$$x \mapsto D_i \log \pi(x, \mathbf{c}) - \lambda(x, 1)$$

is continuous and strictly decreasing on $[\max_j c_j, +\infty)$. Finally, the solo-hazard analogue of **Assumption 1** (v) holds:

$$\int_0^\infty e^{-rs} \sup_{0 \leq x \leq s} \left\{ e^{-\int_0^x \lambda(u, 1) du} \pi(x, 0, \dots, 0) \right\} ds < +\infty.$$

Say we are in the *additive risk environment* when the above assumptions hold and there is perfect transparency and common knowledge of rationality.

Proposition 10 (The solo hazard bounds a additive risk). Suppose that we are in the additive risk environment. Then every subgame-perfect equilibrium satisfies:

$$\max_{i \in \mathcal{N}} \tau_i \leq \inf \{t \geq 0 : D_i \log \pi(t, \dots, t) \leq \lambda(t, 1)\} \quad \text{almost surely.}$$

Proof. We first solve the last-active firm's problem, then induct on the number of active firms. Let τ^C denote the infimum in the statement. If $\tau^C = +\infty$, the claim is immediate. Suppose $\tau^C < +\infty$. Continuity implies

$$D_i \log \pi(\tau^C, \dots, \tau^C) \leq \lambda(\tau^C, 1).$$

Fix a public history at date t , and suppose that firm i scales alone while its rivals remain frozen at $\mathbf{c} \leq (t, \dots, t)$. Conditional on survival to t , its payoff from stopping at $\sigma \geq t$ is

$$R_{t, \mathbf{c}}^C(\sigma) := \int_t^\sigma e^{-\int_t^s [r + \lambda(u, 1)] du} \pi(s, \mathbf{c}) ds + e^{-\int_t^\sigma [r + \lambda(u, 1)] du} \frac{\pi(\sigma, \mathbf{c})}{r},$$

with the terminal term absent when $\sigma = +\infty$. At every finite $\sigma > t$,

$$\frac{\partial}{\partial \sigma} R_{t, \mathbf{c}}^C(\sigma) = e^{-\int_t^\sigma [r + \lambda(u, 1)] du} \frac{\pi(\sigma, \mathbf{c})}{r} [D_i \log \pi(\sigma, \mathbf{c}) - \lambda(\sigma, 1)].$$

At $t = \tau^C$, log-supermodularity and **Assumption 3** give, for every $x \geq \tau^C$,

$$\begin{aligned} D_i \log \pi(x, \mathbf{c}) - \lambda(x, 1) &\leq D_i \log \pi(x, \tau^C, \dots, \tau^C) - \lambda(x, 1) \\ &\leq D_i \log \pi(\tau^C, \dots, \tau^C) - \lambda(\tau^C, 1) \leq 0, \end{aligned}$$

where the second inequality is strict when $x > \tau^C$. Hence

$$R_{\tau^C, c}^C(\sigma) < R_{\tau^C, c}^C(\tau^C) = \frac{\pi(\tau^C, c)}{r} \quad \text{for every } \sigma > \tau^C,$$

including $\sigma = +\infty$, because the integrability condition makes the terminal term vanish and the objective continuous there. Thus a last active firm uniquely prefers to stop at τ^C .

The additive risk payoff is continuous in the stopping profile. Indeed, along any convergent sequence of stopping profiles, the active-firm count converges at every date other than the finitely many limiting stopping dates, so cumulative hazards converge on every bounded interval. The solo-hazard integrability condition supplies the same dominating envelope as in [Lemma 7](#). Dominated convergence therefore applies. In particular, stopping immediately after an observed stop is payoff-equivalent to stopping simultaneously.

We now claim that no equilibrium continuation from a public history at τ^C carries the frontier above τ^C . We prove the claim by induction on the number of active firms. The preceding calculation proves the case of one active firm. Suppose the claim holds with fewer than m active firms, and consider a history with m active firms. If an equilibrium continuation carries the frontier above τ^C with positive probability, then, because m is finite, some active firm i satisfies

$$\mathbb{P}(\tau_i > \tau^C) > 0.$$

Let c be its rivals' technologies at this history. If firm i stops at τ^C , the successor subgame has $m - 1$ active firms. Subgame perfection and the induction hypothesis make every remaining active firm stop immediately, so the deviation gives firm i the payoff $R_{\tau^C, c}^C(\tau^C)$.

For comparison, fix any realized equilibrium continuation and write σ for firm i 's stopping date. Rival technologies are weakly above c , which weakly lowers firm i 's flow payoff. While firm i remains active, the actual hazard is at least $\lambda(s, 1)$; after it stops, any remaining hazard can only lower its payoff further. Hence, path by path, its continuation payoff is no greater than $R_{\tau^C, c}^C(\sigma)$, which is strictly below $R_{\tau^C, c}^C(\tau^C)$ whenever $\sigma > \tau^C$. The comparison holds before averaging over equilibrium randomization. Since firm i delays with positive probability, stopping at τ^C strictly raises its expected payoff, contradicting sequential rationality. This proves the induction claim.

If an equilibrium frontier exceeded τ^C with positive probability, its public continuation at the finite date τ^C , conditional on survival, would contradict the claim. Therefore the frontier is at most τ^C almost surely. $\square$

II.2 Accumulative risk. Keep the baseline hazard λ , but suppose that risk persists at the peak technology. For a stopping profile τ , let

$$T(\tau) := \max_{j \in \mathcal{N}} \tau_j.$$

The date- s hazard is now $\lambda(s \wedge T(\tau))$. Thus firm i 's payoff is

$$\int_0^\infty \exp \left\{ -rs - \int_0^s \lambda(u \wedge T(\tau)) du \right\} \pi(s \wedge \tau_i, (s \wedge \tau_j)_{j \neq i}) ds.$$

If $T(\tau) < +\infty$, the hazard remains equal to $\lambda(T(\tau))$ after the last stop. Its survival factor is no larger than the baseline survival factor, so **Assumption 1(v)** keeps payoffs finite.

The marginal condition below need not be single-crossing under **Assumption 1**. We therefore impose the analogue of **1 (iv)**:

Assumption 4 (Persistent-risk single-crossing). In addition to **Assumption 1**, for every firm i and fixed rival technology vector c , the map

$$x \mapsto D_i \log \pi(x, c) - \frac{\lambda'(x)}{r + \lambda(x)}$$

is strictly decreasing on $[\max_j c_j, +\infty)$.

Say we are in the *accumulative risk environment* when the above assumptions hold and there is perfect transparency and common knowledge of rationality.

Proposition 11 (Marginal persistent risk bounds the frontier). Suppose we are in the accumulative risk environment. Then every subgame-perfect equilibrium satisfies

$$\max_{i \in \mathcal{N}} \tau_i \leq \inf \left\{ t \geq 0 : D_i \log \pi(t, \dots, t) \leq \frac{\lambda'(t)}{r + \lambda(t)} \right\} \quad \text{almost surely.}$$

Proof. Again, we solve the last-active firm's problem, then induct. Let τ^A denote the infimum in the statement. If $\tau^A = +\infty$, the claim is immediate. Suppose $\tau^A < +\infty$. Continuity implies that the defining inequality holds at τ^A .

Fix a history at date t , and suppose that firm i scales alone while its rivals remain frozen at $c \leq (t, \dots, t)$. Conditional on survival to t , its payoff from stopping at $\sigma \geq t$ is

$$R_{t,c}^A(\sigma) := \int_t^\sigma e^{-\int_t^s [r + \lambda(z)] dz} \pi(s, c) ds + e^{-\int_t^\sigma [r + \lambda(z)] dz} \frac{\pi(\sigma, c)}{r + \lambda(\sigma)},$$

with the terminal term absent when $\sigma = +\infty$. At every finite $\sigma > t$,

$$\frac{\partial}{\partial \sigma} R_{t,c}^A(\sigma) = e^{-\int_t^\sigma [r+\lambda(z)] dz} \frac{\pi(\sigma, \mathbf{c})}{r + \lambda(\sigma)} \left[D_i \log \pi(\sigma, \mathbf{c}) - \frac{\lambda'(\sigma)}{r + \lambda(\sigma)} \right].$$

At $t = \tau^A$, log-supermodularity and **Assumption 4** give, for every $x \geq \tau^A$,

$$\begin{aligned} D_i \log \pi(x, \mathbf{c}) - \frac{\lambda'(x)}{r + \lambda(x)} &\leq D_i \log \pi(x, \tau^A, \dots, \tau^A) - \frac{\lambda'(x)}{r + \lambda(x)} \\ &\leq D_i \log \pi(\tau^A, \dots, \tau^A) - \frac{\lambda'(\tau^A)}{r + \lambda(\tau^A)} \leq 0, \end{aligned}$$

where the second inequality is strict when $x > \tau^A$. Therefore

$$R_{\tau^A, \mathbf{c}}^A(\sigma) < R_{\tau^A, \mathbf{c}}^A(\tau^A) = \frac{\pi(\tau^A, \mathbf{c})}{r + \lambda(\tau^A)} \quad \text{for every } \sigma > \tau^A,$$

including $\sigma = +\infty$ by **Assumption 1(v)**. Thus a last active firm uniquely prefers to stop at τ^A .

The accumulative-risk payoff is continuous in the stopping profile. Along any convergent sequence, the peak path $s \mapsto s \wedge T(\boldsymbol{\tau})$ and its cumulative hazard converge on every bounded interval. The persistent survival factor is bounded above by the baseline one, so **Assumption 1(v)** supplies an integrable envelope. Dominated convergence applies, and immediate successor stopping is payoff-equivalent to simultaneous stopping.

We claim that no equilibrium continuation from a public history at τ^A carries the frontier above τ^A . Induct on the number of active firms. The last-active calculation proves the base case. Suppose the claim holds with fewer than m active firms, and consider a history with m active firms. If the equilibrium continuation carries the frontier beyond τ^A with positive probability, some active firm i satisfies $\mathbb{P}(\tau_i > \tau^A) > 0$. Let $\mathbf{c}$ be its rivals' current technologies. If firm i stops at τ^A , the successor subgame has $m - 1$ active firms; by subgame perfection and the induction hypothesis, every remaining active firm stops immediately. The deviation therefore gives firm i the payoff $R_{\tau^A, \mathbf{c}}^A(\tau^A)$.

Fix any realized equilibrium continuation and write σ for firm i 's stopping date. Rival technologies are weakly above $\mathbf{c}$, which weakly lowers firm i 's flow payoff. Moreover, the realized peak is always weakly above the path $s \mapsto s \wedge \sigma$ generated when firm i scales alone, so the actual persistent hazard is weakly higher than the hazard in $R_{\tau^A, \mathbf{c}}^A(\sigma)$. Firm i 's realized continuation payoff is thus no greater than this frozen-rival benchmark, which is strictly below $R_{\tau^A, \mathbf{c}}^A(\tau^A)$ whenever $\sigma > \tau^A$, including $\sigma = +\infty$. This comparison is pathwise and therefore covers mixed strategies. Because firm i delays with positive

probability, immediate stopping is a profitable deviation, a contradiction. This proves the induction claim.

If an equilibrium frontier exceeded τ^A with positive probability, its public continuation at the finite date τ^A , conditional on survival, would contradict the claim. Hence the frontier is at most τ^A almost surely. $\square$

III FURTHER RESULTS

III.1 Tight monitoring thresholds under stationarity. We now formalize the claim in the main text that the threshold $\bar{\omega}$ is exact in the stationary environment we saw in [Example 4](#) and [Section 6](#). To recap, flow payoffs and the hazard are

$$\pi(t_i, t_j) = e^{\alpha t_i - \beta t_j}, \quad \lambda(t) \equiv \lambda,$$

where $\alpha, \beta > 0$, $\lambda > \alpha$, and $r > \alpha - \beta$. We maintain common knowledge of rationality. The inequality $\lambda > \alpha$ places the game in the coordination phase: once a firm knows that its rival has stopped, it strictly prefers to stop immediately.

We call this environment *stationary* because a common translation of both technology paths multiplies every continuation payoff by the current diagonal flow $\pi(t, t) = e^{(\alpha - \beta)t}$. Values per unit of current flow therefore do not depend on the date or technology level. At a symmetric quiet history, the normalized value of racing forever is

$$V_R = \frac{1}{r + \lambda - \alpha + \beta}.$$

The normalized value of stopping first is

$$V_S(\omega) = V_S(0) + \frac{\omega [1/r - V_S(0)]}{\omega + r + \lambda + \beta}, \quad V_S(0) = \frac{1}{r + \lambda + \beta}.$$

Thus

$$\Delta(t, \omega) = \pi(t, t) [V_S(\omega) - V_R],$$

and the bracketed difference is positive exactly when

$$\omega > \bar{\omega} := \frac{r\alpha}{\lambda - \alpha + \beta}.$$

Proposition 12 (Exact racing-to-ruin threshold in the stationary environment). In the stationary environment:

- (i) If $\omega \leq \bar{\omega}$, there is an equilibrium in which both firms race forever, so ruin occurs with probability one.
- (ii) If $\omega > \bar{\omega}$, every $\tau \in \text{PBE}(\omega)$ satisfies $\max_i \tau_i < +\infty$ almost surely.

Proof. We construct a ruin equilibrium at and below the threshold, then rule out ruin above it.

Part (i). Consider the assessment used in [Proposition 7](#): each firm races forever at every quiet history and stops immediately after detecting its rival's stop. Against a rival that races until detection, any finite stop at σ yields, by [Lemma 6](#), the payoff from racing forever plus a strictly positive multiple of $\Delta(\sigma, \omega)$. Stationarity and $\omega \leq \bar{\omega}$ give $\Delta(\sigma, \omega) \leq 0$ at every date, so no finite stop is profitable. Immediate stopping after a signal is uniquely optimal because $\lambda > \alpha$. The assessment is therefore a perfect Bayesian equilibrium, including at $\omega = \bar{\omega}$, where the stopping deviation is payoff-equivalent to racing forever.

Part (ii). Fix an equilibrium with outcome (τ_1, τ_2) and let $E := \{\tau_1 = \tau_2 = +\infty\}$.

Step 1: reduce an infinite frontier to joint racing. Because $\lambda > \alpha$, a sole survivor's payoff against any frozen rival is strictly decreasing in its stopping date. Immediate stopping is therefore uniquely optimal at every signal history under every admissible belief. If one firm stops in finite time, its rival detects the stop in finite time almost surely and then stops immediately. Hence $\{\max_i \tau_i = +\infty\} = E$ up to a null event.

Step 2: show that joint racing has zero probability. Suppose instead that $\mathbb{P}(E) > 0$. Fix firm i and its quiet history H_t at date t , meaning that it is active and has received no signal. Since $E \subseteq H_t$, we have $\mathbb{P}(H_t) \geq \mathbb{P}(E) > 0$. Conditional on H_t , its rival is in one of three states: it has stopped but remains undetected; it is active and will stop at a finite date; or it is active and never stops. The conditional probability of the second state is at most

$$\frac{\mathbb{P}(t < \tau_j < +\infty)}{\mathbb{P}(E)} \rightarrow 0.$$

The joint conditional probability that both firms' residual plans never stop is at least $\mathbb{P}(E \mid H_t) \geq \mathbb{P}(E)$.

On the last two states, the rival's technology remains weakly above t . Thus, at every future date $s \geq t$, the discounted, survival-weighted flow from firm i 's never-stopping plan is bounded by

$$\pi(t, t)e^{-(r+\lambda-\alpha)(s-t)}.$$

Its continuation payoff is consequently at most $\pi(t, t)/(r + \lambda - \alpha)$.

Now modify firm i 's plan at H_t only on the device realizations for which it would

never stop: stop at t on those realizations and leave every other plan unchanged. If the rival has already stopped, the change is a weak gain. If both residual plans never stop, it replaces $V_R\pi(t, t)$ with $V_S(\omega)\pi(t, t)$, a gain of $[V_S(\omega) - V_R]\pi(t, t) > 0$. This gain receives conditional weight at least $\mathbb{P}(E)$. The only possible loss comes when the rival is active but will stop at a finite date; its conditional probability tends to zero, and its magnitude is bounded by $\pi(t, t)/(r + \lambda - \alpha)$. The conditional expected gain is therefore positive for all sufficiently large t , contradicting sequential rationality at H_t . Hence $\mathbb{P}(E) = 0$. Step 1 now gives $\max_i \tau_i < +\infty$ almost surely. $\square$

Trust for the case $\omega < \beta$ We now consider the other case [Section 6](#) excluded there: news that is slow relative to payoff erosion: $\omega < \beta$. The setup is otherwise identical to that of [Section 6](#). Define

$$\underline{\ell}_\beta := \frac{V_R - V_S(0)}{1/r - V_W(\beta)}, \quad \text{noting} \quad \frac{1}{r} - V_W(\beta) = \frac{\lambda - \alpha}{r(r + \lambda - \alpha + \beta)} = \frac{(\lambda - \alpha)V_R}{r} :$$

this is the verification threshold $\underline{\ell}$ of [Section 6](#) evaluated at news rate β , and it does not depend on ω .

Proposition 13 (Trust with slow news). In the stationary coordination environment with types as above, suppose $0 < \omega < \beta$ and maintain $\underline{\ell}_\beta < \bar{\ell}$. Then [Proposition 9](#) holds with $\underline{\ell}$ replaced by $\underline{\ell}_\beta$:

- (i) *Low trust.* If $\ell < \underline{\ell}_\beta$, then every $\tau \in \text{PBE}(\omega, \ell)$ has $\tau_1 = \tau_2 = +\infty$ almost surely.
- (ii) *Medium trust.* If $\underline{\ell}_\beta \leq \ell \leq \bar{\ell}$, then $\text{PBE}(\omega, \ell)$ contains both the outcome in which rational firms stop immediately and the never-stopping outcome.
- (iii) *High trust.* If $\ell > \bar{\ell}$, then every $\tau \in \text{PBE}(\omega, \ell)$ satisfies $\mathbb{P}(\tau_1 = \tau_2 = +\infty \mid \text{both firms rational}) \leq (\bar{\ell}/\ell)^2$, and stopping at zero remains an equilibrium outcome.

Moreover, rational firms stopping at a common date is an equilibrium if and only if $\ell \geq \underline{\ell}_\beta$, and the binding deviation is a marginal delay rather than indefinite verification: for $\ell < \underline{\ell}_\beta$, the most profitable plan against the bunched profile is to delay by $\sigma^* = \log(\underline{\ell}_\beta/\ell)/(\beta - \omega)$.

Proof. As in the proof of [Proposition 9](#), abbreviate $\delta := r + \lambda - \alpha + \beta = 1/V_R$ and $\rho := r + \lambda + \beta = 1/V_S(0)$, and write $N := V_R - V_S(0)$; the displayed identity for $1/r - V_W(\beta)$ is a direct computation. That proof used $\omega \geq \beta$ in two places: in part (i), to

damp the undetected-stop cell, and in part (ii)(a), to conclude that the deviation payoff W is maximized at an endpoint of $[0, \infty]$. Both steps are replaced by an analysis of marginal delay.

The marginal-delay deviation. Fix a quiet history of firm i at date t and, for $h > 0$, let $\text{DELAY}(h)$ denote the plan: continue until $t + h$, stopping immediately upon a signal, and stop at $t + h$ if none has arrived. Per unit of current flow, we compute the derivative at $h = 0$ of its payoff, net of stopping at once, against each rival state in the posterior decomposition of part (i) of the proof of [Proposition 9](#).

Crazy. $\text{DELAY}(h)$ yields $(1 - e^{-\delta h})V_R + e^{-\delta h}V_S(0)$: postponing an unreciprocated stop gains at rate $\delta N > 0$.

Rational, stopped at $c \leq t$, undetected. The delaying firm races beside a frozen rival under the live hazard, and $\alpha < \lambda$ makes stopping at once optimal on this cell: the deviation loses at rate $e^{\beta(t-c)}(\lambda - \alpha)/r = e^{\beta(t-c)}\delta[1/r - V_W(\beta)]$.

Rational, active, with residual quiet-path stopping time τ . Stopping now ends the risk at $X \wedge \tau$, where $X \sim \text{Exp}(\omega)$ is the detection lag, and is worth $V_S(0) + [1/r - V_S(0)]\mathbb{E}[e^{-\rho(X \wedge \tau)}]$; on $\{\tau > h\}$, delaying replaces this by $(1 - e^{-\delta h})V_R + e^{-\delta h}\{V_S(0) + [1/r - V_S(0)]\mathbb{E}[e^{-\rho(X \wedge (\tau-h))}]\}$. Differentiating at $h = 0$ and using $\mathbb{E}[e^{-\rho(X \wedge \tau)}] = \frac{\omega}{\omega + \rho} + \frac{\rho}{\omega + \rho}e^{-(\omega + \rho)\tau}$, the deviation gains at rate $-\delta\psi(\tau)$, where

$$\psi(\tau) := \left[\frac{1}{r} - V_S(0)\right] \left[\frac{\omega}{\omega + \rho} + \rho\left(\frac{1}{\omega + \rho} - \frac{1}{\delta}\right)e^{-(\omega + \rho)\tau}\right] - N \quad \text{for } \tau \in (0, \infty],$$

and $\psi(0) := 1/r - V_W(\beta)$ at an atom at zero, where the rival stops at t itself and the previous cell applies with $c = t$. (A residual stop at $t + u$ with $u \in (0, h]$ contributes $\alpha u/r - (\lambda - \alpha)(h - u)/r + o(h) \geq -\delta\psi(0)h + o(h)$, so such mass is covered by the bound below.) Since $\omega + \rho - \delta = \omega + \alpha > 0$, the coefficient of $e^{-(\omega + \rho)\tau}$ is negative: ψ is strictly increasing on $(0, \infty)$, from $\psi(0^+) = -N - \alpha[1/r - V_S(0)]/\delta < 0$ to $\psi(\infty) = V_S(\omega) - V_R$. Hence, using the maintained $\underline{\ell}_\beta < \bar{\ell}$,

$$\sup_{[0, \infty]} \psi = \max \left\{ \frac{1}{r} - V_W(\beta), V_S(\omega) - V_R \right\} = \frac{N}{\min\{\underline{\ell}_\beta, \bar{\ell}\}} = \frac{N}{\underline{\ell}_\beta}.$$

A marginal delay is thus costly only against a rival that has just stopped—the slight-delay exposure that defines $\underline{\ell}_\beta$ —or against one that will never stop spontaneously, forgoing the

prize $V_S(\omega) - V_R$ of stopping first.

Part (i). Suppose, for contradiction, that some rational firm stops at a quiet history with positive probability. With F_j and S_j as in part (i) of the proof of [Proposition 9](#), let $t_0 := \inf\{t \geq 0 : F_1(t) + F_2(t) > 0\} < \infty$, and fix $\varepsilon > 0$ with $e^{(\beta-\omega)\varepsilon} < \underline{\ell}_\beta/\ell$. At a quiet history of firm i at any date $t \leq t_0 + \varepsilon$, every undetected stop has age $t - c \leq \varepsilon$. Weighting the three cells in proportion to 1, $\ell S_j(t)$, and $\ell e^{-\omega(t-c)} dF_j(c)$, the net gain rate of DELAY is proportional to

$$\begin{aligned} & \delta \left\{ N - \ell \left(S_j(t) \mathbb{E}[\psi(\tau)] + \left[\frac{1}{r} - V_W(\beta) \right] \int_{[t_0, t]} e^{(\beta-\omega)(t-c)} dF_j(c) \right) \right\} \\ & \geq \delta N \left[1 - \frac{\ell}{\underline{\ell}_\beta} e^{(\beta-\omega)\varepsilon} \right] \\ & > 0, \end{aligned}$$

by $\mathbb{E}[\psi(\tau)] \leq N/\underline{\ell}_\beta$, $e^{(\beta-\omega)(t-c)} \leq e^{(\beta-\omega)\varepsilon}$, and $S_j(t) + \int_{[t_0, t]} dF_j \leq 1$. A small delay therefore strictly improves on stopping at every quiet history at dates $t \leq t_0 + \varepsilon$, so no rational firm stops there, under any pure or mixed rule—contradicting the definition of t_0 . Hence no rational firm ever stops at a quiet history; since the first stop in any equilibrium must occur at a quiet history, no stop ever occurs and $\tau_1 = \tau_2 = +\infty$ almost surely.

Part (ii). The never-stopping profile is an equilibrium for $\ell \leq \bar{\ell}$ by part (ii)(b) of the proof of [Proposition 9](#), which nowhere compares ω with β . For the bunched profile, the deviation payoff $W(\sigma)$ is as computed in part (ii)(a) there, and

$$e^{\delta\sigma} W'(\sigma) = -p \frac{\lambda - \alpha}{r} e^{(\beta-\omega)\sigma} + (1-p) \delta N.$$

For $\omega < \beta$ the right side is strictly decreasing in σ , so W' changes sign at most once, from + to −, and W is maximized at $\sigma = 0$ if and only if $W'(0) \leq 0$, i.e.

$$p \frac{\lambda - \alpha}{r} \geq (1-p) \delta N \iff \ell \geq \frac{N}{(\lambda - \alpha)V_R/r} = \underline{\ell}_\beta.$$

Hence for $\ell \geq \underline{\ell}_\beta$ no plan improves on stopping at once and the bunched profile is an equilibrium; stationarity extends this to any common date. Conversely, for $\ell < \underline{\ell}_\beta$ the unique maximizer of W is the interior delay $\sigma^* = \log(\underline{\ell}_\beta/\ell)/(\beta - \omega) > 0$, a strictly profitable deviation—so no equilibrium has rational firms stopping at a common date. This proves part (ii) and the final claim.

Part (iii). The exclusion of never-stopping for $\ell > \bar{\ell}$ uses only the immediate-stop gamble, and zero ruin is attained by the bunched profile, an equilibrium since $\ell > \bar{\ell} > \underline{\ell}_\beta$; both are as in **Proposition 9**. For the $(\bar{\ell}/\ell)^2$ bound, the proof there lets a delaying plan gain at most $1/r$ per unit of the vanishing undetected-stop weight; when $\omega < \beta$ this bound fails, since payoffs on that cell grow like $e^{\beta(t-c)}$ while its weight decays only like $e^{-\omega(t-c)}$. But no bound is needed: because $\alpha < \lambda$, stopping at once attains the cell-maximum $e^{\beta(t-c)}/r$ on the undetected cell, so, relative to stopping at once, a continuation plan gains nothing there and can profit only on the crazy and active cells, where the argument of **Proposition 9** applies verbatim. The conclusion follows unchanged. $\square$

Together with **Proposition 9**, the simultaneous-coordination threshold at every news rate is $[V_R - V_S(0)] / [1/r - V_W(\omega \vee \beta)]$: constant in ω below β —the vertical segment in **Figure 8**—and rising thereafter. The two regimes fit together continuously: as $\omega \uparrow \beta$ the binding delay σ^* lengthens without bound, deforming into the indefinite-verification deviation that binds when $\omega \geq \beta$. Transparency below β therefore neither helps nor hurts simultaneous coordination—the binding temptation is instantaneous, and its cost and benefit are both ω -free—while faster news still lowers the sequential threshold $\bar{\ell}$. If instead $\underline{\ell}_\beta > \bar{\ell}$, the remarks following **Proposition 9** apply with $\underline{\ell}_\beta$ in place of $\underline{\ell}$.