COOPERATING AGAINST CATASTROPHE *

Drew Fudenberg
MIT

Andrew Koh
Columbia

October 21, 2026

Abstract

We study continuous-time games where competing firms choose R&D speeds while a safety threshold advances. There is a common disaster hazard whenever capabilities exceed the threshold. At low risk, racing is a Markov perfect equilibrium (MPE) while permanent pacing is not; at intermediate risk, both coexist; at high risk, pacing is an MPE but racing is not. At low risk, high discounting forces perpetual racing from safety ties in all subgame-perfect equilibria (SPE) and we bound disaster risk across all SPE. Delayed observation of capabilities weakens pacing incentives. Endogenous safety research can lower the effective hazard needed to sustain pacing.

JEL codes: C73, D43, O32

Keywords: AI competition, R&D races, self-enforcing restraint, pacing the frontier, dynamic public goods, observation lags

*Fudenberg: MIT Department of Economics drew.fudenberg@gmail.com; Koh: Columbia University Department of Economics andrew.koh@columbia.edu. We thank Allan Dafoe, Daniel Kokotajlo, Harry Pei, David Rand, and Phil Trammell for helpful conversations and National Science Foundation grant SES-2417162 for financial support.

1 INTRODUCTION

In September 2026, Dario Amodei, CEO of Anthropic, advocated for *pacing the frontier*: slowing the development of AI capabilities to give AI safety research and institutions time to catch up (Amodei, 2026). Other leaders of American AI firms, including Sam Altman, Demis Hassabis, and Elon Musk, soon voiced their support. A few months earlier, more than 1,300 employees of frontier AI companies signed an open letter urging the US government to support an international effort to develop the tools needed to pace frontier AI development (Pacing the Frontier Signatories, 2026). At the same time, critics have warned that government intervention could primarily protect incumbent firms and their rents (Doctorow, 2026; Horowitz, 2026). One obstacle to pacing the frontier is that each firm may reason that if it slows down, a rival will press ahead.[1] Yet strategic rivalry can also produce mutual restraint when the parties share an interest in avoiding a bad outcome (Schelling, 1960). When can self-interested, profit-maximizing firms slow AI research to a safe pace without government intervention?

To answer this question, we analyze a continuous-time game where two firms each choose a development speed at every date, from zero (pausing) to one (racing at maximum speed).[2] Each firm's technological capability (e.g., that of its AI model) is the integral of past speeds; the capability frontier is the higher of the two firms' capabilities. An exogenous safety threshold (e.g., the state of AI alignment research, or institutional preparedness) advances at a constant, exogenously fixed rate strictly between zero and one.

When the frontier exceeds safety, firms are exposed to a positive risk of disaster that permanently terminates both firms' profits.[3] The hazard governing disaster risk is a nondecreasing and concave function of the distance above safety, and might include a discrete jump at the boundary. Before disaster, each firm's flow profit increases with its own capability and decreases with its rival's. Each firm's expected payoff is the present discounted value of flow profits, weighted by the probability of survival up to each date. There is no direct cost of development.

Most of our paper focuses on the Markov Perfect Equilibria (MPE) of this game,[4] and

[1] Versions of this argument have been advanced by, among others, Andreessen (2023); Aschenbrenner (2024); Cowen (2026).

[2] We interpret players as firms that maximize profits, but we think our results might also apply to US-China AI competition.

[3] We assume that this is the only effect of disaster on the firms. If they internalize any social losses, it would be easier to support pacing the frontier and reduce the set of parameters where both firms always race.

[4] Here the state is the safety threshold and each firm's capability.

their relation to two simple strategy profiles that capture the possibilities of unrestrained development and endogenous restraint. The *racing profile* prescribes undertaking R&D at the maximum speed at every state. The *pacing profile* prescribes that a firm that is strictly behind ('follower') always races, and the firm ahead ('leader') follows an optimal stopping rule: race to an optimally chosen distance above safety, and then pause until either safety or its rival catches up. Tied firms pace if they are exactly at the threshold, pause if they are above it but not too far, and race otherwise.

Whether either or both of these two profiles is an MPE depends only on the "effective hazard rate," which is the constant disaster rate that, if it applied throughout joint racing from a safety tie, would give the same payoff as the model's actual hazard profile. At low effective hazard, racing is an MPE, but permanent pacing from a safety tie is not. At intermediate hazard, racing and pacing are both MPEs. At high hazard, pacing is an MPE but global racing is not. The lower cutoff is where a firm is indifferent between racing and safe pacing from a tie; the upper cutoff is where racing ceases to be a best response to a racing rival. A steeper hazard profile can reduce equilibrium disaster probability if it induces the leader to accommodate rather than race, so the behavioral response can outweigh the direct increase in exposure.

We also study subgame-perfect equilibrium (SPE), which allows strategies to depend on the full history of play. When effective hazard is below the lower cutoff and the discount rate is sufficiently high, every SPE outcome from a safety tie is perpetual joint racing. A recursive payoff guarantee bounds disaster risk across SPE. With constant hazard, history-dependent strategies can support turn-taking that gives both firms strictly more than the pacing MPE. We give an exact condition for permanent pacing to be an SPE outcome under more general payoff functions, and show that a grim trigger strategy implements pacing whenever it is sustainable. With our baseline payoffs, and when the hazard rate never exceeds the upper cutoff, this condition is the same as the lower cutoff for pacing in Markov strategies, so history dependence does not enlarge the set of parameters where permanent pacing can be sustained.

Next, to reflect the gap between AI firms' deployed (public) and internal (often private) capabilities (Model Evaluation and Threat Research, 2026), we suppose that internal capabilities become observable and generate profits and risk only after a deployment lag. We show how this lag shapes the sustainability of pacing: a longer lag allows a firm to build a larger lead before its rival can respond, making a deviation from pacing more attractive.

Our final extension allows firms to allocate research capacity between capability devel-

opment and publicly shared safety research. This creates an incentive to free ride on a rival's safety effort, but a firm that pulls ahead induces its rival to race to catch up and pause contributions to safety. When disaster risk grows linearly in distance above safety, we show that more productive safety research lowers the effective-hazard cutoff above which global racing is not an equilibrium: slowing development becomes more attractive because it also accelerates safety progress. Under conditions ensuring that followers prefer to race, more productive safety research also lowers the effective hazard needed to sustain pacing. Firms can then develop faster while remaining on the safety threshold, raising the payoff from cooperation.

Related work. *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. In our model, unlike any of these, investment generates a common disaster hazard. Fudenberg and Tirole (1983) constructs equilibria in which firms stop investing at lower levels than would be profit-maximizing if their opponent's capital level were fixed because further investment would induce their rival to invest as well. Our pacing equilibrium similarly sustains restraint through fear of provoking the rival to invest more, but here firms slow development to limit a common disaster hazard while safety catches up, rather than stopping at fixed capital stocks. Leininger (1989) studies restraint sustained by threats of further escalation in a dollar auction; our model instead links restraint to a common disaster hazard.

AI races. Dafoe (2018), Askell et al. (2019),Bostrom et al. (2020), and Zwetsloot and Dafoe (2019) analyze AI competition without developing formal models. Armstrong et al. (2016) and Bueno de Mesquita et al. (2026) analyze how incentives to finish first can induce developers to sacrifice safety in static games where firms make once-and-for-all decisions, and Emery-Xu et al. (2024) studies how incomplete information about heterogeneous capabilities affects investment in safety. Hendrycks et al. (2025) and AI Futures Project (2026) outline how US-China AI cooperation might be sustained via the threat of sabotaging or destroying a rival's data centers. Stafford et al. (2022) studies a dynamic race to implement a risky technology, where waiting permits further learning. Fudenberg and Koh (2026) studies a two-firm R&D race where firms observe and respond

to information about their opponent's progress and stopping is irreversible; it focuses on the effects of the speed of monitoring and beliefs about opponent's rationality.

Methodology of Continuous time games. The earliest papers on continuous-time games where each player controls a rate, such as Isaacs (1965), assume that equilibrium strategies are continuous. Fudenberg and Tirole (1983) studies discontinuous investment strategies in a model of strategic deterrence using the concept of *state space equilibrium*, which was subsequently generalized by Maskin and Tirole (2001) to games where the state is not a primitive of the model. In Bolton and Harris (1999) and Keller et al. (2005), each player chooses an intensity in $[0,1]$ at every date, the objective is linear in the intensity, and the authors must say which feedback rules generate a well-defined path. Simon and Stinchcombe (1989) and Bergin and MacLeod (1993) discuss the technical issues in defining strategies on continuous-time histories.

Observation Lags Information structures in which signals reach players with a delay have mostly been studied in repeated games. Abreu et al. (1991) show how bundling observations and publicly releasing them affects cooperation. Fudenberg and Olszewski (2011), Fudenberg et al. (2014), and Obara and Bhaskar (2019) study repeated games in which imperfect signals of opponent actions arrive with privately known stochastic lags. In our model, as in Fudenberg and Koh (2026), every action is observed after the same fixed lag, with no bundling and no noise; the lag matters because it gives a deviator a head start in a potential race.

Dynamic public goods. In our extension to endogenous safety, safety is a stock built up by voluntary, privately costly contributions. Fershtman and Nitzan (1991) study Markov-perfect contributions to a public-good stock and show that free riding leads to underprovision. Admati and Perry (1991) and Marx and Matthews (2000) study sequential contributions to a joint project and show that the prospect of future contributions can reduce free riding. In our paper, a firm that pulls ahead induces its rival to race and withdraw its safety contribution, which increases the probability of disaster.

2 MODEL

2.1 Development, safety, and disaster risk. Time is continuous, $t \geq 0$. There are two firms $i \in \{1,2\}$. Firm i chooses a development speed $a_{it} \in [0,1]$, which it can change at any time. Its total capability is $A_{it} = A_{i0} + \int_0^t a_{is}\ \mathrm{d}s$.[5] We call speed one *racing* and speed zero *pausing*. Pausing leaves the firm's capability unchanged, and the firm can resume

development at any time.

The *capability frontier* A_t is the higher of the two firms' capabilities: $A_t := \max\{A_{1t}, A_{2t}\}$. The *safety threshold* X_t advances at a constant speed $g \in (0,1)$: $X_t = X_0 + gt$. The capability frontier's distance above safety is $D_t := A_t - X_t$. Thus when $D > 0$ frontier capabilities are ahead of the safety threshold and vice versa for $D < 0$. If one firm has a strict lead and develops at speed a_L, then $\dot{D} = a_L - g$.

The disaster risk λ depends on distance above safety, with

$$\lambda(D) = \begin{cases} 0, & D \leq 0, \\ \lambda_0 + \phi(D), & D > 0, \end{cases} \tag{RISK}$$

where ϕ is continuous on $[0, \infty)$, nondecreasing, concave, and twice continuously differentiable on $(0, \infty)$, $\phi(0) = 0$, and $\lambda_0 \geq 0$. Since $\lambda(0) = 0$, pacing exactly along the safety threshold incurs no risk. Note that risk is generated by the capability frontier's position relative to safety (what Trammell and Aschenbrenner (2026) call "state risk") so existing capabilities that remain strictly above safety continue to generate risk even during a pause.

Disaster permanently ends both firms' future profits. Let $\Lambda_t := \int_0^t \lambda(D_s) ds$. The *disaster time* is the random variable τ_D with survival function $\Pr(\tau_D > t) = \exp(-\int_0^t \lambda(D_s)\ \mathrm{d}s) := \exp(-\Lambda_t)$.

2.2 Payoffs. To simplify the analysis, we use a functional form for the payoffs that implies the only source of non-stationarity in the model is the evolution of the safety threshold. Specifically, before disaster, firm i receives flow payoff $\pi_i(A_i, A_j) = \exp\{\alpha A_i - \beta A_j\}$,where $\alpha > \beta > 0$ and there is no direct cost of development. Own development raises profit at rate α; rival development lowers it at rate β. Since $\alpha > \beta$, advancing both capabilities equally raises both firms' profits. Firm i's expected payoff is

$$V_i = \mathbb{E}\left[\int_0^{\tau_D} e^{-rt} e^{\alpha A_{it} - \beta A_{jt}}\ \mathrm{d}t\right] = \int_0^{\infty} e^{-rt - \Lambda_t}\, e^{\alpha A_{it} - \beta A_{jt}}\ \mathrm{d}t,$$

where the equality uses the fact, explained next, that the speed paths are deterministic while the game is ongoing. We assume $r > \alpha - \beta$ throughout, so a tied pair has finite discounted payoffs even if both firms race forever without disaster.

[5]We will later impose measurability conditions that ensure this integral is well defined.

2.3 Histories, information, and strategies. The state S_t records both firms' capabilities and the safety threshold, $S_t = (A_{1t}, A_{2t}, X_t) \in \mathcal{S} := \mathbb{R}^2 \times \mathbb{R}$. Since $a_{it} \in [0,1]$, each firm's capability path is absolutely continuous, nondecreasing, and 1-Lipschitz. Conversely, every absolutely continuous path with derivative in $[0,1]$ almost everywhere can be written as $A_{it} = A_{i0} + \int_0^t a_{is}\, \mathrm{d}s$ for a measurable speed path $a_i : [0,\infty) \to [0,1]$, unique almost everywhere.

For each initial state $S = (A_1, A_2, X) \in \mathcal{S}$, let

$$\Omega_S(S) := \left\{ (S_t)_{t\geq 0} \in \mathcal{C}([0,\infty), \mathcal{S}) : \begin{array}{l} S_0 = S, \\ X_t = X_0 + gt \text{ for every } t \geq 0, \\ A_1, A_2 \text{ are absolutely continuous} \\ \dot{A}_{1t}, \dot{A}_{2t} \in [0,1] \text{ a.e.} \end{array} \right\}.$$

Let $\Omega_S := \bigcup_{S\in\mathcal{S}} \Omega_S(S)$ be the collection of feasible state paths from arbitrary initial states, and let $\Omega := \Omega_S \times [0,\infty]$ be the full outcome space, whose second coordinate is the disaster time τ_D. The information in the public history up to time t is given by the filtration $\mathcal{F}_t := \sigma(S_{t'},\ \mathbf{1}\{\tau_D \leq t'\} : t' \leq t)$; $\mathcal{F} := (\mathcal{F}_t)_{t\geq 0}$ tracks this information over time.

Definition 1 (Strategies)**.** A *pure strategy* for firm i is a map $\sigma_i : [0,\infty) \times \Omega \to [0,1]$ that is progressively measurable with respect to $\mathcal{F}$,[6] where $\sigma_i(t,\omega)$ is the prescribed speed at date t given the public history in ω up to that date.

A pure *Markov strategy* depends on the history only through the current state: it is a Borel measurable map $\sigma_i : \mathcal{S} \to [0,1]$ assigning a prescribed speed to each state $S = (A_1, A_2, X)$.[7]

2.4 Paths and payoffs at discontinuities. A Markov strategy prescribes a speed at each state, and the state evolves according to these prescribed speeds. The following example shows that strategies need not determine play.

Example 1 (Non-unique paths)**.** Suppose each firm races except when its capability equals safety, where it paces:

$$\sigma_i(S) = \begin{cases} g, & A_i = X, \\ 1, & A_i \neq X. \end{cases}$$

Start both capabilities at zero, with $X_0 > 0$. Both firms race until the time T where $A_{it} = T = X_t$. Because investment at precisely time T has no direct impact on the integral

[6] That is, for every $T < \infty$, its restriction to $[0,T] \times \Omega$ is $\mathcal{B}([0,T]) \otimes \mathcal{F}_T$-measurable.

[7] We are abusing notation here, because Markov strategies depend only on the most recent state.

defining the evolution of states, both the path $A_{it} = X_t = T + g(t - T)$ and the path $A_{it} = t$ are feasible paths.

We address this ambiguity by (i) specifying which paths are admissible under a strategy profile and following deviations; and (ii) assigning payoffs over those paths.

For any pair of Markov policies $\sigma = (\sigma_1, \sigma_2)$, define

$$K[\sigma](S) = \bigcap_{\varepsilon>0} \overline{\text{co}}\Big\{(\sigma_1(S'), \sigma_2(S')) : S' \in \mathcal{S},\ \|S' - S\| < \varepsilon\Big\}. \tag{K}$$

The set $K[\sigma](S) \subseteq [0,1]^2$ contains the joint speeds allowed at state S. Here $\overline{\text{co}}$ is the closed convex hull; it allows averages of policy values that are arbitrarily close to the state. This is the *Krasovskii regularization* (Krasovskii and Subbotin (1988)) generated by σ.[8] For a single Markov policy σ_j, we write $K[\sigma_j](S)$ for its allowed speeds under regularization. This set is defined by (K) with the scalar policy σ_j in place of the pair.

An *admissible path* from S is an absolutely continuous state path starting at S and satisfying, almost everywhere,

$$(\dot{A}_{1t}, \dot{A}_{2t}) \in K[\sigma](S_t) \quad \text{and} \quad \dot{X}_t = g. \tag{A}$$

Let $\mathcal{P}(\sigma; S)$ denote the set of paths starting at S satisfying (A); since this condition is a differential inclusion, $\mathcal{P}(\sigma; S)$ is its solution set from S, with speeds given by the path derivatives. The regularized map has nonempty compact convex values, is upper semi-continuous, and is bounded, so Filippov's existence theorem for differential inclusions (Theorem 2.1 in Filippov (1988)) supplies a global path from every state.

Risk and payoffs. Write $V_i[(S_t)_{t\geq 0}]$ for the original discounted payoff evaluated on a state path. Define $V_i(\sigma, S) = \sup_{(S_t)_{t\geq 0} \in \mathcal{P}(\sigma;S)} V_i[(S_t)_{t\geq 0}]$. Each firm values a profile by its best admissible continuation. This ensures that firms evaluate the "pacing profile" (defined later) under the pacing outcome. Note also that the suprema need not be attained. Payoffs may be $+\infty$ for arbitrary profiles, but our assumption that $r > \alpha - \beta$ ensures finiteness for tied paths. The equilibria constructed below have finite payoffs.

Deviations and equilibrium. A *deviation* replaces firm i's strategy by an alternative strategy σ_i', holding fixed the opponent's strategy σ_j. Write (σ_i', σ_j) for the profile with only

[8] Recent work in differential games e.g., Jaakkola and Wagener (2026) uses a regularization based on Filippov solutions. We use Krasovskii regularization to retain the speeds prescribed at exact ties and safety boundaries, even though these sets have zero volume in the state space. By contrast, Filippov regularization disregards prescriptions confined to such sets and applying it to Example 1 would rule out pacing. Bivas et al. (2020) compares the two regularizations.

firm i's strategy replaced. When discussing continuation play, a *history* is a path of states at which disaster has not occurred and we write $h_t = (S_s)_{0 \le s \le t}$ for such a history. It is *feasible* if its state path is absolutely continuous and $\dot{A}_{is} \in [0,1]$ for $i = 1,2$ and $\dot{X}_s = g$ almost everywhere on $[0,t]$.

Write $\mathcal{P}(\sigma; h_t)$ for the admissible continuations after h_t. Define the continuation payoff from history h_t under the strategy profile σ by

$$V_i(\sigma, h_t) := \sup_{(S_s) \in \mathcal{P}(\sigma; h_t)} \int_t^\infty e^{-r(s-t)-(\Lambda_s - \Lambda_t)} \pi_i(S_s)\, \mathrm{d}s.$$

Definition 2 (Markov-perfect equilibrium)**.** A *Markov profile* σ is a Markov-perfect equilibrium (MPE) if, after every feasible history h_t starting from any initial state, neither firm can gain from a deviation: for each firm i and every σ_i', $V_i((\sigma_i', \sigma_j), h_t) \le V_i(\sigma, h_t)$.

Although deviations can depend on the full history of past play, a Markov best response always exists under our path and payoff rules.

Lemma 1 (Markov best responses)**.** If the opponent uses a Markov strategy, then there is a Markov strategy that is a best response. Consequently, a Markov profile σ is an MPE if and only if, for every state S, firm i, and alternative Markov strategy σ_i', $V_i((\sigma_i', \sigma_j), S) \le V_i(\sigma, S)$.

This lemma would be immediate in a discrete-time version of this game, but requires a proof here because of regularization; the proof is in the Supplemental Appendix

Definition 3. An *MPE outcome* is an admissible path under an MPE profile that attains both firms' assigned continuation payoffs.

The separate definition of MPE outcome is needed here because the firms' best admissible continuations need not coincide, so the the equilibrium conditions alone do not ensure that one path gives both firms their assigned payoffs.[9] Our conclusions about low-risk racing and disaster probabilities concern equilibrium outcomes—paths attaining both assigned payoffs—not every admissible continuation of every equilibrium profile.

[9] To illustrate, assign the four joint speed prescriptions $(0,0)$, $(0,1)$, $(1,0)$, and $(1,1)$ on four disjoint countable dense subsets of state space, and prescribe joint pausing elsewhere. Every neighborhood of every state contains all four prescriptions and their convex hull is $[0,1]^2$, so Krasovskii regularization allows every feasible development path. Each firm's assigned payoff is consequently its highest payoff over all feasible paths so the profile satisfies the equilibrium inequalities, but the assigned payoffs can be incompatible. Suppose disaster risk is zero, $r > \alpha$, and both firms start tied at capability and safety zero. Each firm's maximum payoff is $1/(r-\alpha)$, attained only by racing forever while its rival pauses forever. These two individually optimal continuations cannot occur on the same path.

3 EQUILIBRIUM

3.1 Racing and pacing strategies. We normalize each firm's continuation payoff by its current flow profit. At a tie $A_1 = A_2 = A$, this profit is $e^{(\alpha-\beta)A}$. Let $P(D)$ be the normalized payoff when tied firms reach safety by pausing if $D > 0$ or racing if $D < 0$, and then pace forever. At $D = 0$ they pace, so $P(0) = \frac{1}{r-g(\alpha-\beta)}$. For $D > 0$, the wait lasts D/g. Profits and disaster risk continue during the pause, so the payoff is

$$P(D) = \int_0^{D/g} \exp\left\{-rt - \int_0^t \lambda(D - gs)\ \mathrm{d}s\right\} \mathrm{d}t + \exp\left\{-rD/g - \frac{1}{g}\int_0^D \lambda(s)\ \mathrm{d}s\right\} P(0).$$

The first term gives discounted profits while waiting for safety to catch up. The second is the discounted pacing value, weighted by survival until safety.

For $D < 0$, the firms race without risk for $-D/(1-g)$ units of time before pacing. Their normalized payoff is thus

$$P(D) = \frac{1 - e^{(r-\alpha+\beta)D/(1-g)}}{r - \alpha + \beta} + e^{(r-\alpha+\beta)D/(1-g)} P(0).$$

Let $R(D)$ be the normalized payoff when both firms race forever from a tie at D:

$$R(D) = \int_0^\infty \exp\left\{-(r - \alpha + \beta)t - \int_0^t \lambda\big(D + (1-g)s\big)\ \mathrm{d}s\right\} \mathrm{d}t.$$

Define the *effective hazard* as the constant hazard that gives the payoff from racing starting at the safety boundary

$$\lambda^{\text{eff}} := \frac{1}{R(0)} - (r - \alpha + \beta).$$

It equals λ_0 when $\phi = 0$ and rises when the risk profile ϕ rises pointwise. It also depends on the safety speed g and the discounted growth rate $r - \alpha + \beta$.

Definition 4 (Racing)**.** The *racing profile* $\sigma^R = (\sigma_1^R, \sigma_2^R)$ prescribes $\sigma_i^R(S) := 1$ for each firm i at every state S.

We construct the pacing profile when $\lambda^{\text{eff}} \geq (1-g)(\alpha - \beta)$, or equivalently when $R(0) \leq P(0)$. Write $G := |A_1 - A_2|$ for the *capability gap* between firms. When this gap is strict, the leader has capability $A^L = A$ and the follower has capability $A^F = A - G$.

If $P(0) = R(0)$, set $\bar{D} = 0$. If $P(0) > R(0)$, let $\bar{D} > 0$ be the unique solution to $P(\bar{D}) =$

$R(\bar{D})$.[10] For tied firms, define

$$a^*(D) := \begin{cases} 1, & D < 0, \\ g, & D = 0, \\ 0, & 0 < D \leq \bar{D}, \\ 1, & D > \bar{D}. \end{cases}$$

Definition 5 (Pacing)**.** If $\lambda^{\text{eff}} \geq (1-g)(\alpha-\beta)$, the *pacing profile* $\sigma^P = (\sigma_1^P, \sigma_2^P)$ prescribes, for firm i with opponent j,

$$\sigma_i^P(S) := \begin{cases} a^*(D), & A_i = A_j, \\ 1, & A_i < A_j, \\ b(D,G), & A_i > A_j. \end{cases}$$

Here $b(D,G)$ is the leader's best response when the follower races; we characterize it below.

Since the follower races until it catches up (if it does), the leader faces a well-defined optimization problem. Let $V_L(D,G)$ denote the leader's optimal discounted payoff divided by $e^{(\alpha-\beta)A}$, the flow profit each firm would earn if they were both at the current capability frontier. With this normalization, the leader's current flow profit is $e^{\beta G}$ and its Bellman equation is

$$\underbrace{rV_L}_{\text{discounting}} = \underbrace{e^{\beta G}}_{\text{current profit}} - \underbrace{\lambda(D)V_L}_{\text{disaster loss}} + \max_{a\in[0,1]} \Big\{ \underbrace{(a-g)\partial_D V_L}_{\substack{\text{change in distance} \\ \text{above safety}}} + \underbrace{(a-1)\partial_G V_L}_{\substack{\text{change in} \\ \text{the lead}}} + \underbrace{(\alpha-\beta)aV_L}_{\substack{\text{growth of} \\ \text{frontier flow}}} \Big\}.$$

The first term in the continuation value captures how current speeds translate into the distance above safety (since $\dot{D} = a - g$); the second term captures how current speeds change the lead over the leader's rival (since $\dot{G} = a - 1$); the last term accounts for growth of the normalizing factor. Because the pacing profile determines speeds when $G = 0$, the boundary condition is $V_L(D,0) = \max\{P(D), R(D)\}$.

We show that the leader attains its optimal payoff by racing until a chosen date and

[10] This cutoff exists and is unique; see Lemma 6 in Appendix A.

then *accommodating*: pausing until the follower catches up, after which both firms choose speeds a^*. The optimal accommodation date may be zero, finite and positive, or infinite, in which case the leader races forever.

Suppose $P(0) > R(0)$ and $0 \leq D \leq \bar{D}$. The leader chooses between racing forever and racing to a point $d \in [D, \bar{D}]$ before accommodating. The gain from the latter, relative to racing forever, is

$$\exp\left[-\frac{1}{1-g}\int_D^d \big(r-\alpha+\beta+\lambda(s)\big)\ \mathrm{d}s\right]\ \big[C(d,G)-R(d)\big].$$

Here $C(d, G)$ is the leader's normalized payoff from accommodating immediately at (d, G); its formula is given in the Supplemental Appendix. The choice $d = D$ means immediate accommodation. The leader accommodates if the maximum gain is nonnegative, selecting the smallest maximizing d; otherwise it races forever. If $D < 0$, it races to safety before applying this rule. If $D > \bar{D}$, it races forever.

The leader's best response to a racing follower is

$$b(D,G) := \begin{cases} 1, & P(0) = R(0), \\ g, & P(0) > R(0),\ D = 0,\ \text{and } V_L(0,G) = e^{\beta G}C(0,G), \\ 0, & P(0) > R(0),\ 0 < D \leq \bar{D},\ \text{and } V_L(D,G) = e^{\beta G}C(D,G), \\ 1, & \text{otherwise.} \end{cases}$$

The equality $V_L(D, G) = e^{\beta G}C(D, G)$ means that immediate accommodation attains the leader's best payoff. When it holds, the leader paces at speed g along the safety threshold and pauses when $0 < D \leq \bar{D}$. Otherwise it strictly prefers to keep racing, either forever or until a later accommodation date. Figure 1 is a phase diagram under the pacing strategy σ^P, where the x-axis is distance D of the capability frontier above safety and the y-axis is the gap G between leader and follower. Racing moves the state right: $(\dot{D}, \dot{G}) = (1 - g, 0)$. For $D < 0$ it leads to safety; in the red region at $D > 0$ it continues forever. Pausing in blue moves the state down and left: $(\dot{D}, \dot{G}) = (-g, -1)$. On the green line the leader paces and the gap closes vertically.

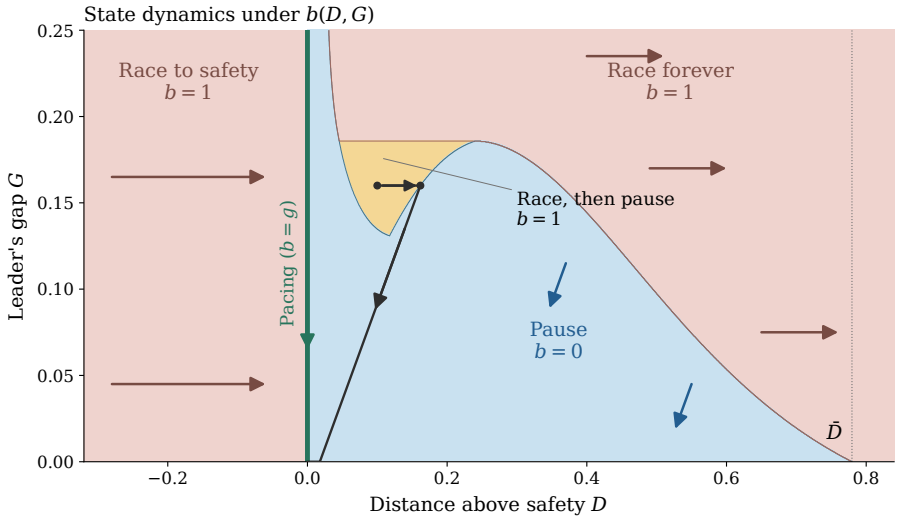

Figure 1: Phase diagram induced by σ^P under linear risk

3.2 Markov Perfect Equilibria.

Theorem 1 (Pacing and racing)**.**

(a) *Racing only.* If $\lambda^{\text{eff}} < (1-g)(\alpha-\beta)$, σ^R is an MPE, and no MPE outcome paces forever from a tie on safety.

(b) *Racing and pacing.* If $(1-g)(\alpha-\beta) \leq \lambda^{\text{eff}} \leq (1-g)\alpha$, σ^R and σ^P are both MPE.

(c) *Pacing only.* If $\lambda^{\text{eff}} > (1-g)\alpha$, σ^P is an MPE, and σ^R is not an MPE.

Theorem 1 does not rule out other MPE. However, from a safety tie, every MPE outcome gives each firm at least $R(0)$ (Lemma 5), and every MPE outcome on which the firms remain tied gives each firm at most $\max\{P(0), R(0)\}$ (Lemma 6 with $(\eta, \rho) = (\alpha-\beta, r)$). Under the hypotheses of Theorem 2(b), the disaster-probability bound applies to every SPE outcome, and hence to every MPE outcome.

Figure 2 illustrates the theorem. The lower cutoff rules out perpetual pacing in part (a). Start with both firms at the safety threshold, $A_{1,0} = A_{2,0} = X_0 = A_0$. If firm i races forever, the speed limit implies $A_{it} = A_0 + t$ and $A_{jt} \leq A_0 + t$. Firm i therefore determines the frontier, so $D_t = (1-g)t$, exactly as under joint racing. Its flow profit satisfies $e^{\alpha A_{it} - \beta A_{jt}} \geq e^{(\alpha-\beta)(A_0+t)}$. The deviation thus faces the same disaster risk and guarantees normalized payoff $R(0)$, whatever the rival does. Perpetual joint pacing gives $P(0)$ so when $R(0) > P(0)$, either firm can profitably deviate from pacing by racing

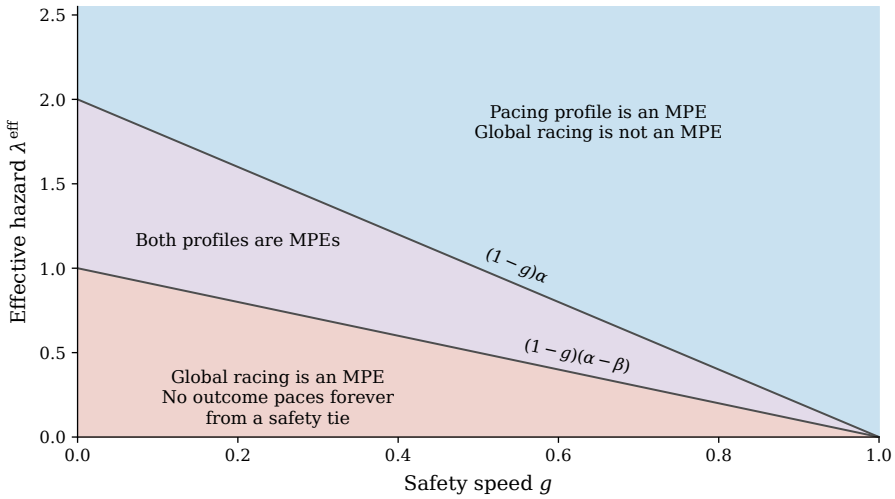

Figure 2: Regimes of cooperation

forever. This inequality is equivalent to $\lambda^{\text{eff}} < (1-g)(\alpha-\beta)$. In part (b), perpetual joint racing and permanent joint pacing are both equilibrium outcomes from a safety tie: When $(1-g)(\alpha-\beta) < \lambda^{\text{eff}} < (1-g)\alpha$, each firm prefers joint pacing, but neither gains by slowing down alone from joint racing.

Figure 3 illustrates coexistence of these equilibria.[11] It starts at $A_{10} = X_0 = 0.20$ and $A_{20} = 0.10$. In both panels, the leader paces while the follower races until catch-up. Under the pacing profile, both firms then pace, so the leader's initial slowdown leads to a permanently safe continuation. Panel (b) shows an unprofitable deviation against the racing profile: the leader tries pacing until catch-up and then resumes racing, while the rival races throughout. The slowdown provides only a temporary safe interval and sacrifices the lead. It gives the leader 0.3396 per unit of its initial profit flow, below the 0.3448 from racing immediately. The follower also has an incentive to race, because slowing down makes immediate accommodation less attractive to the leader.

[11] We use linear risk with $\alpha = 2$, $\beta = 1$, $r = 3$, $g = 1/2$, and $\kappa \simeq 6.4124$. The slope is chosen to give $\lambda^{\text{eff}} = 0.9$, strictly between the two cutoffs.

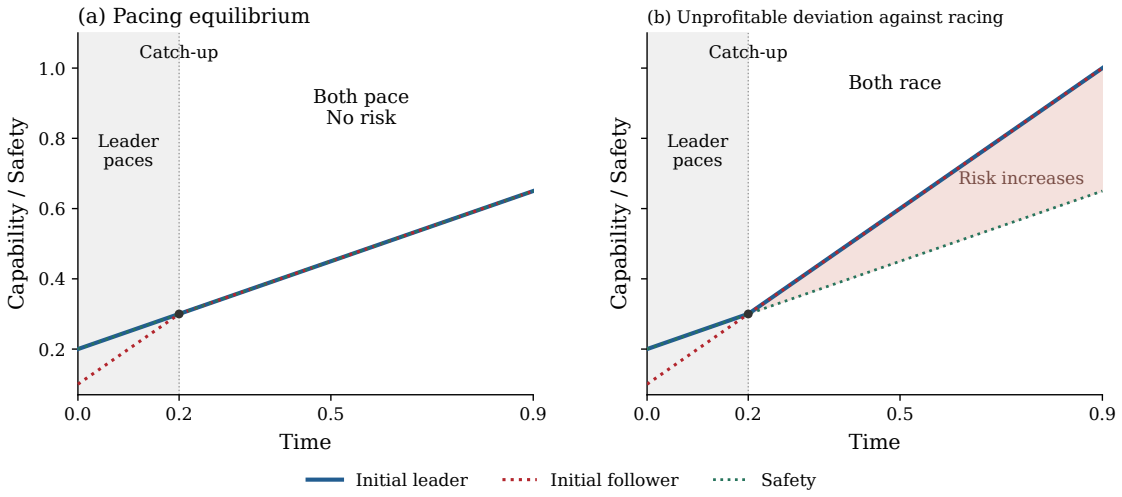

Figure 3: Slowdown with different continuations

Note that the pacing profile specifies behavior at every state, but it need not lead to pacing from every initial state. With safety far behind, the leader races forever.

In part (c), pacing remains an equilibrium but global racing does not. Above the upper cutoff, a leader facing a racing rival prefers to pace briefly before resuming racing: the temporary reduction in disaster risk outweighs the cost of letting the follower narrow the gap. Hence global racing cannot be an equilibrium.

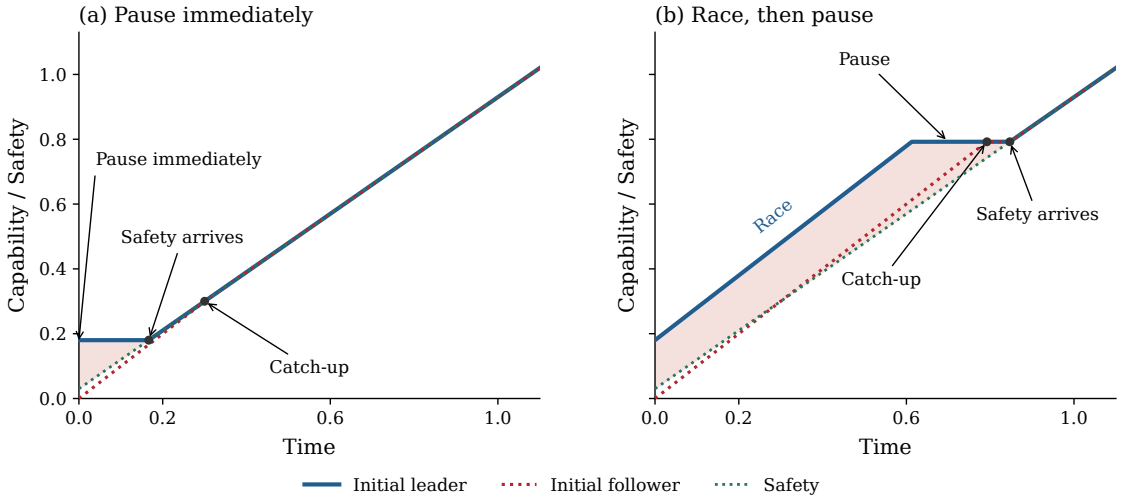

Figure 4: Different accommodation timing; same linear-risk parameters as Figure 3 (see Footnote 11).

When the leader starts above safety, it might prefer a short race before accommodating. Figure 4 (a) illustrates the leader's strategy of pausing immediately until safety catches up, and then pacing the frontier while the follower develops at maximum speed. Eventually, the follower catches up, and both firms pace the frontier. Panel (b) illustrates the leader's strategy of racing before pausing.

Figure 5 enlarges the part of Figure 1 containing the path in Figure 4(b). Racing moves the state to the right at a fixed gap: $(\dot{D}, \dot{G}) = (1 - g, 0)$. At the boundary of the pausing region, the leader stops while the follower keeps racing, so $(\dot{D}, \dot{G}) = (-g, -1)$ and the path turns down and to the left. Catch-up occurs before safety arrives. Both firms then pause until $D = 0$ and then pace together.

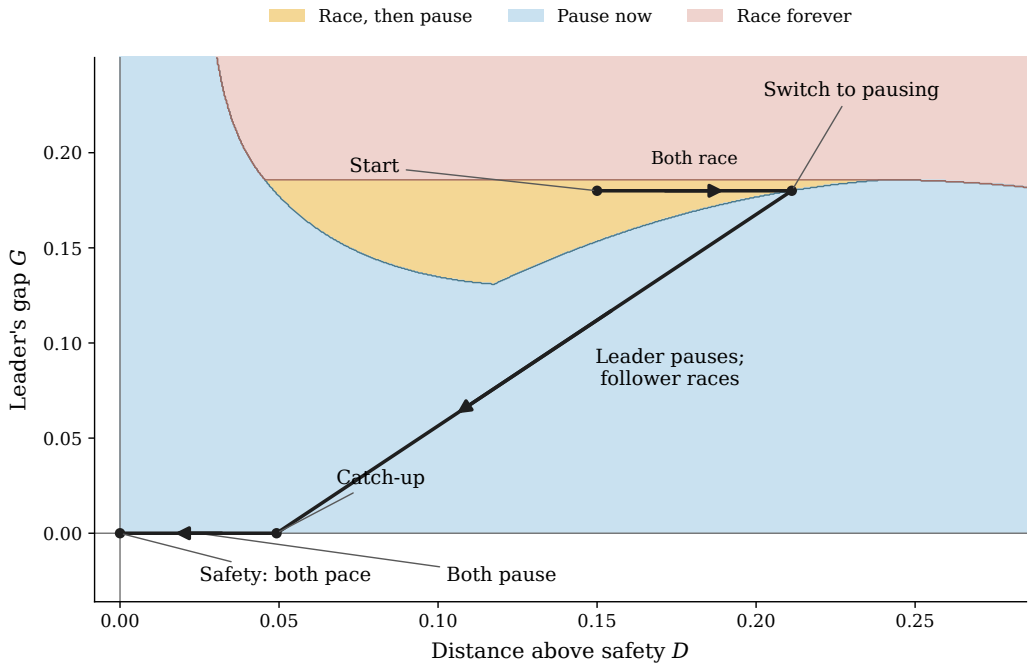

Figure 5: The race-then-pause path in state space

3.3 Comparative statics. We now flesh out some simple implications of Theorem 1.

Proposition 1 (Safety progress and patience favor pacing)**.** With other primitives fixed, the ratio $\lambda^{\text{eff}}/(1 - g)$ is:

(a) weakly increasing in safety speed g;

(b) weakly increasing under pointwise increases of the hazard profile λ;

(c) weakly decreasing in the discount rate r, and strictly if ϕ is nonconstant.

Thus faster safety, higher risk, or greater patience weakly expands the region where σ^P is an MPE and weakly shrinks the region where σ^R is an MPE.

Faster safety reduces hazard exposure along a racing path, but it also reduces racing's profit-growth advantage over pacing. Under concave risk, the reduction in exposure is smaller, proportionally, than the reduction in that growth advantage, so pacing becomes more attractive. Patience works through a different channel: It puts more weight on the later, riskier stages of a race, an effect that is absent under constant risk. The proof is in the Supplemental Appendix.

Under constant risk strictly above safety, effective hazard equals λ_0. The two cutoffs are $(1-g)(\alpha-\beta)$ and $(1-g)\alpha$, both decreasing in g and independent of r. The leader chooses between immediate accommodation and perpetual racing: any gain from eventually accommodating is largest when accommodation starts immediately. Under linear risk, $\lambda_0 = 0$ and $\phi(D) = \kappa D$, the racing payoff from safety is

$$R(0) = \int_0^\infty e^{-(r-\alpha+\beta)t-\kappa(1-g)t^2/2}\,\mathrm{d}t.$$

Let the lower and upper cutoff slopes be $\underline{\kappa}$ and $\overline{\kappa}$, defined respectively by $R(0) = P(0)$ and $R(0) = 1/(r+\beta-g\alpha)$. Pacing exists for $\kappa \geq \underline{\kappa}$, global racing exists for $\kappa \leq \overline{\kappa}$, and both exist between the cutoffs, including their endpoints.

A higher hazard can reduce the eventual probability of disaster by inducing restraint. Fix an initial state with $D > 0$ and $G > 0$, and suppose the pacing equilibrium prescribes perpetual racing under λ, but racing to a finite distance $d \geq D$ followed by accommodation under a pointwise higher hazard $\lambda' \geq \lambda$. Under λ, cumulative hazard is infinite, so disaster occurs with probability one. Under λ', distance above safety first rises from D to d at speed $1-g$, then falls to zero at speed g so the eventual disaster probability is

$$1 - \exp\left\{-\frac{1}{1-g}\int_D^d \lambda'(x)\,\mathrm{d}x - \frac{1}{g}\int_0^d \lambda'(x)\,\mathrm{d}x\right\} < 1.$$

Thus, the switch from perpetual racing to accommodation outweighs the direct increase in hazard, although the disaster arrival times are typically not ordered in first-order stochastic dominance.

3.4 Subgame Perfect Equilibria. The multiplicity in Theorem 1 raises the question of whether richer history dependence permits more cooperation, or if low risk inevitably

produces racing in all subgame-perfect equilibria.

Definition 6 (Subgame-perfect equilibrium)**.** A pure profile σ is a subgame-perfect equilibrium (SPE) if, after every feasible history h_t, for each firm i and every alternative history-dependent strategy σ_i', $V_i((\sigma_i', \sigma_j), h_t) \leq V_i(\sigma, h_t)$. An *SPE outcome* is an admissible continuation that attains both assigned payoffs on the same path.

We establish that when risk is low and the future is discounted sufficiently, firms perpetually race across all SPE. We also bound the probability of disaster up to a fixed time. To show this, we first construct a lower-bound $\underline{V}_\infty$ on firms' payoffs across all SPE. We do so recursively, as illustrated in Figure 6. At the rightmost safety tie, each firm can guarantee its racing payoff $\underline{V}_0 = R(0)$. Thus, at the middle tie, firm 2 (red) can pause which gives firm 1 (blue) the option to wait and then pace to the rightmost tie, where it can guarantee $\underline{V}_0$ The logic here is that because disaster ends both firms' profits, firm 1's incentive to avoid disaster also protects firm 2. This gives firm 2 the improved payoff guarantee $\underline{V}_1$. Now working backward, at the initial tie, firm 1 can in turn pause and—with the roles reversed—use firm 2's continuation guarantee $\underline{V}_1$ to obtain $\underline{V}_2$. Repeating the argument gives guarantees converging to $\underline{V}_\infty$.

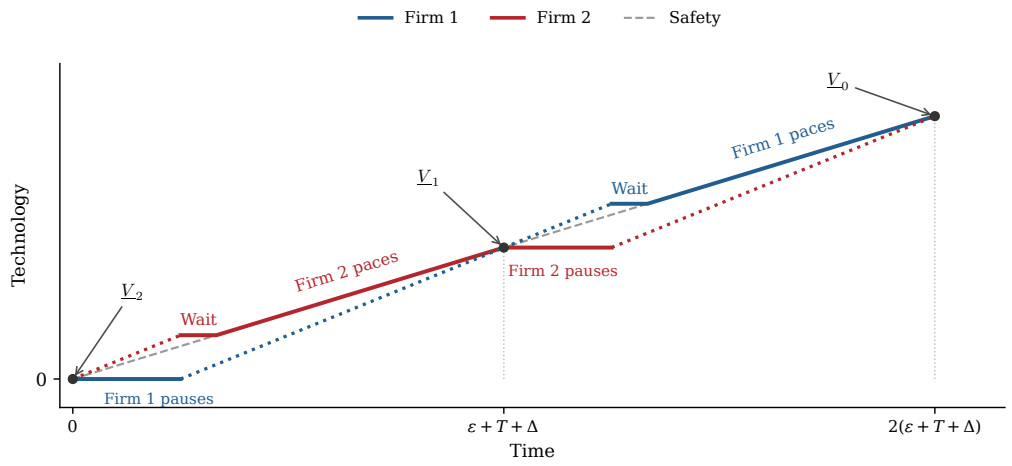

Figure 6: Constructing the recursive payoff guarantee

Next, we construct an upper bound $\overline{V}(t, z)$ on average discounted profits across all SPE when cumulative disaster exposure reaches z by date t. To do so, we first observe that cumulative disaster by time t is upper-bounded by $\Lambda_t^R := \int_0^t \lambda\big((1-g)s\big)\, ds$, since $(1-g)s \geq D_s$. Next notice that $F(t) := \max\left\{e^{(\alpha-\beta)t}, \frac{e^{\alpha t}+e^{-\beta t}}{2}\right\}$ is the largest average flow

profit by date t.[12] Thus define

$$\overline{V}(t,z) := \underbrace{\int_0^t e^{-rs} F(s) \exp\Big[-\max\{0, z - \Lambda_t^R + \Lambda_s^R\}\Big]\, \mathrm{d}s}_{\substack{\text{Profits through } t \\ \text{with exposure postponed as long as possible}}} + \underbrace{e^{-z} \int_t^\infty e^{-rs} F(s)\, \mathrm{d}s}_{\substack{\text{Profits after } t \\ \text{with no further exposure}}}.$$

which is decreasing in z. Thus, the exposure ceiling

$$\overline{\Lambda}(t) := \max\left\{z \in [0, \Lambda_t^R] : \overline{V}(t,z) \geq \underline{V}_\infty\right\}$$

gives the largest cumulative exposure consistent with our SPE payoff guarantees.

Theorem 2 (Racing and risk bounds in every SPE)**.**

(a) *Low risk forces racing.* If $\lambda^{\text{eff}} < (1-g)(\alpha-\beta)$ and $r + \lambda^{\text{eff}} > \frac{\alpha^2+\beta^2}{\alpha-\beta}$, every SPE outcome starting from a safety tie is perpetual joint racing.

(b) *Upper bound on risk.* If $r > \alpha$ and $g > 1/2$ and firms start from a tie on the safety threshold, then every SPE outcome satisfies $\Pr(\tau_D \leq t) \leq 1 - e^{-\overline{\Lambda}(t)}$ for every $t \geq 0$. Also, every SPE gives each firm normalized payoff at least the amount $\underline{V}_\infty$.

Part (a) of Theorem 2 shows that firms cannot cooperate on pacing the frontier when risks are sufficiently low. The underlying idea is that each firm can guarantee its racing payoff $R(0)$ against any response, so an SPE must deliver at least $2R(0)$ in total profits. But under part (a)'s assumptions, low risk makes delay costly, while sufficient discounting limits the future gains from creating a capability lead. Hence, the total payoff from racing of $2R(0)$ attains the upper bound while every other outcome gives strictly less and so cannot be an SPE.

The next result shows that history-dependent cooperation can give both firms strictly more than the pacing payoff in the region where racing and pacing coexist as MPEs.

Proposition 2 (An SPE with higher payoffs than pacing)**.** Suppose $\lambda(D) = \lambda$ for every $D > 0$, so that $\lambda^{\text{eff}} = \lambda$, and that $(1-g)(\alpha-\beta) \leq \lambda \leq (1-g)\alpha$ and $r + \lambda < \alpha$. Starting from $A_{10} = A_{20} = X_0 = 0$, there is an SPE outcome giving both firms strictly more than the pacing-MPE payoff $P(0) = 1/[r - g(\alpha-\beta)]$.

The proof is in the Supplemental Appendix. When $r + \lambda < \alpha$, a firm's discounted

[12] The first term is if it maximizes total profits for both firms to be tied, and the second is if one firm races while the other's capability is at 0.

flow profit rises while it races and its rival pauses. The construction uses alternating development turns to give both firms more than permanent pacing.

Finally, we analyze when the threat of future racing can sustain permanent pacing under more general profit functions. For this result we allow profits to be any positive, continuously differentiable functions $\pi_i(A_i, A_j)$ such that for each firm i $R_i(S) < \infty$ for every state S and $P_i(x) < \infty$ for every $x \in \mathbb{R}$, and at every capability pair A_i, A_j, $\dfrac{\partial \log \pi_i(A_i, A_j)}{\partial A_i} \geq \alpha > 0$ and $\dfrac{\partial \pi_i(A_i, A_j)}{\partial A_j} \leq 0$. Here α is a common lower bound on the proportional return to own development. We also assume $\sup_{D>0} \lambda(D) =: \bar{\lambda} \leq (1-g)\alpha$.[13]

From a state $S = (A_1, A_2, X)$, with $D = \max\{A_1, A_2\} - X$, define the normalized mutual-racing payoff $R_i(S)$ and the normalized joint-pacing payoff $P_i(x)$ from a safety tie at capability x by

$$R_i(S) := \int_0^\infty e^{-rs - \int_0^s \lambda(D+(1-g)t)\, \mathrm{d}t} \frac{\pi_i(A_i + s, A_j + s)}{\pi_i(A_i, A_j)}\, \mathrm{d}s,$$
$$P_i(x) := \int_0^\infty e^{-rs} \frac{\pi_i(x + gs, x + gs)}{\pi_i(x, x)}\, \mathrm{d}s.$$

Under exponential profits, the safety-tie values reduce to $R_i(x, x, x) = R(0)$ and $P_i(x) = P(0)$.

Proposition 3 (Pacing is sustainable if and only if it beats racing)**.** When $\bar{\lambda} \leq (1-g)\alpha$, permanent joint pacing from $A_{10} = A_{20} = X_0 = A_0$ is an SPE outcome if and only if for each firm i and each $x \geq A_0$, $P_i(x) \geq R_i(x, x, x)$. When this holds, a grim trigger implements pacing: both firms pace while the history follows the prescribed pacing path, and race forever after any deviation.

The idea driving Proposition 3 is that when $\bar{\lambda} \leq (1-g)\alpha$, mutual racing is a credible continuation that attains both firms' lowest equilibrium payoffs at every state. With exponential profits, the payoff comparison reduces to $P(0) \geq R(0)$, or equivalently $\lambda^{\text{eff}} \geq (1-g)(\alpha - \beta)$ which is exactly the condition for permanent pacing in Markov strategies in Theorem 1.

We use the grim strategy for convenience, but whenever deviating from equilibrium generates a strict loss, a milder punishment would suffice. This is of interest because when there is small noise in the system, the grim punishments might be imposed even though neither player actually deviated. By analogy with the literature on repeated games with

[13] The argument holds for more general increasing hazards that satisfy the same bound.

imperfect public monitoring (Abreu et al., 1990; Fudenberg et al., 1994), [14] we conjecture that strategies that respond to deviations of the capability frontier from its projected level with mutual racing for a finite duration can often enforce with higher equilibrium payoffs.

3.5 Monitoring lags: internal vs deployed models. In the baseline model, a firm's capability is observed immediately. We now separate internal capability from deployed capability: development is observed by the rival only after a fixed lag L, when it is deployed and begins affecting profits and risk. We show that low effective hazard rules out pacing, intermediate risk can sustain pacing when deployment is sufficiently prompt, and high effective hazard sustains pacing at every deployment lag.

The two firms develop capability at speeds in $[0, 1]$. An internal model becomes public and is deployed after a fixed lag $L > 0$. We assume the initial commonly known history is that both firms developed at speed g: $A_{is} = X_s + gL, s \in [-L, 0], i = 1, 2$,so $A_0 = X_0 + gL$. At date t, deployed capabilities determine profits and risk: $\pi_{it} := e^{\alpha A_{i,t-L} - \beta A_{j,t-L}}$ and $\lambda_t = \lambda(A_{t-L} - X_t)$. We retain the hazard profile in (RISK) and assume a finite ceiling $\bar{\lambda} := \sup_{D>0} \lambda(D)$. The initial histories place deployed capabilities at the safety threshold during $[0, L]$, so profits over that interval are fixed and hazard is zero.

Firm i observes its own development through date t and its rival's capability through $\max\{0, t - L\}$; both firms observe survival and the safety threshold $X_t = X_0 + gt$.

Strategies and equilibrium. Firm i's private history h_{it} records its own development, its observations of its rival's capabilities, and survival through t. A private history is feasible if there exists a pair of absolutely continuous capability histories extending the initial histories, with speeds in $[0, 1]$ almost everywhere, that reproduces the firm's own development and its observed rival deployments. As in the baseline model, we regularize joint speed prescriptions on feasible full histories and evaluate each profile by each firm's best admissible path. The formal definition, including the treatment of firms' different private histories, is in the Supplemental Appendix.

Definition 7 (Pure Strategy Nash Equilibrium)**.** A pure strategy profile σ is a *Nash equilibrium* if, for each firm i and every pure strategy σ_i', $V_i(\sigma) \geq V_i((\sigma_i', \sigma_j))$. A *Nash equilibrium outcome* is an admissible path attaining both firms' assigned payoffs on the same path, with play stopped at disaster.

Definition 8 (Joint pacing under delayed deployment)**.** An equilibrium outcome sustains *joint pacing* if, conditional on survival, $A_{1,t-L} = A_{2,t-L} = X_t$ at every date $t \geq 0$.

[14] See e.g. the T period punishments used in Green and Porter (1984).

Deployment lags that sustain pacing. If a firm races from time 0, the capability of its deployed technology at time $L+s$ exceeds safety by $(1-g)s$ so $\Lambda_s^R := \int_0^s \lambda((1-g)t)\,\mathrm{d}t$ is the associated cumulative disaster exposure through time $L+s$.

Let $B(L)$ be the normalized payoff from racing against a rival that paces for L units of time and then races forever:

$$B(L) = \underbrace{\int_0^L e^{-(r-\alpha+\beta g)s-\Lambda_s^R}\,\mathrm{d}s}_{\text{rival paces before detection}} + \underbrace{e^{\beta(1-g)L}\int_L^\infty e^{-(r-\alpha+\beta)s-\Lambda_s^R}\,\mathrm{d}s}_{\text{rival races after detection}}.$$

As before, pacing gives value $P(0) = 1/[r-g(\alpha-\beta)]$.[15] Let $\widehat{B}(L)$ be the largest normalized payoff a firm can obtain against a rival that paces for L units of time and then races forever. Unlike $B(L)$, this value allows the firm to choose any development path; its formal definition is in the Supplemental Appendix.

Proposition 4 (Risk and delay determine sustainable pacing)**.** For every lag $L > 0$, joint racing is a pure Nash equilibrium outcome. Joint pacing is a pure Nash equilibrium outcome if and only if $\widehat{B}(L) \le P(0)$. In particular:

(a) If $\lambda^{\text{eff}} < (1-g)(\alpha-\beta)$, then no pure Nash equilibrium sustains joint pacing for any $L > 0$.

(b) If $\lambda^{\text{eff}} \ge (1-g)(\alpha-\beta)$ and $\bar{\lambda} < (1-g)\alpha$, then joint pacing is a pure Nash equilibrium outcome if and only if $L \le L^*$, where $L^* \ge 0$ is the unique finite solution to $B(L^*) = P(0)$.

(c) If $\lambda^{\text{eff}} \ge (1-g)\alpha$, then joint pacing is a pure Nash equilibrium outcome for every $L > 0$.

Parts (a)–(c) are not exhaustive: When $(1-g)(\alpha-\beta) \le \lambda^{\text{eff}} < (1-g)\alpha \le \bar{\lambda}$, pacing is sustainable if and only if $\widehat{B}(L) \le P(0)$, but neither the simple lag-cutoff characterization in (b) nor the all-lags conclusion in (c) applies.

The strategy that sustains pacing in part (b) prescribes speed g while the firm has always paced and its observations of its rival are consistent with pacing, and speed 1 otherwise. During the detection lag, a firm that deviates to racing builds a lead of $(1-g)L$, which

[15] Starting from total payoffs, we first subtract profits over $[0, L]$ (that cannot be changed by speed choices from 0) and then divide by e^{-rL} since current speed choices only affect profits after L units of time, and divide by $e^{(\alpha-\beta)A_0}$ which is the profit scale from initial capabilities. Hence, the pacing payoff under this normalization is $P(0)$.

persists once both firms race, so a longer lag increases the payoff from this deviation and makes pacing harder to sustain.

Proposition 4 is stated for Nash equilibrium, making the impossibility of sustaining pacing in part (a) stronger. The construction in part (b) makes racing optimal after a detected deviation, independently of beliefs about hidden rival development. The strategy in part (c) prescribes speed g after every history. It is not sequentially rational after every history: a firm that observes a rival's strict lead would prefer to race toward catch-up.

Proposition 4 implies that firms can benefit from committing to revealing internal progress before deployment. This might reflect, for instance, the use of a third-party 'embedded evaluator' like Model Evaluation and Threat Research (METR) that tracks the state of internal R&D. Doing so can shorten or eliminate the lag between internal development and public observation, which allows pacing to be sustained for a wider set of parameters.

3.6 Endogenous safety. Our main model takes safety progress as exogenous. In practice, firms can advance the safety threshold through research and tools that they share publicly (Delaney et al., 2024; Anthropic, 2025). To do so, they face a tradeoff between (private) capabilities quickly and contributing to (public) safety research that advances the safety threshold.[16] This introduces a dynamic public good dimension to our analysis, and we analyze how incentives to contribute vs free ride might aid or hinder cooperation.

Firm i chooses development speed $a_i \in [0,1]$ and devotes the remaining research capacity $1 - a_i$ to safety. The safety frontier now evolves according to

$$\dot{X} = g + \eta\big[(1-a_1) + (1-a_2)\big], \tag{RES}$$

where $g \in (0,1)$ is background safety progress and $\eta \in [0, 1-g]$ measures the productivity of firms' safety research, so $\eta = 0$ recovers the model with exogenous safety. Profits and information remain as in the baseline model. We will assume the linear hazard $\lambda(D) = \kappa \max\{D, 0\}$, where $\kappa > 0$ is the hazard slope and $D = \max\{A_1, A_2\} - X$ is distance from safety. Admissible paths, payoffs, and equilibrium concepts are defined analogously, replacing $\dot{X} = g$ with (RES). We retain the joint speed regularization in (HK). Since the new safety velocity is affine in the development speeds, this is equivalent to jointly regularizing all three state velocities.

For both firms to remain tied on safety, their common development speed must be

[16] In July 2023 when OpenAI pledged to devote 20% of the compute it had secured to alignment over the following four years (Leike and Sutskever, 2023), but Kahn reported that the 'Superalignment team' did not receive the promised compute.

$v := \dfrac{g+2\eta}{1+2\eta}$. A strict leader that stays on safety while its follower races develops at speed $(g+\eta)/(1+\eta) < 1$, so the follower catches up.

Let $R(D)$ be each firm's payoff per unit of its current profit under perpetual joint racing. At a nonnegative distance from safety,

$$R(D) = \int_0^\infty e^{-[r-\alpha+\beta+\kappa D]t-\kappa(1-g)t^2/2}\,\mathrm{d}t.$$

As before, effective hazard is $\lambda^{\text{eff}} = 1/R(0) - (r-\alpha+\beta)$. Let $P(D)$ be the normalized payoff when tied firms pause until safety arrives and then jointly pace. For $D \geq 0$,

$$P(D) = \int_0^{D/(g+2\eta)} e^{-rt-\kappa[Dt-(g+2\eta)t^2/2]}\,\mathrm{d}t + \frac{e^{-rD/(g+2\eta)-\kappa D^2/[2(g+2\eta)]}}{r-(\alpha-\beta)v}.$$

The pacing payoff on safety is $P(0) = 1/[r-(\alpha-\beta)v]$.

The *racing profile* prescribes speed one at every state. To define the *pacing profile*, suppose first that $P(0) > R(0)$, and let $\bar{D}$ be the unique positive crossing of P and R. Tied firms race below safety and above $\bar{D}$, pause when $0 < D \leq \bar{D}$, and pace at speed v on safety. A strict follower races. A strict leader races to an optimally chosen distance from safety and then accommodates: it pauses above safety and then stays on the safety threshold until catch-up. As in our construction of the pacing profile in the main analysis, the leader chooses either the the earliest optimal accommodation date or never accommodates. When $P(0) = R(0)$, the pacing profile prescribes speed one except at safety ties, where both firms choose v.

For the pacing result, we impose

$$r > \alpha \quad \text{and} \quad \eta\kappa \leq \min\left\{\alpha(r-\alpha),\ \frac{\beta(r-\alpha+\beta)(g+2\eta)}{1-g}\right\}. \tag{ES}$$

Proposition 5 (Endogenous safety)**.** If the hazard is $\lambda(D) = \kappa \max\{D, 0\}$, then

(i) The racing profile is an MPE if and only if $\lambda^{\text{eff}} \leq \dfrac{\alpha(1-g)}{1+\eta}$.

(ii) If (ES) also holds, then perpetual joint pacing from a safety tie is an MPE outcome if and only if $\lambda^{\text{eff}} \geq \dfrac{(\alpha-\beta)(1-g)}{1+2\eta}$.

The logic of Proposition 5 is that endogenous safety provision creates a temptation to free ride, but also strengthens the response to a deviation. A firm can try to develop faster

while benefiting from its rival's safety effort, but trying to pull ahead can induce the rival to race to catch up and withdraw its safety contribution, and the prospect of losing the rival's future contributions can help discipline the incentive to free ride today. Under (ES), this response is credible: racing is optimal for the follower given the proposed leader strategy.[17]

More productive safety research raises the joint pacing speed v, allowing firms to develop faster while remaining safe. The higher pacing payoff lowers the effective hazard needed to sustain pacing under (ES). More productive safety research also makes slowing down more attractive to a leader facing a racing rival, because the resources released from development accelerate safety progress. This lowers the effective-hazard cutoff above which global racing is not an equilibrium.

REFERENCES

D. Abreu, D. Pearce, and E. Stacchetti. Toward a theory of discounted repeated games with imperfect monitoring. *Econometrica*, 58(5):1041–1063, 1990.

D. Abreu, P. Milgrom, and D. Pearce. Information and timing in repeated partnerships. *Econometrica*, 59(6):1713–1733, 1991.

A. R. Admati and M. Perry. Joint projects without commitment. *Review of Economic Studies*, 58(2):259–276, 1991.

AI Futures Project. Ai 2040: Plan a, 2026. URL https://ai-2040.com/. Accessed: 2026-09-21.

D. Amodei. We must pace the frontier, Sept. 2026. URL https://darioamodei.com/post/we-must-pace-the-frontier.

M. Andreessen. Why AI will save the world. Andreessen Horowitz, June 2023. URL https://a16z.com/ai-will-save-the-world/.

Anthropic. Open-sourcing circuit-tracing tools, May 2025. URL https://www.anthropic.com/research/open-source-circuit-tracing.

S. Armstrong, N. Bostrom, and C. Shulman. Racing to the precipice: A model of artificial intelligence development. *AI & Society*, 31(2):201–206, 2016.

L. Aschenbrenner. Situational awareness: The decade ahead, June 2024. URL https://situational-awareness.ai/.

A. Askell, M. Brundage, and G. Hadfield. The role of cooperation in responsible ai development. *arXiv preprint arXiv:1907.04534*, 2019.

[17] A follower can also reduce risk by slowing down while behind. Higher η strengthens this insurance benefit which can make catching up less attractive. The point of (ES) is to rule this out.

J. Bergin and W. B. MacLeod. Continuous time repeated games. *International Economic Review*, 34(1):21–37, 1993.

M. Bivas, A. Daniilidis, and M. Quincampoix. Characterization of Filippov representable maps and Clarke subdifferentials. arXiv:2003.00436, 2020.

P. Bolton and C. Harris. Strategic experimentation. *Econometrica*, 67(2):349–374, 1999.

N. Bostrom, A. Dafoe, and C. Flynn. Public policy and superintelligent AI: A vector field approach. In S. M. Liao, editor, *Ethics of Artificial Intelligence*, pages 293–326. Oxford University Press, 2020.

E. Bueno de Mesquita, W. Dziuda, and M. Polborn. The agi race and existential risk. Technical report, National Bureau of Economic Research, 2026.

T. Cowen. A doomsday scenario for american AI, sep 2026.

A. Dafoe. AI governance: A research agenda. Technical report, Centre for the Governance of AI, Future of Humanity Institute, University of Oxford, 2018. URL https://cdn.governance.ai/GovAI-Research-Agenda.pdf.

P. Dasgupta and J. E. Stiglitz. Uncertainty, industrial structure, and the speed of R&D. *The Bell Journal of Economics*, 11(1):1–28, 1980.

O. Delaney, O. Guest, and Z. Williams. Mapping technical safety research at ai companies: A literature review and incentives analysis. *arXiv preprint arXiv:2409.07878*, 2024.

C. Doctorow. *The Reverse Centaur's Guide to Life After AI: How to Think About Artificial Intelligence—Before It's Too Late*. Tor Books, 2026.

N. Emery-Xu, A. Park, and R. Trager. Uncertainty, information, and risk in international technology races. *Journal of Conflict Resolution*, 68(10):2019–2047, 2024.

C. Fershtman and S. Nitzan. Dynamic voluntary provision of public goods. *European Economic Review*, 35(5):1057–1067, 1991.

A. F. Filippov. *Differential Equations with Discontinuous Righthand Sides*. Kluwer, Dordrecht, 1988.

D. Fudenberg and A. Koh. Racing to ruin. *arXiv preprint arXiv:2607.27638*, 2026.

D. Fudenberg and W. Olszewski. Repeated games with asynchronous monitoring of an imperfect signal. *Games and Economic Behavior*, 72(1):86–99, 2011.

D. Fudenberg and J. Tirole. Capital as a commitment: Strategic investment to deter mobility. *Journal of Economic Theory*, 31(2):227–250, 1983.

D. Fudenberg, R. Gilbert, J. Stiglitz, and J. Tirole. Preemption, leapfrogging and competition in patent races. *European Economic Review*, 22(1):3–31, 1983.

D. Fudenberg, D. K. Levine, and E. Maskin. The folk theorem with imperfect public information. *Econometrica*, 62(5):997–1039, 1994.

D. Fudenberg, Y. Ishii, and S. D. Kominers. Delayed-response strategies in repeated games with observation lags. *Journal of Economic Theory*, 150:487–514, 2014.

E. J. Green and R. H. Porter. Noncooperative collusion under imperfect price information. *Econometrica: Journal of the Econometric Society*, pages 87–100, 1984.

G. Haddad. Functional viability theorems for differential inclusions with memory. *Annales de l'Institut Henri Poincaré C, Analyse non linéaire*, 1(3):179–204, 1984.

C. Harris and J. Vickers. Perfect equilibrium in a model of a race. *The Review of Economic Studies*, 52(2):193–209, 1985.

C. Harris and J. Vickers. Racing with uncertainty. *The Review of Economic Studies*, 54(1): 1–21, 1987.

D. Hendrycks, E. Schmidt, and A. Wang. Superintelligence strategy: Expert version. *arXiv preprint arXiv:2503.05628*, 2025.

B. Horowitz. The fight over open source AI, 2026. URL https://a16z.com/the-fight-over-open-source-ai/. Accessed: 2026-09-13.

R. Isaacs. *Differential Games*. Wiley, New York, 1965.

N. Jaakkola and F. Wagener. Markov-perfect equilibria in differential games—with an application to climate policy. *Review of Economic Studies*, 2026. Advance article.

J. Kahn. Openai promised 20% of its computing power to combat the most dangerous kind of AI—but never delivered, sources say. Fortune, May 2024.

G. Keller, S. Rady, and M. Cripps. Strategic experimentation with exponential bandits. *Econometrica*, 73(1):39–68, 2005.

N. N. Krasovskii and A. I. Subbotin. *Game-Theoretical Control Problems*. Springer-Verlag, New York, 1988.

T. Lee and L. L. Wilde. Market structure and innovation: A reformulation. *The Quarterly Journal of Economics*, 94(2):429–436, 1980.

J. Leike and I. Sutskever. Introducing superalignment. OpenAI, July 2023. URL https://openai.com/index/introducing-superalignment/.

W. Leininger. Escalation and cooperation in conflict situations: The dollar auction revisited. *Journal of Conflict Resolution*, 33(2):231–254, 1989.

G. C. Loury. Market structure and innovation. *The quarterly journal of economics*, 93(3): 395–410, 1979.

L. M. Marx and S. A. Matthews. Dynamic voluntary contribution to a public project. *Review of Economic Studies*, 67(2):327–358, 2000.

E. Maskin and J. Tirole. Markov perfect equilibrium: I. Observable actions. *Journal of Economic Theory*, 100(2):191–219, 2001.

Model Evaluation and Threat Research. Frontier risk report (february to march 2026): A pilot assessment of misalignment risks from internal ai agent deployments. Technical report, METR, May 2026. URL https://metr.org/blog/2026-05-19-frontier-risk-report/. Published May 19, 2026.

I. Obara and V. Bhaskar. Reputation and information lag. Working paper, 2019.

Pacing the Frontier Signatories. Pacing the frontier: A statement from 1,386 employees of frontier AI companies, July 2026. URL https://www.pacingthefrontier.com/. Accessed: 2026-09-13.

J. F. Reinganum. A dynamic game of R and D: Patent protection and competitive behavior. *Econometrica*, 50(3):671–688, 1982.

T. C. Schelling. *The Strategy of Conflict*. Harvard University Press, Cambridge, MA, 1960.

L. K. Simon and M. B. Stinchcombe. Extensive form games in continuous time: Pure strategies. *Econometrica*, 57(5):1171–1214, 1989.

E. Stafford, R. F. Trager, and A. Dafoe. Safety not guaranteed: International races for risky technologies. Working paper, November, 2022. URL https://cdn.governance.ai/International_Races_for_Risky_Technologies_DRAFT_NOV_2022.pdf.

P. Trammell and L. Aschenbrenner. Existential risk and growth. Working paper, January 5, 2026. URL https://philiptrammell.com/static/Existential_Risk_and_Growth.pdf.

R. Zwetsloot and A. Dafoe. Thinking about risks from AI: Accidents, misuse and structure. Lawfare, Feb. 2019. URL https://www.lawfaremedia.org/article/thinking-about-risks-ai-accidents-misuse-and-structure. February 11.

A PROOFS

A.1 Histories and verification.

Admissible continuations after histories. We use the histories and pure strategies of Section 2.3, writing the prescribed speed pair at history h_t as $\sigma(h_t)$. To extend regularization to histories, we keep each history constant after its terminal date[18] and measure distance with the *stopped-history metric* $d(h_t, \widetilde{h}_{t'}) := |t - t'| + \sup_{s \geq 0} \|h_t(s \wedge t) - \widetilde{h}_{t'}(s \wedge t')\|$, where $\|\cdot\|$ is the Euclidean norm on $\mathbb{R}^3$. The supremum is finite because the two stopped histories are constant beyond their terminal dates. Define the joint Krasovskii set by

$$K[\sigma](h_t) := \bigcap_{\varepsilon > 0} \overline{\mathrm{co}}\{\sigma(\widetilde{h}_{t'}) : d(h_t, \widetilde{h}_{t'}) < \varepsilon\}. \tag{HK}$$

An admissible continuation after h_t extends that fixed history by an absolutely continuous

[18] That is, $h_t(t') = S_{t'}$ for $t' \leq t$ and $h_t(t') = S_t$ for $t' \geq t$.

state path satisfying, almost everywhere for $s > t$,

$$(\dot{A}_{1s}, \dot{A}_{2s}) \in K[\sigma](h_s) \quad \text{and} \quad \dot{X}_s = g. \tag{HA}$$

Write $\mathcal{P}(\sigma; h_t)$ for these continuations, where h_s includes the original history and the extension through date s.

Lemma 2 (Regularity on histories)**.** For every pure strategy profile σ, the correspondence $K[\sigma]$ defined by (HK) has nonempty compact convex values, is globally bounded, and is upper semicontinuous in the stopped-history metric.

Proof. For $\varepsilon > 0$, let $K_\varepsilon[\sigma](h)$ be the closed convex hull of the prescriptions at feasible histories within ε of h. It is a nonempty compact convex subset of $[0,1]^2$. As ε decreases, these sets are nested, and $K[\sigma](h) = \bigcap_{n\geq 1} K_{1/n}[\sigma](h)$. Thus $K[\sigma](h)$ is nonempty, compact, and convex. The inclusion $K[\sigma](h) \subseteq [0,1]^2$ gives global boundedness.

It remains to prove upper semicontinuity. Suppose $h^n \to h$ and $v^n \to v$, with $v^n \in K[\sigma](h^n)$. For any $\varepsilon > 0$, eventually $d(h^n, h) < \varepsilon/2$, and hence $K[\sigma](h^n) \subseteq K_\varepsilon[\sigma](h)$. Closedness then gives $v \in K_\varepsilon[\sigma](h)$. Since this holds for every $\varepsilon > 0$, $v \in K[\sigma](h)$, so the correspondence has a closed graph. A closed-graph correspondence whose values lie in the common compact set $[0,1]^2$ is upper semicontinuous. □

Lemma 3 (Paths after histories)**.** Fix a pure strategy σ and a feasible history h_t.

(i) *Existence.* There is an admissible continuation after h_t: $\mathcal{P}(\sigma; h_t) \neq \varnothing$.

(ii) *Restriction and concatenation.* If $S \in \mathcal{P}(\sigma; h_t)$, then, for every $s \geq t$, $S|_{[s,\infty)} \in \mathcal{P}(\sigma; h_s)$, where h_s is the full history generated by h_t and S. If $\widetilde{S} \in \mathcal{P}(\sigma; h_s)$, then the joined path

$$\widehat{S}_v := \begin{cases} S_v, & t \leq v \leq s, \\ \widetilde{S}_v, & v > s \end{cases}$$

satisfies $\widehat{S} \in \mathcal{P}(\sigma; h_t)$.

(iii) *Consistency with Markov paths.* If σ is Markov, then history regularization equals state regularization: $K[\sigma](h_s) = K[\sigma](S_s)$, and hence (HA) coincides with (A).

(iv) *Deviations against a Markov opponent.* If σ_j is Markov, then, for every history-dependent

deviation σ_i',

$$S \in \mathcal{P}((\sigma_i', \sigma_j); h_t) \quad \Longrightarrow \quad \begin{cases} \dot{A}_{is} = a_i(s), \\ \dot{A}_{js} \in K[\sigma_j](S_s), \\ \dot{X}_s = g, \end{cases} \qquad \text{for almost every } s > t,$$

for some measurable speed path $a_i : [t, \infty) \to [0, 1]$. These are the conditions in (D).

The proof is in the Supplemental Appendix.

For verifying that a profile is an MPE, we develop a test using paths instead of alternative strategies. Fix any measurable speed path $a_i : [0, \infty) \to [0, 1]$ and consider every absolutely continuous state path starting at $S_0 = S$ that satisfies, almost everywhere,

$$\dot{A}_{it} = a_i(t), \ \dot{A}_{jt} \in K[\sigma_j](S_t), \text{ and } \dot{X}_t = g. \tag{D}$$

Condition (D) restricts the resulting state path, given the chosen speed path and the opponent's strategy; it imposes no optimality requirement on a_i. By Lemma 3, every history-dependent deviation against a Markov opponent satisfies (D) with its realized own speed as the control.

Lemma 4 (Verification with several paths)**.** Suppose finite candidate values $\widehat{V}_i(S)$ satisfy the following conditions at every state S:

(i) *Attaining the candidate payoffs.* One common admissible path under σ gives each firm $\widehat{V}_i(S)$.

(ii) *No profitable deviation.* Every path from S satisfying (D), for every measurable control a_i, gives firm i at most $\widehat{V}_i(S)$.

Then $V_i(\sigma, S) = \widehat{V}_i(S)$, and σ is an MPE.

Proof. Condition (i) gives $V_i(\sigma, S) \geq \widehat{V}_i(S)$. Every joint admissible path satisfies (D) with $a_i = \dot{A}_i$, because projection of the joint Krasovskii set onto the opponent's coordinate lies in $K[\sigma_j](S)$. Condition (ii) therefore bounds every such path by $\widehat{V}_i(S)$, giving equality. Now fix a feasible history h_t and a deviation σ_i' that may depend on the full history. By Lemma 3, every continuation under (σ_i', σ_j) satisfies (D) from S_t. Condition (ii) bounds its payoff by $\widehat{V}_i(S_t)$. Because Markov continuation payoffs depend only on the current state, taking suprema gives $V_i((\sigma_i', \sigma_j), h_t) \leq \widehat{V}_i(S_t) = V_i(\sigma, h_t)$. □

Lemma 5 (Racing guarantee)**.** If a firm starts as a weak leader at distance D above safety and then races, it earns at least $R(D)$ per unit of its initial flow against every admissible response.

Proof. The firm remains a weak leader because its rival cannot develop faster. Its frontier is its initial capability plus elapsed time, so its hazard is exactly $\lambda(D + (1-g)t)$. Its flow grows at least at rate $\alpha - \beta$, since its rival's speed is at most one. Integrating gives $R(D)$. Prescribing speed one at every state or history fixes the realized speed at one almost everywhere under both (A) and (HA). □

A.2 Proof of Theorem 1.

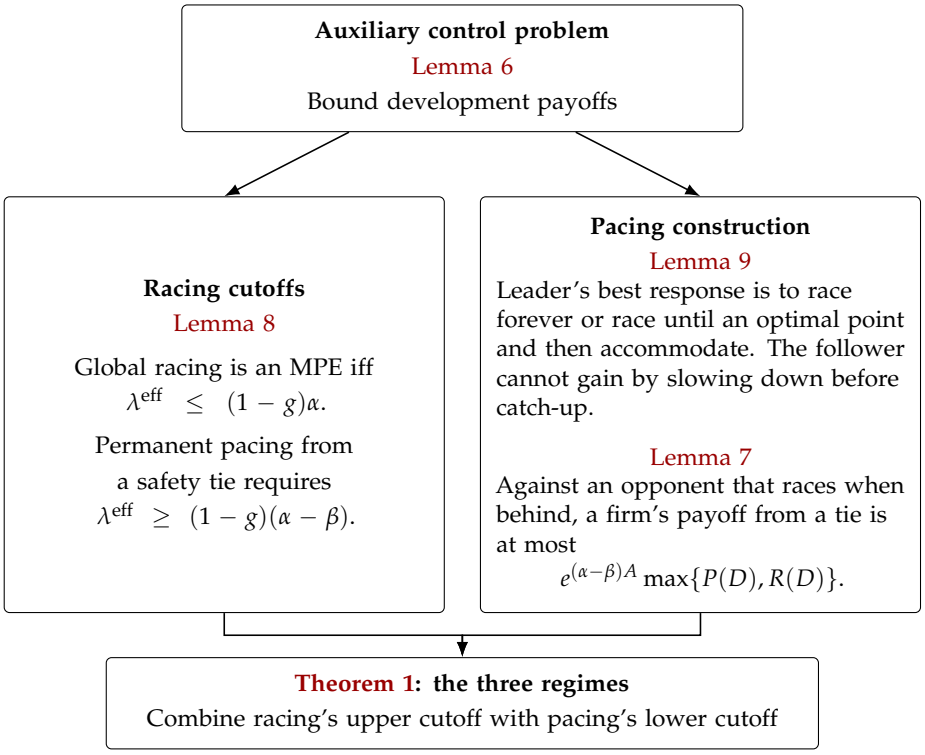

Figure 7: Proof structure for Theorem 1.

We start by introducing an auxiliary problem where a single decision-maker chooses a speed path $(a_t)_t$. Fix an initial distance from safety $D \in \mathbb{R}$, a flow growth coefficient $\eta > 0$, and a discount rate $\rho > 0$. Assume either $\rho > \eta$, or that the hazard has a finite ceiling $\bar{\lambda}$

and $\rho - g\eta > 0$ and $\rho - \eta + \bar{\lambda} > 0$. The auxiliary problem is:

$$\max_{\substack{a:[0,\infty)\to[0,1] \\ a \text{ measurable}}} \quad \int_0^\infty \exp\left\{-\rho t + \eta \int_0^t a(s)\ \mathrm{d}s - \int_0^t \lambda(D_s)\ \mathrm{d}s\right\} \mathrm{d}t$$
$$\text{subject to} \quad \dot{D}_t = a(t) - g \quad \text{for almost every } t \geq 0, D_0 = D.$$

The value of pacing on safety is $P_\eta(0) = 1/(\rho - g\eta)$. For this auxiliary problem, denote the racing and waiting-then-pacing values by R_η and P_η (fixing ρ):

$$R_\eta(D) = \int_0^\infty \exp\left\{-(\rho-\eta)t - \int_0^t \lambda(D+vs)\,\mathrm{d}s\right\}\mathrm{d}t,$$
$$P_\eta(D) = \int_0^{D/g} \exp\left\{-\rho t - \int_0^t \lambda(D-gs)\,\mathrm{d}s\right\}\mathrm{d}t + \exp\left\{-\rho D/g - \frac{1}{g}\int_0^D \lambda(s)\,\mathrm{d}s\right\} P_\eta(0).$$

The racing value is defined at every D; the waiting formula applies for $D \geq 0$.

Lemma 6 (When to race and when to wait)**.**

(a) If $R_\eta(0) \geq P_\eta(0)$, then perpetual racing is optimal for the auxiliary problem from every initial D. If the inequality is strict, every optimal path races almost everywhere. At equality, pacing is also optimal when starting on safety.

(b) If $R_\eta(0) < P_\eta(0)$, then the racing and waiting values have a unique positive crossing. The optimal value in the auxiliary problem above safety is $\max\{P_\eta(D), R_\eta(D)\}$, attained by the following choices:

Initial position	**Optimal continuation**
Below safety	Race to safety, then pace.
On safety	Pace.
Above safety but below the crossing	Pause until safety arrives, then pace.
At the crossing	Either pause until safety arrives and then pace, or race forever.
Above the crossing	Race forever.

Proof. We verify the candidate value and characterize the crossing.

Step 1: the payoff inequality. To show that F is an upper bound, we verify that, at every differentiable state $1 + (a-g)F'(D) + (\eta a - \rho - \lambda(D))F(D) \leq 0 \quad$ for every $a \in [0,1]$. If the LHS is nonpositive for every speed, integrating along any feasible path bounds that path's payoff by $F(D)$. For $D > 0$, the equations for racing and waiting are $vR'_\eta =$

$(\rho - \eta + \lambda)R_\eta - 1 \quad \text{and} \quad gP'_\eta = 1 - (\rho + \lambda)P_\eta$. Both values are nonincreasing, and P_η is strictly decreasing on $(0, \infty)$. For waiting, increasing D reduces development and weakly increases exposure at every date, with a strict loss over a positive interval; for racing, monotonicity follows directly from the integral. For each t, the integrand of $R_\eta(D)$ is log-convex in D because λ is concave. Hence, by Hölder's inequality, $R_\eta(\theta D_1 + (1-\theta)D_2) \leq R_\eta(D_1)^\theta R_\eta(D_2)^{1-\theta}$. Thus $\log R_\eta$ is convex, so R'_η / R_η is nondecreasing.

Step 2: when racing is optimal everywhere. When $R_\eta(0) \geq P_\eta(0)$, at $D = 0+$, $\dfrac{(\rho - g\eta + \lambda_0)R_\eta(0) - 1}{v} \geq 0$. Log convexity then implies $R'_\eta + \eta R_\eta \geq 0$ for all $D > 0$. letting $z = \exp\{(\rho - \eta)D/v\}$, the speed coefficient is $\dfrac{\eta(1-z)}{\rho - \eta} + \dfrac{z}{v}\left[(\rho - g\eta)R_\eta(0) - 1\right] \geq 0$. The first term is $-\eta D/v$ when $\rho = \eta$. On safety the pacing residual is $1 - (\rho - g\eta)R_\eta(0) \leq 0$. These inequalities establish that perpetual racing is optimal from every initial distance D. If $R_\eta(0) = P_\eta(0)$, pacing on safety is also optimal.

Step 3: when waiting and racing have a positive crossing. Suppose $R_\eta(0) < P_\eta(0)$. Above safety, we show that the value is $\max\{P_\eta(D), R_\eta(D)\}$, with a unique positive crossing. Let $Q_\eta = P'_\eta + \eta P_\eta$ be the speed coefficient for waiting. Differentiation gives $gQ'_\eta = -(\rho + \lambda)Q_\eta + \eta - \lambda' P_\eta$ and $Q_\eta(0+) = -\lambda_0 P_\eta(0)/g \leq 0$.

The term $\eta - \lambda' P_\eta$ is nondecreasing: concavity makes λ' nonnegative and nonincreasing, while P_η decreases. It converges to η. Hence any zero of Q_η is crossed from negative to positive, and Q_η eventually becomes positive: for large enough D, $\eta - \lambda'(D)P_\eta(D) \geq \eta/2$, so whenever $Q_\eta(D) \leq 0$ we have $Q'_\eta(D) \geq \eta/(2g)$. Thus Q_η crosses zero exactly once.

Define $M_\eta(D) = \exp\{[(\rho - \eta)D + \int_0^D \lambda(s)\,\mathrm{d}s]/v\}$. Subtracting the two control equations yields $\left((P_\eta - R_\eta)/M_\eta\right)' = Q_\eta/(vM_\eta)$. Both P_η and R_η are bounded, while $M_\eta(D) \to \infty$: this is immediate if $\rho > \eta$; otherwise $\lambda(D) \uparrow \bar{\lambda}$ and $\log M_\eta(D)/D \to (\rho - \eta + \bar{\lambda})/v > 0$. Hence $(P_\eta - R_\eta)/M_\eta \to 0$. Since its derivative has the sign of Q_η, the ratio first decreases and then increases toward zero. It starts positive, so it crosses zero exactly once, before Q_η does; therefore $Q_\eta < 0$ at the crossing. *Step 4: integrate the bound along every path.* The candidate is locally Lipschitz and continuously differentiable except at safety and, when present, the positive crossing. Its composition with any absolutely continuous state path is absolutely continuous. Away from those states, the chain rule gives the verified payoff inequality. At almost every date when the path is at either of these states, its state derivative is zero, so its speed is $a = g$. On safety, the safety residual applies with hazard zero. At the positive crossing, the residual is $1 - (\rho - g\eta + \lambda)P_\eta = gQ_\eta \leq 0$, so the inequality holds almost everywhere along every path. Let $Z_t = \exp\{-\rho t + \eta \int_0^t a_s\,\mathrm{d}s - \int_0^t \lambda(D_s)\,\mathrm{d}s\}$ denote the integrand in the objective, and note that $\int_0^T Z_t\,\mathrm{d}t + Z_T F(D_T) \leq$

$F(D_0)$. Since $F \geq 0$, letting $T \to \infty$ bounds every path and the stated policies attain the bound. When $R_\eta(0) > P_\eta(0)$, the racing speed coefficient is strictly positive off safety and the safety residual is strictly negative. Any path attaining the bound must therefore race almost everywhere. □

Lemma 7. If a Markov opponent races whenever it is strictly behind, then a firm starting weakly behind remains weakly behind on every path satisfying (D). Its continuation payoff is at most $e^{cA_i} \max\{P(A_i - X), R(A_i - X)\}$.

Proof. Whenever the deviator is strictly ahead, its opponent's speed is one even after scalar regularization. Thus the positive part of their capability difference has nonpositive derivative almost everywhere and starts at zero, so it remains zero. Along every continuation, the deviator's flow is at most e^{cA_i} and its hazard is at least $\lambda(A_i - X)$. Applying Lemma 6 with $(\eta, \rho) = (c, r)$ to its own capability path gives the bound. □

For the leader facing a racing follower, write $B = r + \beta - g\alpha$ for the discount rate net of profit growth while the leader paces. It is positive, and $B > r - gc = 1/P(0)$.

Lemma 8 (The racing cutoffs)**.** Global racing is an MPE if and only if $R(0) \geq 1/B$. If $R(0) > P(0)$, no MPE outcome paces forever from a tie on safety.

Proof. We rule out pacing when its payoff is too low, then check when racing is a best response to a racing rival. If $R(0) > P(0)$, the racing guarantee in Lemma 5 gives a profitable deviation from perpetual pacing at a tie on safety.

Apply Lemma 6 with $(\eta, \rho) = (\alpha, r + \beta)$. The racing value is the same R and the safe value is $1/B$. This is the leader's problem when its follower never catches up. If $R(0) \geq 1/B$, global racing is an MPE. Against a rival that always races, an initially leading firm's payoff per unit of its initial flow is bounded by this auxiliary problem: its flow grows at rate $\alpha a - \beta$, and actual hazard is at least $\lambda(A_i - X)$. Racing attains the bound $R(D)$. An initially trailing firm cannot catch a racing rival; full speed maximizes its flow without changing the frontier or hazard. At a tie the same bound or pointwise comparison applies.

Conversely, if $R(0) < 1/B$, start with a strict leader on safety. Against global racing it gains by pacing for $0 < h < G/v$ and then resuming racing. Per unit of the leader's initial flow, the gain is $(1 - e^{-Bh})\big[1/B - R(0)\big] > 0$. The pacing interval incurs zero hazard, including when $\lambda_0 > 0$. Thus global racing is an MPE exactly when $\lambda^{\text{eff}} \leq v\alpha$. □

Lemma 9 (Pacing passes the verification test)**.** If $R(0) \leq P(0)$, then σ^P is an MPE with finite payoffs and a common attaining path from every state.

Proof sketch. The scalar-control lemma bounds deviations at ties and along racing segments. At a strict gap, the leader chooses its earliest optimal accommodation point; the stopping formula implies that accommodation remains optimal once begun, and the associated value satisfies the verification inequalities for both racing and waiting. A follower cannot improve by slowing down before catch-up, and after catch-up Lemma 7 applies. The prescribed common path attains both bounds, so Lemma 4 establishes the result. The full proof is in the Supplemental Appendix. □

Proof of Theorem 1. By Lemma 8, global racing is an MPE exactly when $\lambda^{\text{eff}} \leq (1-g)\alpha$, and no equilibrium outcome paces forever from a tie on safety when $\lambda^{\text{eff}} < (1-g)(\alpha-\beta)$. By Lemma 9, pacing is an MPE whenever $\lambda^{\text{eff}} \geq (1-g)(\alpha-\beta)$, including equality. The two lemmas give the three regimes and paths attaining both payoffs. From a tie on safety, the pacing path incurs zero hazard and gives each firm $P(0)$. □

A.3 Proof of Theorem 2.

Proof of Theorem 2(a). Part (a) says that every SPE outcome from a safety tie is perpetual joint racing. We prove this in two steps. First, we show that along any feasible continuation, the firms' combined payoff is at most $2R(0)$, their combined payoff under joint racing. Second, in any SPE each firm can guarantee itself at least $R(0)$ by racing, so an SPE outcome must attain the aggregate bound $2R(0)$. We show that equality is possible only if the firms remain tied and race almost everywhere.

Put $c = \alpha - \beta$, and subtract A_0 from both capabilities and the safety threshold. This leaves risk unchanged, so we may assume $A_{1,0} = A_{2,0} = X_0 = 0$. Then the hypotheses of part (a) imply $R(0) > P(0)$ and $2\alpha\beta R(0) < c$. Moreover, $R(D) \leq R(0)$ and, by Lemma 6, $R'(D) + cR(D) > 0$ for $D > 0$. Define $W(G) := 2P(0)(\alpha e^{\beta G} + \beta e^{-\alpha G})/(\alpha+\beta)$ and $\widehat{V}(S) := [R(\max\{D,0\})/P(0)]e^{c\max\{A,X\}}W(\max\{A,X\} - \min_i A_i)$. This function is nonnegative, locally Lipschitz, and equals $2R(0)$ at the initial tie. Along any feasible path, define $\mathcal{E}_t := \pi_1 + \pi_2 + \frac{\mathrm{d}}{\mathrm{d}t}\widehat{V}(S_t) - (r + \lambda(D_t))\widehat{V}(S_t)$. This is the combined flow payoff plus the drift of $\widehat{V}$, net of discounting and disaster risk. Showing that it is nonpositive gives the desired aggregate payoff bound.

On or below safety, put $G = X - \min_i A_i$. Raising the leader's capability to X increases total flow, so $\pi_1 + \pi_2 \leq e^{cX}(e^{\beta G} + e^{-\alpha G})$. Also, $\dot{G} \leq g$, $W' \geq 0$, and $W/P(0) - (e^{\beta G} + e^{-\alpha G}) - gW' = [c - 2\alpha\beta g P(0)](e^{\beta G} - e^{-\alpha G})/(\alpha+\beta) \geq 0$. Since $1/P(0) = r - gc$, $\mathcal{E}_t$ is at most $(1 - R(0)/P(0))(\pi_1 + \pi_2) < 0$. This also applies almost everywhere on the safety boundary, where $\dot{A} = g$.

Above safety, write π_L and π_F for the leader's and follower's flows. Here $\widehat{V} = 2R(D)(\alpha\pi_L + \beta\pi_F)/(\alpha+\beta)$. At a strict lead, $\mathcal{E}_t$ is linear in speeds, with positive coefficient on the leader's speed and nonpositive coefficient on the follower's. Indeed, these coefficients are respectively $2[(R'+cR)(\alpha\pi_L+\beta\pi_F)+\alpha\beta R(\pi_L-\pi_F)]/(\alpha+\beta)$ and $-2\alpha\beta R(\pi_L-\pi_F)/(\alpha+\beta)$, so it is maximized when the leader races and the follower pauses. Substituting $(1-g)R' = (r-c+\lambda)R-1$ bounds it by $[2\alpha\beta R(D)-c](\pi_L-\pi_F)/(\alpha+\beta) < 0$. On an unsafe tie, speeds agree almost everywhere; at common speed a, the same expression equals $-2e^{cA}(1-a)(R'+cR) \le 0$, with equality only when $a=1$. Multiply $\mathcal{E}_t$ by $e^{-rt-\Lambda_t}$ and integrate. Since the terminal value is nonnegative, letting the horizon tend to infinity gives total payoff at most $2R(0)$, with strict inequality if the residual inequality is strict on a set of positive measure. By **??**, each firm can guarantee $R(0)$ from the initial tie, so every SPE must attain this bound. Equality requires zero time on or below safety and no strict capability lead. Thus the firms remain tied and race almost everywhere, and every SPE outcome is perpetual joint racing. □

The payoff guarantee for part (b). For a pause of length $\varepsilon > 0$, define the maximum waiting time T and the guaranteed safe interval length Δ by $T := \frac{1-g}{g}\varepsilon$ and $\Delta := \frac{2g-1}{g(1-g)}\varepsilon$. The first step in Figure 6 illustrates the rival's safe option when it races throughout the initial pause. The payoff bound also covers every other response to that pause.

After the pause, the rival can wait to safety if it is above it, or race to safety if it is below it, and then pace. Along this comparison path, its excess above safety is at most $((1-g)\varepsilon - gs)_+$. Write $\Gamma_\varepsilon(s)$ for the resulting exposure bound, measuring elapsed time s from the end of the pause: $\Gamma_\varepsilon(s) := \frac{1}{g}\int_{((1-g)\varepsilon-gs)_+}^{(1-g)\varepsilon} \lambda(D)\,\mathrm{d}D$, $s \ge 0$. The exposure bound is constant after T. If the rival races throughout the pause, then waits, paces, and then races at $T+\Delta$, its continuation payoff is at least

$$\widetilde{L}_2(\varepsilon) = e^{\alpha\varepsilon}\int_0^T e^{-(r+\beta)s-\Gamma_\varepsilon(s)}\,\mathrm{d}s + e^{\alpha g\varepsilon-\Gamma_\varepsilon(T)}\int_T^{T+\Delta} e^{-(r+\beta-\alpha g)s}\,\mathrm{d}s + e^{-\Gamma_\varepsilon(T)-(r-\alpha+\beta)(T+\Delta)}R(0).$$

Define

$$\begin{aligned} \underline{V}_1 &:= \max\left\{R(0),\ \sup_{\varepsilon>0}\widetilde{L}_1(\varepsilon)\right\} \\ \widetilde{L}_1(\varepsilon) &:= \int_0^\varepsilon e^{-(r+\beta)s-\Lambda_s^R}\,\mathrm{d}s + \frac{r-\alpha}{r+\beta}e^{-(r+\alpha+\beta)\varepsilon-\Lambda_\varepsilon^R}\widetilde{L}_2(\varepsilon). \end{aligned} \tag{RP}$$

The terminal guarantee in $\widetilde{L}_2$ is $R(0)$. Using a continuation guarantee m instead gives the pause bound $\widetilde{L}_1(\varepsilon) + q(\varepsilon)[m-R(0)]$, where

$$q(\varepsilon) := \frac{r-\alpha}{r+\beta}e^{-(r+\alpha+\beta)\varepsilon-\Lambda_\varepsilon^R-\Gamma_\varepsilon(T)-(r-\alpha+\beta)(T+\Delta)}.$$

The racing guarantee gives each firm a normalized continuation payoff of at least $R(0)$ in every SPE after every feasible history ending at a safety tie. We show that, for every $\varepsilon > 0$, pausing for ε and then resuming equilibrium also guarantees at least $\widetilde{L}_1(\varepsilon)$. More generally, if $m \in [R(0), P(0)]$ is an established normalized continuation-payoff guarantee after every safety tie, the same argument yields the guarantee $\widetilde{L}_1(\varepsilon) + q(\varepsilon)[m - R(0)]$. Starting from $R(0)$, we iterate this implication below.

Proof of the general pause bound. Put $c = \alpha - \beta$ and $k = (r - \alpha)/(r + \beta)$. Let $m \in [R(0), P(0)]$ be a normalized payoff guarantee after every history ending at a safety tie. It also applies to a leader on safety: pacing until the first tie at date t and then resuming equilibrium gives at least $P(0) + e^{-t/P(0)}[m - P(0)] \geq m$. If no tie occurs, pacing gives at least $P(0)$. When $m = R(0) > P(0)$, the same guarantee follows instead from Lemma 5.

At any history ending at S, speed limits imply $\pi_i(S_s) \leq \pi_i(S)e^{\alpha s}$ and $\pi_j(S_s) \geq \pi_j(S)e^{-\beta s}$. For $0 < d < d'$, common survival gives $d \int_0^\infty e^{-ds - \Lambda_s}\, ds \leq d' \int_0^\infty e^{-d's - \Lambda_s}\, ds$, because the two sides equal $\mathbb{E}[1 - e^{-d\tau_D}]$ and $\mathbb{E}[1 - e^{-d'\tau_D}]$. Taking $d = r - \alpha$ and $d' = r + \beta$, then taking suprema over the same admissible paths, yields

$$V_j(\sigma, h) \geq k \frac{\pi_j(S)}{\pi_i(S)} V_i(\sigma, h). \tag{PR}$$

Normalize the initial safety tie to zero and let firm 1 pause until ε. Every response ends at $S_\varepsilon = (0, x, g\varepsilon)$ with $x \in [0, \varepsilon]$. Restart elapsed time there and put $H := T + \Delta = g\varepsilon/(1 - g)$. Firm 2 can follow $A^x_{2s} = \max\{x, \min\{x + s, g\varepsilon + gs\}\}$ until H, then resume equilibrium. This path waits or races to safety and then paces. Since $A_{1s} \leq s \leq A^x_{2s}$ for $s \leq H$, firm 2 remains weakly ahead. Exposure is at most $\Gamma_\varepsilon(s)$, and at H firm 2 is on safety with capability H, where its continuation guarantee is $e^{cH}m$.

The comparison path with $x = \varepsilon$ gives firm 2 at least $B_\varepsilon(m) := \widetilde{L}_2(\varepsilon) + e^{-\Gamma_\varepsilon(T) - (r-c)H}[m - R(0)]$. Moreover, $0 \leq A^\varepsilon_{2s} - A^x_{2s} \leq \varepsilon - x$, so $e^{-(\alpha+\beta)x}e^{\alpha A^x_{2s} - \beta s} \geq e^{-(\alpha+\beta)\varepsilon}e^{\alpha A^\varepsilon_{2s} - \beta s}$. The terminal guarantee is independent of x and satisfies the same comparison, so equation (PR) implies $V_1(\sigma, h_\varepsilon) \geq ke^{-(\alpha+\beta)\varepsilon}B_\varepsilon(m)$ after every response to the pause.

During the pause, firm 1's flow is at least $e^{-\beta s}$ and exposure is at most Λ^R_s. Pausing and then resuming equilibrium consequently gives at least $\int_0^\varepsilon e^{-(r+\beta)s - \Lambda^R_s}\, ds + ke^{-(r+\alpha+\beta)\varepsilon - \Lambda^R_\varepsilon}B_\varepsilon(m) = \widetilde{L}_1(\varepsilon) + q(\varepsilon)[m - R(0)]$. By Lemma 3, each prescribed segment can be joined to a continuation approaching the relevant firm's assigned payoff. The initial pacing argument uses the same concatenation at a first tie along any admissible response, and the argument applies symmetrically after every safety tie. □

Iterating the guarantee. Starting from the racing guarantee, define

$$\begin{aligned} \underline{V}_0 &:= R(0), \\ \underline{V}_{n+1} &:= \max\left\{R(0), \sup_{\varepsilon>0}\left[\widetilde{L}_1(\varepsilon) + q(\varepsilon)(\underline{V}_n - R(0))\right]\right\}. \end{aligned} \tag{RF}$$

The first iterate is (RP). The limit has the formula

$$\underline{V}_\infty = \max\left\{R(0), \sup_{\varepsilon>0}\frac{\widetilde{L}_1(\varepsilon) - q(\varepsilon)R(0)}{1-q(\varepsilon)}\right\}. \tag{LF}$$

Lemma 10 (Recursive guarantees tighten the risk bound)**.** If the hypotheses of Theorem 2(b) hold, then the floors $\underline{V}_n$ increase weakly to a finite limit $\underline{V}_\infty$, given by (LF). After every feasible history ending at a safety tie, every SPE gives each firm normalized payoff at least $\underline{V}_\infty \geq \underline{V}_1$. The inequality is strict whenever $\underline{V}_1 > R(0)$. If $R(0) \leq P(0)$, then $\underline{V}_\infty \leq P(0)$ and

$$0 \leq \underline{V}_\infty - \underline{V}_n \leq \left(\frac{r-\alpha}{r+\beta}\right)^n [P(0) - R(0)].$$

If $R(0) > P(0)$, then $\underline{V}_n = \underline{V}_\infty = \underline{V}_1 = R(0)$ for every n.

Proof of Lemma 10. Step 1: Establish the payoff lower bounds. If $R(0) \leq P(0)$, the pause bound maps any valid floor $\underline{V}_n \leq P(0)$ into the valid floor $\underline{V}_{n+1}$. The pacing MPE in Theorem 1 is an SPE giving exactly $P(0)$ from a safety tie, so $\underline{V}_{n+1} \leq P(0)$. Induction from $\underline{V}_0 = R(0)$ proves the guarantee after every feasible history ending at a safety tie. If $R(0) > P(0)$, the racing MPE gives exactly $R(0)$ at a safety tie. The first pause floor must therefore equal $R(0)$, and all subsequent iterates remain there.

Step 2: Identify the limit and the improvement. If $R(0) > P(0)$, Step 1 establishes all the remaining claims. Suppose henceforth that $R(0) \leq P(0)$, and write $c := (r-\alpha)/(r+\beta) \in (0,1)$. Define

$$T(x) := \max\left\{R(0), \sup_{\varepsilon>0}\left[\widetilde{L}_1(\varepsilon) + q(\varepsilon)(x - R(0))\right]\right\}.$$

Since $0 < q(\varepsilon) \leq c$, the map T is nondecreasing and satisfies $|T(x) - T(y)| \leq c|x-y|$. As $\underline{V}_1 \geq \underline{V}_0$, induction gives $\underline{V}_{n+1} \geq \underline{V}_n$ for every n. Step 1 bounds this sequence above by $P(0)$, so it converges to some $\underline{V}_\infty \leq P(0)$. Continuity of T gives $T(\underline{V}_\infty) = \underline{V}_\infty$, so $0 \leq \underline{V}_\infty - \underline{V}_n \leq c^n(\underline{V}_\infty - R(0)) \leq c^n[P(0) - R(0)]$. At every feasible history ending at a safety tie, each firm's normalized payoff in every SPE is at least $\underline{V}_n$ for every n, and hence at least $\underline{V}_\infty$.

To establish (LF), set

$$M := \max\left\{R(0), \sup_{\varepsilon>0} \frac{\widetilde{L}_1(\varepsilon) - q(\varepsilon)R(0)}{1-q(\varepsilon)}\right\}.$$

Because $q(\varepsilon) < 1$, the fixed-point identity implies $\underline{V}_\infty \geq M$, so M is finite. The definition of M also gives $T(M) \leq M$. Since $\underline{V}_0 = R(0) \leq M$ and T is nondecreasing, induction yields $\underline{V}_n \leq M$ for every n. Taking limits gives $\underline{V}_\infty \leq M$, proving (LF).

Suppose now that $\underline{V}_1 > R(0)$. The function $\widetilde{L}_1$ is continuous on $(0,\infty)$, with $\lim_{\varepsilon\downarrow 0} \widetilde{L}_1(\varepsilon) = cR(0) < R(0)$ and $\lim_{\varepsilon\to\infty} \widetilde{L}_1(\varepsilon) = \int_0^\infty e^{-(r+\beta)s-\Lambda_s^R}\,\mathrm{d}s < R(0)$. For the second limit, the integral term converges by dominated convergence. Ignoring hazard bounds the rival's comparison payoff by $e^{\alpha\varepsilon}/(r-\alpha) + R(0)$. Thus, for some constant C independent of ε, the absolute value of the continuation term is at most

$$Ce^{-(r+\alpha+\beta)\varepsilon}\left[\frac{e^{\alpha\varepsilon}}{r-\alpha} + R(0)\right],$$

which converges to zero. The strict inequality follows from

$$\int_0^\infty e^{-(r+\beta)s-\Lambda_s^R}\,\mathrm{d}s < \int_0^\infty e^{-(r-\alpha+\beta)s-\Lambda_s^R}\,\mathrm{d}s = R(0).$$

Choose $\varepsilon_0 > 0$ such that $\widetilde{L}_1(\varepsilon_0) > R(0)$. The endpoint limits imply that there exist $0 < a < \varepsilon_0 < b < \infty$ such that $\widetilde{L}_1(\varepsilon) < \widetilde{L}_1(\varepsilon_0)$ whenever $\varepsilon \notin [a,b]$. Continuity therefore implies that $\widetilde{L}_1$ attains its global maximum at some $\varepsilon^* \in [a,b]$, with $\widetilde{L}_1(\varepsilon^*) = \underline{V}_1$. It follows that $\underline{V}_2 \geq \widetilde{L}_1(\varepsilon^*) + q(\varepsilon^*)(\underline{V}_1 - R(0)) > \underline{V}_1$, so $\underline{V}_\infty > \underline{V}_1$. If instead $\underline{V}_1 = R(0)$, then $T(R(0)) = R(0)$, so every iterate equals $R(0)$. □

Proof of Theorem 2(b). Fix an SPE outcome from a safety tie, and normalize the initial common capability to zero. At date s, both capabilities lie in $[0,s]$. Their summed flow is convex, so it is maximized at a corner of this square. The maximizing corner has both capabilities equal to s, or one equal to s and the other zero. Let $F(s) := \max\left\{e^{(\alpha-\beta)s}, \frac{e^{\alpha s}+e^{-\beta s}}{2}\right\}$. This gives the flow ceiling $2F(s)$. Since an SPE outcome attains both firms' assigned payoffs on the same path, Lemma 10 implies

$$\underline{V}_\infty \leq \frac{V_1(\sigma,h_0) + V_2(\sigma,h_0)}{2} \leq \int_0^\infty e^{-rs-\Lambda_s} F(s)\,\mathrm{d}s.$$

Fix a finite date t and write $z = \Lambda_t$. Since $D_s \leq (1-g)s$, cumulative exposure by any date s is at most Λ_s^R. Therefore, for $s \leq t$, the exposure that must already have occurred

by s is at least $\max\{0,\ z-(\Lambda_t^R-\Lambda_s^R)\}$. For $s\geq t$, cumulative exposure is at least z. To bound payoffs from above, we use the most favorable exposure comparison consistent with these bounds: it postpones exposure as long as possible before t and assumes no further exposure after t.[19] Thus $\underline{V}_\infty \leq \overline{V}(t,z)$, where

$$\overline{V}(t,z) := \int_0^t e^{-rs}F(s)\min\{1, e^{-z+\Lambda_t^R-\Lambda_s^R}\}\ \mathrm{d}s + e^{-z}\int_t^\infty e^{-rs}F(s)\ \mathrm{d}s. \tag{UC}$$

For fixed t, $\overline{V}(t,z)$ is continuous and strictly decreasing in z: the first term is nonincreasing, while the positive tail is strictly decreasing. Define $\overline{\Lambda}(t) := \max\left\{z\in[0,\Lambda_t^R] : \overline{V}(t,z)\geq \underline{V}_\infty\right\}$. The set is nonempty. At least one of the racing and pacing profiles is an SPE, and applying the preceding payoff bound to an attaining outcome of that equilibrium produces some $z\in[0,\Lambda_t^R]$ for which $\overline{V}(t,z)\geq\underline{V}_\infty$. Since $\overline{V}(t,\Lambda_t)\geq\underline{V}_\infty$, we have $\Lambda_t\leq\overline{\Lambda}(t)$, so $\Pr(\tau_D\leq t) = 1-e^{-\Lambda_t}\leq 1-e^{-\overline{\Lambda}(t)}$. □

A.4 Proof of Proposition 3.

Proof. We bound equilibrium payoffs from below, then show that racing attains the bound, and use it to sustain pacing.

Step 1: the payoff guarantee. Fix a feasible history h_t ending at state $S=(A_1,A_2,X)$. If firm i races thereafter, its rival develops no faster than under mutual racing, so firm i's flow profit is at least as high and hazard is no greater, so $V_i(\sigma,h_t)\geq\pi_i(A_i,A_j)R_i(S)$ in every SPE. If pacing is an SPE outcome, then after every history on the pacing path the prescribed continuation must attain its assigned payoff; otherwise, Lemma 3 would allow a better continuation to be concatenated to the original path. Thus, at every safety tie $x\geq A_0$, the racing guarantee implies $P_i(x)\geq R_i(x,x,x)$.

Step 2: racing attains the bound. Restart time at any state $S=(A_1,A_2,X)$ and hold the rival to perpetual racing. For any feasible own path, let $\delta_t := A_i+t-A_{it}$ be its capability shortfall relative to racing. Since speed is at most one, this shortfall is nonnegative and nondecreasing, so $D+(1-g)s-\delta_t\leq D_s\leq D+(1-g)s$ where $0\leq s\leq t$. Writing Λ_t^R and Λ_t for cumulative exposure under racing and the alternative, monotonicity of hazard and a change of variables give

$$\begin{aligned}0\leq\Lambda_t^R-\Lambda_t &\leq \int_0^t\left[\lambda(D+(1-g)s)-\lambda(D+(1-g)s-\delta_t)\right]\ \mathrm{d}s\\ &= \frac{1}{1-g}\left[\int_{D+(1-g)t-\delta_t}^{D+(1-g)t}\lambda(z)\ \mathrm{d}z-\int_{D-\delta_t}^{D}\lambda(z)\ \mathrm{d}z\right]\leq\frac{\bar\lambda}{1-g}\delta_t.\end{aligned}$$

[19] This comparison relaxes the feasible-path restrictions and need not itself be feasible.

Racing raises log flow profit by at least $\alpha\delta_t$, so

$$\frac{e^{-\Lambda_t^R}\pi_i(A_i+t,A_j+t)}{e^{-\Lambda_t}\pi_i(A_{it},A_j+t)} \geq \exp\left[\left(\alpha-\frac{\bar{\lambda}}{1-g}\right)\delta_t\right] \geq 1.$$

Discounting and integrating shows that no feasible own path beats racing. A constant-racing rival retains speed one under joint regularization, so the bound covers every history-dependent deviation. Mutual racing is therefore an SPE whose unique path attains the lower bound for both firms at every state.

Step 3: construct pacing equilibrium. Let

$$\sigma_i(h_t) = \begin{cases} g, & A_{1s} = A_{2s} = X_s = A_0 + gs \text{ for every } s \in [0,t], \\ 1, & \text{otherwise.} \end{cases}$$

This strategy is progressively measurable, and the set of histories following the prescribed path is closed in the stopped-history metric. Once a history departs, all sufficiently nearby histories also record a departure. The nondeviating firm's speed is therefore 1 under regularization after that history and every extension. On the prescribed path, pacing is admissible; we now bound the payoff from deviations.

Fix a history h_t on the prescribed path, any deviation σ_i', and any admissible continuation under (σ_i', σ_j). A continuation that never departs gives the pacing payoff. Otherwise, let $T \geq t$ be the infimum of departure dates. Continuity gives $A_{1T} = A_{2T} = X_T = A_0 + gT$. Every history strictly after T records a departure, so the rival races almost everywhere thereafter. Step 2 bounds the deviator's payoff from T by $\pi_i(X_T, X_T)R_i(X_T, X_T, X_T)$. Before T, profits coincide with pacing and hazard is zero. Its gain over pacing is thus at most $e^{-r(T-t)}\pi_i(X_T, X_T)\big[R_i(X_T, X_T, X_T) - P_i(X_T)\big] \leq 0$. Taking suprema rules out every strategy deviation. The same bound with $\sigma_i' = \sigma_i$ bounds the assigned payoff, while joint pacing attains it for both firms on the same path. After every other history, the profile prescribes perpetual racing, which Step 2 verifies as an equilibrium. Thus σ is an SPE with permanent joint pacing as an outcome. □

SUPPLEMENTAL APPENDIX: NOT FOR PUBLICATION

Proof of Lemma 1.

Proof. Fix a Markov opponent σ_j.

Step 1: Bound the speeds available to deviations. From Lemma 3, the opponent's speed along any admissible deviation path lies in $K[\sigma_j](S_t)$, the set of speeds allowed by its Markov strategy under regularization. Thus $\dot{A}_{it} \in [0,1]$ and $\dot{A}_{jt} \in K[\sigma_j](S_t)$ almost everywhere.

Step 2:Constructing the rectangle of admissible speeds. We seek a measurable Markov strategy σ_i^* and a countable exceptional set of states E such that, ordering speeds as firm i then firm j,

$$K[(\sigma_i^*, \sigma_j)](S) = [0,1] \times K[\sigma_j](S), \qquad S \notin E. \tag{RECT}$$

The idea is to place prescriptions of both 0 and 1 arbitrarily close to the opponent's graph. Regularization then fills in every intermediate own speed, while preserving the opponent's allowed speeds. Consider the opponent's graph $\{(S, \sigma_j(S)) : S \in \mathcal{S}\}$ in state–speed space. List all open balls with rational centers and rational radii as $U_1, U_2, \ldots$. These sets form a *countable base*: every neighborhood of any point contains a listed ball that also contains that point.

At stage n, inspect the graph points inside the ball U_n. If there are only finitely many, skip that ball. Otherwise, there are infinitely many distinct states to choose from, because the opponent prescribes just one speed at each state. Earlier stages have used only finitely many states: $\#\{\text{states chosen before stage } n\} \leq 2(n-1) < \infty$. We can therefore choose two graph points in U_n whose states have not been chosen earlier. Put the first state in E_0, the set receiving prescription zero, and the second in E_1, the set receiving prescription one. After processing all balls, the sets E_0 and E_1 are countable and disjoint. Let $\sigma_i^*(S) = \mathbf{1}\{S \in E_1\}$ be the Markov strategy that prescribes one on E_1 and zero everywhere else, including on E_0. It is measurable because E_1 is countable.

For every accumulation point (S, b) of the opponent's graph and each own prescription $a \in \{0,1\}$, there are states $S_n \in E_a$ such that $S_n \longrightarrow S$ and $\big(\sigma_i^*(S_n), \sigma_j(S_n)\big) \longrightarrow (a, b)$. Indeed, every neighborhood of an accumulation point contains infinitely many graph points, so a listed ball containing that point and contained in the neighborhood receives a state of each type. Take E to be the state coordinates of isolated graph points. This set is countable because each isolated point has a listed ball containing no other graph point. Outside E, every point over S in the graph's closure is an accumulation point.

It remains to show that the construction gives (RECT). Fix a state $S \notin E$. The opponent's allowed speeds form an interval with lowest speed b_- and highest speed b_+: $K[\sigma_j](S) = [b_-, b_+]$. The endpoints are the limits of the infima and suprema of prescribed speeds on shrinking state neighborhoods, so each is approached by opponent prescriptions at states converging to S. The construction gives all four corners: $\{0,1\} \times \{b_-, b_+\} \subseteq K[(\sigma_i^*, \sigma_j)](S)$. For either endpoint b, convex combinations allow every own speed $a \in [0,1]$. Combining the two opponent endpoints then fills the rectangle, including when $b_- = b_+$. Conversely, every allowed joint speed has its first coordinate in $[0,1]$ and its opponent coordinate in $K[\sigma_j](S)$. This proves (RECT).

Step 3: Compare paths and payoffs. Fix a feasible history h_t and a deviation σ_i' that may depend on the full history. Take any capability path allowed by that deviation. From Lemma 3, this path obeys the speed bounds in Step 1 and safety continues to advance at speed g, as stated in (D). Hence $X_s = X_t + g(s-t), s \geq t$. Since $g > 0$, the path visits each full state at most once and reaches the countable set E at only countably many dates. At almost every other date, (RECT) allows the same speeds under (σ_i^*, σ_j). From Lemma 3, the history-based and state-based path rules coincide for Markov profiles. The path we have just checked using the state-based rule is therefore also admissible after the original history h_t. Thus our strategy allows the same path: $\mathcal{P}((\sigma_i', \sigma_j); h_t) \subseteq \mathcal{P}((\sigma_i^*, \sigma_j); h_t)$. Profits and disaster risk are the same on the same path. Taking the supremum over allowed paths gives $V_i((\sigma_i', \sigma_j), h_t) \leq V_i((\sigma_i^*, \sigma_j), h_t) = V_i((\sigma_i^*, \sigma_j), S_t)$. Since the deviation was arbitrary, our Markov strategy is a best response after the history h_t. The same strategy works after every feasible history.

Finally, if a Markov profile σ passes every Markov deviation test, then

$$V_i((\sigma_i', \sigma_j), h_t) \leq V_i((\sigma_i^*, \sigma_j), S_t) \leq V_i(\sigma, S_t) = V_i(\sigma, h_t),$$

where the second inequality is the test applied to σ_i^*. Thus the profile is an MPE. The converse holds because every Markov deviation is also an allowed history-dependent deviation. □

Proof of Proposition 2.

Proof. Put $\delta := r - \alpha + \beta > 0$ and $q := \alpha - r - \lambda > 0$. Choose $\rho \in (0,1)$ sufficiently close to one that

$$r + \beta + 2\log\rho > \alpha - r, \qquad \frac{\rho^2}{1-\rho^2}\frac{1 - e^{-(r+\beta+\lambda)}}{r+\beta+\lambda} > \frac{1}{\delta}. \tag{*}$$

Such a ρ exists because $r + \beta > \alpha - r$ and the left-hand side of the second inequality diverges as $\rho \uparrow 1$.

Let $M_n := e^{\alpha - r}\rho^n$ for $n \geq -1$. Construct a reference path as follows. Firms alternate: in even block n, firm 1 races and firm 2 waits; in odd block n, firm 2 races and firm 1 waits. Let $T_0 = 0$, and end block n at the first date T_{n+1} at which the active firm's discounted flow $f_i(t) := e^{-rt-\Lambda_t}\pi_i(S_t)$ reaches M_n.

While active, $\mathrm{d}\log f_i/\,\mathrm{d}t \in [q, \alpha - r]$; while inactive, $\mathrm{d}\log f_i/\,\mathrm{d}t \in [-(r+\beta+\lambda), -(r+\beta)]$. Hence every target is reached in finite time. Writing $\Delta_n := T_{n+1} - T_n$, we have $\Delta_0 \geq 1$ and

$$\Delta_{n+1} \geq \frac{(r+\beta)\Delta_n + 2\log\rho}{\alpha - r}.$$

The first inequality in (*) implies inductively that $\Delta_n \geq 1$ for all n, so the switching dates have no finite accumulation point.

Payoffs are finite. An active block ending at M_n contributes at most M_n/q, and the following inactive block contributes at most $M_n/(r+\beta)$. Since each firm's own targets decline geometrically by the factor ρ^2, these bounds are summable.

Fix a date $t \in [T_n, T_{n+1})$. Both current discounted flows are at most M_{n-1}. Each firm's next target is at least $\rho^2 M_{n-1}$, and after every future target it waits for at least one unit of time. During that unit its discounted flow is at least $M_m e^{-(r+\beta+\lambda)s}$. Thus, by the second inequality in (*), $\int_t^\infty f_i(s)\,\mathrm{d}s > \frac{M_{n-1}}{\delta} \geq \frac{\max_j f_j(t)}{\delta}$ for $i = 1, 2$. Restarting discounting and survival at t gives

$$V_i^*(h_t) > \frac{\max_j \pi_j(S_t)}{r - \alpha + \beta}, \qquad i = 1, 2. \tag{**}$$

To support this path, prescribe the reference actions as long as the public history exactly coincides with the reference history. After any deviation, prescribe perpetual racing forever. Perpetual racing is an SPE in the stated hazard range: against a racing opponent, every measurable deviation yields at most $\max_j \pi_j(S)/\delta$. By (**), at every reference history each firm's prescribed continuation is strictly better than this off-path bound. Hence no deviation, nor any alternative admissible continuation that leaves the reference path, is profitable. Because perpetual racing is a SPE so it this strategy profile.

Initially, (**) gives $V_i^* > 1/(r - \alpha + \beta)$. Since $g < 1$, $\frac{1}{r-\alpha+\beta} > \frac{1}{r-g(\alpha-\beta)} = P$. Thus $V_i^* > P(0)$ for both firms. □

Proof of Lemma 3.

Proof. (i) Existence. We follow the approximation and compactness argument of Haddad (1984, proof of Theorem 1.1). By Lemma 2, the history correspondence $h \mapsto K[\sigma](h)$ has nonempty compact convex values, is globally bounded, and is upper semicontinuous.

We construct a continuation directly fixing the history h_t. For each positive integer n, partition the future at dates $t + k/n$, $k = 0, 1, \ldots$. Denote the constructed state path by S^n and its development speed pair by a^n. Start by setting $S^n_{t'} = S_{t'}$ for $0 \leq t' \leq t$. Once the path has been constructed through $t + k/n$, let $h^n_{t+k/n}$ be its full history and choose $a^n(t + k/n) \in K[\sigma](h^n_{t+k/n})$. Extend the state path to the next partition date by

$$S^n_{t'} = S^n_{t+k/n} + \left(t' - t - \frac{k}{n}\right) \begin{pmatrix} a^n_1(t + k/n) \\ a^n_2(t + k/n) \\ g \end{pmatrix}, \qquad t' \in \left[t + \frac{k}{n}, t + \frac{k+1}{n}\right].$$

Set $a^n(t') = a^n(t + k/n)$ on the corresponding half-open interval $[t + k/n, t + (k+1)/n)$. Repeating this construction for $k = 0, 1, \ldots$ defines the entire path and its piecewise constant development speeds. Each path is absolutely continuous on every finite interval: it agrees with the feasible history before t and has finitely many linear pieces on each bounded interval after t. Since the chosen development speeds lie in $[0, 1]^2$, every constructed history is feasible and

$$\|\dot{S}^n_{t'}\|^2 = (\dot{A}^n_{1t'})^2 + (\dot{A}^n_{2t'})^2 + g^2 \leq 2 + g^2 \quad \text{for almost every } t' \geq 0.$$

Here the same bound holds before t by feasibility of the fixed past. Absolute continuity therefore gives, for every $0 \leq a < b$,

$$\|S^n_b - S^n_a\| \leq \int_a^b \|\dot{S}^n_{t'}\| \, \mathrm{d}t' \leq \sqrt{2 + g^2}\,(b - a).$$

Thus the paths share a Lipschitz constant independent of n.

Fix a finite horizon $T > t$. Since all paths start at the same state S_0, the preceding inequality implies $\sup_{0 \leq t' \leq T} \|S^n_{t'}\| \leq \|S_0\| + \sqrt{2 + g^2}\, T$, so the restrictions to $[0, T]$ are uniformly bounded and equicontinuous. By Arzelà–Ascoli, they are relatively compact in $\mathcal{C}([0, T], \mathbb{R}^3)$ with the uniform norm. Extract a subsequence converging uniformly on $[0, t + 1]$, then from it a further subsequence converging uniformly on $[0, t + 2]$, and continue in this way. The diagonal subsequence converges uniformly on $[0, t + m]$ for every positive integer m, and hence locally uniformly on $[0, \infty)$, so these limits define a

single continuous path S on $[0,\infty)$.Relabel this subsequence with n. The limit inherits the common Lipschitz bound and agrees with h_t. Passing to the limit in the capability and safety increments gives, for $t \le a < b$, $0 \le A_{ib} - A_{ia} \le b - a$ and $X_b - X_a = g(b-a)$, so the limit is feasible. It remains to show that the limiting path's development speeds belong to $K[\sigma](h_{t'})$ at almost every date, where $h_{t'}$ is the full history generated by that path.

Let $\eta_n(t')$ be the partition date immediately preceding t' (inclusive of t' whenever it is a partition date). For every finite $T > t$ and $t' \in [t,T]$, the stopped histories satisfy

$$d(h^n_{\eta_n(t')}, h_{t'}) \le \sup_{0\le t''\le T} \|S^n_{t''} - S_{t''}\| + \frac{1+\sqrt{2+g^2}}{n} \longrightarrow 0.$$

By construction, $a^n(t') \in K[\sigma](h^n_{\eta_n(t')})$. Fix $q \in \mathbb{Q}^2$ and a date t'. Choose a subsequence along which $q \cdot a^n(t')$ converges to its limsup. Compactness of $[0,1]^2$ gives a further subsequence along which $a^n(t')$ converges to some vector a. Since $h^n_{\eta_n(t')} \to h_{t'}$, the closed-graph property implies that $a \in K[\sigma](h_{t'})$. It follows that $\limsup_{n\to\infty} q \cdot a^n(t') \le \max_{v\in K[\sigma](h_{t'})} q \cdot v$.

To verify measurability of the right-hand side, define the support function $H_q(h) := \max_{v\in K[\sigma](h)} q \cdot v$. This function is upper semicontinuous on the space of stopped histories. The common Lipschitz bound gives $d(h_s, h_{s'}) \le (1+\sqrt{2+g^2})|s-s'|$, so $s \mapsto h_s$ is continuous in the stopped-history metric, so $s \mapsto H_q(h_s)$ is upper semicontinuous and hence measurable. The functions $q \cdot a^n(t')$ and $H_q(h_{t'})$ are uniformly bounded. Applying reverse Fatou on any interval $[a,b] \subseteq [t,T]$ and using uniform convergence of the capability paths gives

$$q \cdot \begin{pmatrix} A_{1b} - A_{1a} \\ A_{2b} - A_{2a} \end{pmatrix} = \lim_{n\to\infty} \int_a^b q \cdot a^n(t')\, \mathrm{d}t' \le \int_a^b \max_{v\in K[\sigma](h_{t'})} q \cdot v \, \mathrm{d}t'.$$

Since this integral inequality holds on every subinterval, Lebesgue differentiation gives the corresponding inequality between the integrands almost everywhere: $q \cdot (\dot{A}_{1t'}, \dot{A}_{2t'}) \le \max_{v\in K[\sigma](h_{t'})} q \cdot v$ for almost every $t' \in [t,T]$. Since $\mathbb{Q}^2$ is countable, these inequalities hold simultaneously outside a null set. At each remaining date, both sides are continuous in q, so density extends the inequality to every $q \in \mathbb{R}^2$. Thus, in every direction q, the speed pair projects no farther than the allowed set does. By Lemma 2, $K[\sigma](h_{t'})$ is closed and convex. The separating-hyperplane theorem therefore implies that the realized speed pair belongs to this set: $(\dot{A}_{1t'}, \dot{A}_{2t'}) \in K[\sigma](h_{t'})$ for almost every $t' \in [t,T]$. Taking a countable sequence of horizons increasing to infinity proves (HA) on the entire continuation. This

establishes existence.

(ii) Restriction and concatenation. Fix $S \in \mathcal{P}(\sigma; h_t)$ and a later date $t' \geq t$. By admissibility, there is a null set $N \subset (t, \infty)$ such that $(\dot{A}_{1t''}, \dot{A}_{2t''}) \in K[\sigma](h_{t''})$ and $\dot{X}_{t''} = g$ for every $t'' \in (t, \infty) \setminus N$. For restriction, retain $h_{t'}$ as the fixed past. Joining this past to $S|_{[t',\infty)}$ reproduces the original full history $h_{t''}$ at every date $t'' \geq t'$. The displayed conditions therefore hold on (t', ∞) outside the null set $N \cap (t', \infty)$. Absolute continuity is preserved under restriction, so $S|_{[t',\infty)} \in \mathcal{P}(\sigma; h_{t'})$.

For concatenation, take $\widetilde{S} \in \mathcal{P}(\sigma; h_{t'})$ and form $\widehat{S}$ as in the statement, with capability coordinates $\widehat{A}_1, \widehat{A}_2$ and safety coordinate $\widehat{X}$. Let $\widetilde{h}_{t''}$ denote the full history of the appended continuation, including the fixed past $h_{t'}$, and let $\widehat{h}_{t''}$ denote the full history of the joined path. These histories satisfy

$$\widehat{h}_{t''} = \begin{cases} h_{t''}, & t \leq t'' \leq t', \\ \widetilde{h}_{t''}, & t'' > t'. \end{cases}$$

The pieces agree at the joining date, since $\widetilde{S}_{t'} = S_{t'}$. Thus $\widehat{S}$ is continuous and, being absolutely continuous on each side of t', is absolutely continuous on every finite interval. Its derivative equals that of S before t' and that of $\widetilde{S}$ after t', wherever those derivatives exist.

Let $\widetilde{N} \subset (t', \infty)$ be a null set outside which the appended continuation satisfies (HA). The derivative and history identities give $(\dot{\widehat{A}}_{1t''}, \dot{\widehat{A}}_{2t''}) \in K[\sigma](\widehat{h}_{t''})$ and $\dot{\widehat{X}}_{t''} = g$ for every $t'' > t$ outside $N \cup \widetilde{N} \cup \{t'\}$. This union is a null set, and the joined path extends h_t. Hence $\widehat{S} \in \mathcal{P}(\sigma; h_t)$.

(iii) Consistency with Markov paths. Fix a history $h_{t'}$ ending at $S_{t'}$. For any nearby history $\widetilde{h}_{t''}$, its terminal state satisfies $\|\widetilde{S}_{t''} - S_{t'}\| \leq d(h_{t'}, \widetilde{h}_{t''})$. If the profile is Markov, its prescription at that history is its prescription at $\widetilde{S}_{t''}$. Taking closed convex hulls and then intersecting over neighborhood sizes gives $K[\sigma](h_{t'}) \subseteq K[\sigma](S_{t'})$. Conversely, for any state S' near $S_{t'}$, translate every state of $h_{t'}$ by the constant vector $S' - S_{t'}$. The translated history is feasible, ends at S', and has stopped-history distance $\|S' - S_{t'}\|$ from $h_{t'}$. Such translations are allowed because (HK) does not fix the initial state of nearby histories. Every prescription in a state neighborhood is therefore also obtained in the corresponding history neighborhood. Taking closed convex hulls and intersections gives the reverse inclusion, and hence $K[\sigma](h_{t'}) = K[\sigma](S_{t'})$.

(iv) Deviations against a Markov opponent. Fix a history-dependent deviation σ_i' and a Markov

opponent σ_j. The endpoint bound in part (iii) implies that the opponent's coordinate of every prescription in an ε history neighborhood belongs to its prescriptions in the ε state neighborhood. Projection preserves convex combinations and limits. Taking closed convex hulls and then intersecting over $\varepsilon > 0$ therefore gives $\operatorname{proj}_j K[(\sigma_i', \sigma_j)](h_{t'}) \subseteq K[\sigma_j](S_{t'})$. Along any admissible deviation path, set $a_i(t') = \dot{A}_{it'}$ wherever the derivative exists and set $a_i(t') = 0$ on the remaining null set. This defines a measurable speed path in $[0, 1]$. The projection inclusion gives $\dot{A}_{jt'} \in K[\sigma_j](S_{t'})$ almost everywhere, while feasibility gives $\dot{X}_{t'} = g$. These are exactly the conditions in (D). □

Proof of the Pacing Verification Lemma.

Proof. If $R(0) = P(0)$, prescribe speed one at strict gaps and ties off safety, and speed g at safety ties. A strict leader's deviations are bounded by $R(D)$ up to the first tie by Lemma 6, applied with $(\eta, \rho) = (\alpha, r + \beta)$, and thereafter by Lemma 7. A follower cannot catch a racing leader, so racing maximizes its flow at the given hazard. At ties, Lemma 7 again gives the bound $R(D)$. The prescribed paths attain these bounds, so Lemma 4 applies. This is the construction with $\bar{D} = 0$.

Henceforth suppose $R(0) < P(0)$, so $\bar{D} > 0$. Write P_α for the waiting-then-pacing value in Lemma 6 with $(\eta, \rho) = (\alpha, r + \beta)$.

Step 1: accommodation values and the stopping formula. The values $e^{-\beta G} V_L(D, G)$ and $C(D, G)$ are both measured per unit of the leader's current flow. For $0 \le D \le \bar{D}$, put $T = D/g$ and $E_D(t) = \exp\{-(r + \beta)t - \int_0^t \lambda(D - gs)\, \mathrm{d}s\}$. Immediate accommodation gives

$$
\begin{aligned}
C(0, G) &= \frac{1}{B} + \left(P(0) - \frac{1}{B}\right) e^{-BG/v}, \\
C(D, G) &= \int_0^G E_D(t)\, \mathrm{d}t + E_D(G) P(D - gG), && G \le T, \\
C(D, G) &= \int_0^T E_D(t)\, \mathrm{d}t + E_D(T) C(0, G - T), && G \ge T.
\end{aligned}
$$

The two formulas for $C(D, G)$ and their first derivatives agree at $G = D/g$. Moreover, C decreases in both arguments: its integrand is

$$
e^{-rt + c(gt - D)_+ - \beta \min\{G, t - (gt - D)_+\} - \int_0^t \lambda(D - gs)\, \mathrm{d}s},
$$

which decreases with either D or G. Also $C(D, G) \ge P_\alpha(D)$.

Define

$$M(D) = \exp\left\{\frac{(r-c)D + \int_0^D \lambda(s)\,\mathrm{d}s}{v}\right\}, \qquad H(D,G) = \frac{C(D,G) - R(D)}{M(D)}.$$

The stopping formula in Section 3.1 becomes

$$e^{-\beta G} V_L(D,G) = R(D) + M(D)\max\left\{0, \max_{D \le d \le \bar{D}} H(d,G)\right\}.$$

Choose the earliest finite maximizer whenever one is optimal.

Along any path in the stopping construction, the leader's payoff per unit of initial flow has integrand $e^{-rt+c(A_t-A_0)-\Lambda_t} e^{-\beta \min\{G, t-(A_t-A_0)\}}$. Dropping the second factor and applying Lemma 6, with Lemma 7 after catch-up, gives the upper bound below. Racing and immediate accommodation give the lower bound:

$$\max\{P_\alpha(D), R(D)\} \le e^{-\beta G} V_L(D,G) \le \max\{P(D), R(D)\}.$$

For $D \ge \bar{D}$, racing attains the upper bound $R(D)$.

Step 2: monotonicity in the gap and the safety residual. The waiting equation and its speed coefficient are

$$\begin{aligned} 0 &= 1 - gC_D - C_G - (r + \beta + \lambda(D))C, \\ Q_C &= C_D + C_G + \alpha C = \frac{1 - vC_G - (B + \lambda(D))C}{g}. \end{aligned}$$

For $0 < D < \bar{D}$, differentiation gives

$$(Q_C)_G = \begin{cases} \dfrac{\beta E_D(G)}{g}\left[(r - gc + \lambda(D))P(D - gG) - v\right], & G < D/g, \\ -\dfrac{\lambda(D)}{g} C_G, & G > D/g. \end{cases}$$

Both expressions are positive. For the first, the scalar waiting inequality gives $(r - gc + \lambda(d))P(d) \ge 1$ at $d = D - gG$, and $\lambda(D) \ge \lambda(d)$. For the second, $C_G < 0$ and $\lambda(D) > 0$; concavity and monotonicity imply the latter because $R(0) < P(0)$ excludes identically zero risk. Since Q_C is continuous at $G = D/g$, $H_D = Q_C/(vM)$, $H_{DG} > 0$ wherever defined, and $H_G < 0$.

When $\phi = 0$, differentiation also gives $(Q_C)_D > 0$ for $G > 0$, so $H(\cdot, G)$ has no interior

local maximum. Since $C(\bar{D}, G) < P(\bar{D}) = R(\bar{D})$, a finite optimal stop must then be immediate, recovering the choice between immediate accommodation and perpetual racing.

The stopping region is defined by $H(D, G) \geq 0$ and $H(D, G) \geq H(d, G)$ for every $d \in [D, \bar{D}]$. Since $H(D, G)$ decreases with G and $H(d, G) - H(D, G)$ increases with G for $d > D$, increasing the gap cannot turn a racing state into a stopping state. Thus the rule that pauses in this closed region and races outside is nondecreasing in G.

On safety, define $J(G) = e^{-\beta G} V_L(0, G) - C(0, G)$. The stopping formula gives

$$J(G) = \max \left\{ 0,\ R(0) - C(0, G),\ \max_{0 \leq d \leq \bar{D}} [H(d, G) - H(0, G)] \right\}.$$

Since $C_G < 0$ and H has increasing differences, J is nonnegative, nondecreasing, and locally Lipschitz. Using $1 - vC_G(0, G) - BC(0, G) = 0$, the safety pacing residual is

$$e^{\beta G} - v\partial_G V_L(0, G) - (r - gc)V_L(0, G) = -e^{\beta G}[vJ'(G) + BJ(G)] \leq 0$$

almost everywhere. If $J(G) = 0$, pacing reduces the gap and keeps J zero until the tie. Otherwise the earliest optimal choice is a positive target or perpetual racing. The safety rule is therefore also nondecreasing in G.

If $R(0) \leq 1/B$, safety is a stopping state at every finite gap. Immediate pacing beats perpetual racing because $C(0, G) > 1/B \geq R(0)$. To compare pacing with racing to a positive target d and then accommodating, note that $H(d, G) - H(0, G)$ increases with G. As $G \to \infty$, catch-up disappears and $C(d, G) \to P_\alpha(d)$, so this difference converges to $(P_\alpha(d) - R(d))/M(d) - [1/B - R(0)]$. This is the gain from racing to d and then waiting and pacing, rather than pacing immediately, in the auxiliary problem with $(\eta, \rho) = (\alpha, r + \beta)$. By Lemma 6, that gain is nonpositive. Hence $H(d, G) \leq H(0, G)$ at every finite gap.

Step 3: accommodation remains optimal along the prescribed path. Start at a stopping state (D_0, G_0) with $D_0, G_0 > 0$. While waiting, $\dot{D} = -g$ and $\dot{G} = -1$ until safety or a tie, and

$$\frac{\mathrm{d}}{\mathrm{d}t} Q_C(D_t, G_t) = (r + \beta + \lambda(D_t))Q_C(D_t, G_t) + \lambda'(D_t)C(D_t, G_t) - \alpha.$$

At $G = D/g$, the derivative along the waiting path is given by the same expression from either formula for C. Both $C(D_t, G_t)$ and $\lambda'(D_t)$ increase during waiting. Thus the derivative of Q_C multiplied by $\exp\{-\int_0^t (r + \beta + \lambda(D_s))\,\mathrm{d}s\}$ can change sign only from

negative to positive, so the transformed coefficient is maximized at an endpoint. Initially $Q_C \leq 0$ by optimality against larger targets. Its terminal limit is $-\lambda_0 C(0,G)/g \leq 0$ at safety, or $P'(D) + cP(D) \leq 0$ at a tie below $\bar{D}$. Hence $Q_C \leq 0$ throughout waiting.

At a later waiting state (D_1, G_1), the gap when the path passed $d \in [D_1, D_0]$ was $G_1 + (d - D_1)/g \geq G_1$. Since $(Q_C)_G > 0$, $Q_C(d, G_1) \leq 0$, so $H(D_1, G_1) \geq H(d, G_1)$ throughout this interval. For $d \in [D_0, \bar{D}]$, increasing differences and initial optimality give $H(d, G_1) \leq H(D_0, G_1)$, while $H(D_0, G_1) \geq H(D_0, G_0) \geq 0$. Thus (D_1, G_1) remains a stopping state. Continuity extends this conclusion to safety, where monotonicity of J keeps pacing optimal until catch-up.

Step 4: verify the leader's best response. The racing and waiting residuals are, respectively, $e^{\beta G} + v\partial_D V_L - (r - c + \lambda)V_L$ and $e^{\beta G} - g\partial_D V_L - \partial_G V_L - (r + \lambda)V_L$. Their difference is $e^{\beta G}Q$, where $Q = e^{-\beta G}(\partial_D V_L + \partial_G V_L + cV_L)$. In the stopping region, the waiting residual is zero and the racing residual is $e^{\beta G}Q_C \leq 0$. Where $V_L = e^{\beta G}R$, the racing residual is zero; $P_\alpha \leq R$ and Lemma 6 imply $R' + \alpha R \geq 0$, so the waiting residual is nonpositive.

Fix a gap G at which the maximum over targets in the stopping formula is differentiable with respect to G, and consider a racing interval where $V_L > e^{\beta G}\max\{C, R\}$. Such differentiability holds for almost every gap. On this interval, $V_L(D, G) = e^{\beta G}[R(D) + M(D)Z(G)]$, where Z is the maximum of $H(d, G)$ over a fixed target interval bounded away from the current D. The maximum is locally Lipschitz; wherever differentiable, every maximizing target has G-derivative $Z'(G)$. The racing residual is zero.

Let d_* be the optimal target ending the interval. It is interior because $C(\bar{D}, G) < R(\bar{D})$ for $G > 0$. Optimality and the envelope identity give $H_D(d_*, G) = 0$, $\partial_G V_L = e^{\beta G}(C_G + \beta C)$, $Q(d_*, G) = 0$, and $\partial_D V_L(d_*, G) = e^{\beta G}C_D(d_*, G) \leq 0$. Along the racing interval, $v\partial_{DD}V_L = (r - c + \lambda)\partial_D V_L + \lambda' V_L$, and $v\partial_D Q = (r - c + \lambda)Q + \lambda' e^{-\beta G}V_L - \alpha$. Integrating the first equation backward gives $\partial_D V_L \leq 0$. Consequently $\lambda' e^{-\beta G}V_L - \alpha$ is nonincreasing in D, and $\partial_D(Q/M) = \dfrac{\lambda' e^{-\beta G}V_L - \alpha}{vM}$ changes sign at most once, from positive to negative. Thus Q/M is minimized at an endpoint.

Let d_- be the left endpoint. If $d_- > 0$, Step 3 keeps accommodation optimal at $(d_- - gt, G - t)$ for small $t \geq 0$. Racing to the fixed target d_* remains feasible there, so $R(d_- - gt) + M(d_- - gt)H(d_*, G - t) \leq C(d_- - gt, G - t)$, with equality at $t = 0$. Right differentiation shows that this fixed-target payoff has nonpositive waiting residual. Its racing residual is zero, and its derivatives agree with those of the envelope at the chosen

gap, so $Q(d_-, G) \geq 0$. If $d_- = 0$, the safety residual from Step 2 gives

$$vQ(0+, G) = e^{-\beta G}\left\{\lambda_0 V_L(0, G) - \left[e^{\beta G} - v\partial_G V_L(0, G) - (r - gc)V_L(0, G)\right]\right\} \geq 0.$$

Hence $Q \geq 0$ throughout the interval, and the waiting residual is nonpositive. Fubini's theorem gives both residual inequalities almost everywhere on $D > 0, G > 0$.

The payoff residual at any speed $a \in [0, 1]$ is a convex combination of the racing and waiting residuals. To cover paths on switching boundaries, smooth V_L with a nonnegative mollifier of radius ε. The derivative coefficients are constant and the flow and hazard are locally Lipschitz, so both payoff residuals computed using the smoothed value are at most $O(\varepsilon)$ uniformly on compact subsets of $D > 0, G > 0$. Integrating along any absolutely continuous control path and letting $\varepsilon \downarrow 0$ proves the verification inequality. Apply this up to finite dates before a tie, then use continuity.

Below safety, $V_L(D, G) = e^{\beta G}\frac{1-e^{(r-c)D/v}}{r-c} + e^{(r-c)D/v}V_L(0, G)$. Its racing residual is zero, and

$$Q(D, G) = \frac{\alpha(1 - e^{(r-c)D/v})}{r - c} + \frac{e^{(r-c)D/v}}{v}\left\{e^{-\beta G}[(r - gc)V_L(0, G) + v\partial_G V_L(0, G)] - 1\right\} \geq 0$$

almost everywhere by Step 2. The same smoothing argument verifies the inequality along paths in $D < 0$. On $D = 0$, the leader's speed is g almost everywhere, and the safety residual applies; smoothing $V_L(0, \cdot)$ covers its nondifferentiability set. Concatenate across these regions as in Lemma 6. At a tie, Lemma 7 bounds every continuation by $\max\{P(D), R(D)\}$.

Integrate to the first tie or infinity. The bound $V_L(D, G) \leq e^{\beta G}/(r - c)$ makes the discounted terminal term vanish: before a tie the gap cannot increase, the frontier grows at speed at most one, and $r > c$. By Step 3 and the prescribed optimal common-speed path after a tie, the earliest optimal stopping path attains the bound.

Step 5: verify the follower's best response. Use $b(D, G)$ for strict leaders, $a^*(D)$ at ties, and speed one for strict followers.

For an initial strict follower, write $\bar{A}_t$ and $\bar{A}^F_t$ for the prescribed frontier and follower capabilities. The leading firm races to its earliest optimal stopping point, then waits for safety and paces; if a tie occurs while waiting, both firms complete that continuation. If no finite target is optimal, the leading firm races forever. The prescribed follower satisfies $\bar{A}^F_t = \min\{A^F_0 + t, \bar{A}_t\}$.

Fix any unilateral deviation by the follower, with the opponent following its prescribed

strategy, and any continuation path admitted by (D). We first show that the opponent's capability is at least $\bar{A}_t$ while the deviator remains weakly below $\bar{A}_t^F$ and is a strict follower. During prescribed racing, the deviator's capability is at most $A_0^F + t$. Until the opponent departs from the prescribed path, the gap is at least the prescribed gap. By Step 2 the gap stays in the open racing region until the earliest stopping point, including at safety crossings with $J > 0$, so the opponent's regularized speed 1. During the prescribed pause, the opponent's capability cannot decrease. During prescribed pacing, the prescribed frontier lies on the safety boundary, which the opponent cannot cross from above: strictly below safety its speed is one and $\dot{D} = v > 0$. This proves the comparison.

If the prescribed follower never catches up, then $\bar{A}_t^F = A_0^F + t$ forever. The deviator has no higher capability and faces no lower opponent capability or hazard, so its payoff is at most the prescribed payoff.

Otherwise, let T_C be the prescribed catch-up date and let $\tau \geq T_C$ be the first date when the deviator reaches $\bar{A}_t$, with $\tau = \infty$ if it never does. Before T_C, its capability is at most $A_0^F + t = \bar{A}_t^F$; from T_C until τ, it is below $\bar{A}_t = \bar{A}_t^F$. It remains a strict follower before τ, since an earlier tie would require reaching an opponent capability at least $\bar{A}_t$. The comparison therefore applies throughout this interval: its flow payoff is no greater than the prescribed follower's, and its cumulative hazard is no smaller.

If $\tau = \infty$, this proves the payoff bound. If $\tau < \infty$, the deviator has capability $\bar{A}_\tau$ and remains weakly behind. By Lemma 7, its continuation payoff is at most $e^{c\bar{A}_\tau} \max\{P(\bar{A}_\tau - X_\tau), R(\bar{A}_\tau - X_\tau)\}$. The prescribed firms are tied at $\bar{A}_\tau$ and their continuation attains this bound. Combining the continuation bound with the preceding flow and survival comparisons proves that no follower deviation is profitable.

Step 6: one path gives both equilibrium payoffs. The prescribed path, followed by the prescribed optimal common-speed path after a tie, is jointly admissible. Step 4 shows that it attains the leader's value, and Step 5 bounds every follower deviation by its follower payoff. Both payoffs are finite: the leader's value is bounded in Step 4, and the follower's flow is at most $e^{c(A_0^F + t)}$. Thus Lemma 4 establishes the pacing equilibrium. □

Proof of the Comparative-Statics Proposition.

Proof of Proposition 1. We express effective hazard as an average over racing dates, then compare the hazard and the weights. Write the common growth rate as $c = \alpha - \beta$ and the speed at which a race moves away from safety as $v = 1 - g$. Define the density of payoff

weights by

$$w(t) := \frac{1}{R(0)} \exp\left\{-(r-c)t - \int_0^t \lambda(vs)\ \mathrm{d}s\right\}, \qquad t \geq 0.$$

The density integrates to one. Integrating the derivative of the exponential gives

$$\lambda^{\text{eff}} = \int_0^\infty \lambda(vt) w(t)\ \mathrm{d}t. \tag{WH}$$

The boundary terms are one at zero and zero at infinity, since $r > c$. Concavity bounds λ by an affine function at large distances, so all integrals below are finite. The jump at zero does not affect the integrals or the derivative at positive dates.

Step 1: Higher risk. A pointwise increase in λ weakly lowers $R(0)$ and hence weakly raises $\lambda^{\text{eff}}/(1-g)$ at fixed g and r.

Step 2: Faster safety. Compare racing speeds away from safety $0 < v_2 < v_1$, and let w_2 and w_1 be their payoff densities. Their ratio is nondecreasing:

$$\frac{\mathrm{d}}{\mathrm{d}t} \log \frac{w_2(t)}{w_1(t)} = \lambda(v_1 t) - \lambda(v_2 t) \geq 0.$$

If the densities coincide, their averages coincide. Otherwise, because both densities integrate to one, their difference changes sign at most once, from negative to positive. Integrating any nondecreasing function against $w_2 - w_1$ is therefore nonnegative: subtract its value at a crossing, and each side of the crossing has a nonnegative integrand.

Concavity and $\phi(0) = 0$ give, for every $t > 0$,

$$\frac{\lambda(v_2 t)}{v_2} = \frac{\lambda_0 + \phi(v_2 t)}{v_2} \geq \frac{\lambda_0 + \phi(v_1 t)}{v_1} = \frac{\lambda(v_1 t)}{v_1}.$$

The function on the right is nondecreasing in t. Applying the density comparison and (WH) gives

$$\int_0^\infty \frac{\lambda(v_2 t)}{v_2} w_2(t)\ \mathrm{d}t \geq \int_0^\infty \frac{\lambda(v_1 t)}{v_1} w_2(t)\ \mathrm{d}t \geq \int_0^\infty \frac{\lambda(v_1 t)}{v_1} w_1(t)\ \mathrm{d}t.$$

Thus the ratio increases weakly with g. Under linear risk with $\kappa > 0$, the density ratio is strictly increasing and the integrand $\lambda(v_1 t)/v_1 = \kappa t$ is strictly increasing, so the last inequality is strict.

Step 3: Greater patience. At fixed g and hazard profile, increasing r multiplies the unnormalized payoff weights by a strictly decreasing exponential in t. The same crossing argument,

with signs reversed, lowers the average of the nondecreasing function $\lambda(vt)$ in (WH). The decrease is strict when ϕ is nonconstant: the density ratio is strictly decreasing and the hazard differs on intervals of positive length. If $\phi = 0$, (WH) instead gives $\lambda^{\text{eff}} = \lambda_0$. Comparing λ^{eff}/v with c and α in Theorem 1 proves the equilibrium-region claims. □

Proof of result on Monitoring Lags .

Admissible paths with observation lags. Write h_t for a feasible full capability history on $[-L, t]$, conditional on survival, and h_{it} for the private history it induces for firm i. A pure strategy σ_i assigns a speed $\sigma_i(h_{it}) \in [0,1]$ to each feasible private history and is measurable in time along every feasible full history. We define admissible paths by regularizing the joint prescriptions on full histories. Extend each full history constantly after its terminal date, and set $d_L(h_t, \widetilde{h}_{t'}) := |t - t'| + \sup_{s \geq -L} \left\| h_t(s \wedge t) - \widetilde{h}_{t'}(s \wedge t') \right\|$. For a strategy profile σ, let $K_L[\sigma](h_t) := \bigcap_{\varepsilon > 0} \overline{\text{co}} \left\{ \left(\sigma_1(\widetilde{h}_{1t'}), \sigma_2(\widetilde{h}_{2t'})\right) : \widetilde{h}_{t'} \text{ is feasible and } d_L(h_t, \widetilde{h}_{t'}) < \varepsilon \right\}$, where $\widetilde{h}_{it'}$ is firm i's private history induced by the *same* full history $\widetilde{h}_{t'}$. An admissible path from h_0 extends that history and satisfies $(\dot{A}_{1t}, \dot{A}_{2t}) \in K_L[\sigma](h_t)$ for almost every $t > 0$. Write $\mathcal{P}(\sigma; h_0)$ for these paths. As in the baseline model, each firm evaluates a profile by the supremum of its payoff over $\mathcal{P}(\sigma; h_0)$.

Proof of Proposition 4. Subtract profits over $[0, L]$, which are fixed by the initial history, and divide by $e^{-rL + (\alpha - \beta)A_0}$. A model developed at date s is deployed at $s + L$, so the normalized payoff on a path is

$$\int_0^\infty \exp\left\{ -rs + \alpha(A_{is} - A_0) - \beta(A_{js} - A_0) - \int_0^s \lambda(A_t - X_t - gL)\, \mathrm{d}t \right\} \mathrm{d}s.$$

Joint pacing gives $P(0) = 1/[r - g(\alpha - \beta)]$. To sustain racing, prescribe speed 1 at every private history. Under any unilateral deviation, the rival still races. Since both firms start at A_0 and speed cannot exceed one, the deviator cannot overtake its rival. Its choices therefore leave hazard unchanged, while racing maximizes its profits at every date. Both firms attain the normalized payoff $R(0)$, so racing is a Nash equilibrium.

Now suppose $\bar{\lambda} < (1 - g)\alpha$. We will prove the following claim ($\star$): Fix any date $\tau \geq 0$, any feasible history h_τ, and any feasible continuation of the rival's capability extending h_τ (observed by firm i or not). Against this rival path, racing from τ gives firm i a weakly higher normalized payoff than any other feasible own continuation.

To prove $\star$, reindex by development time $s = u - \tau$, writing $A_{is} := A_{i,\tau+s}$, $A_{js} :=$

$A_{j,\tau+s}$, and $X_s := X_\tau + gs$, so that A_{i0} and X_0 are the values at the end of h_τ. Capability developed at time s is deployed at $s + L$, so firm i's choices cannot affect profits or hazard on the first L units of time after τ. Every own continuation therefore has the same profits and the same hazard exposure up to deployment date L, and we compare continuations only from L onward. Let $\delta_s := A_{i0} + s - A_{is} \geq 0$ be the capability that an alternative path gives up relative to racing by date s. It is nondecreasing because $\dot{\delta}_s = 1 - \dot{A}_{is} \geq 0$. Capability developed at time t is deployed at date $L + t$, when safety is $X_t + gL$. The rival's deployed distance from safety is thus $Y_t := A_{jt} - X_t - gL$, and firm i's distance under racing is $D_t^R := A_{i0} + t - X_t - gL = A_{i0} - X_0 - gL + (1-g)t$. At date $L + t$, racing and the alternative have hazards $\lambda(\max\{D_t^R, Y_t\})$ and $\lambda(\max\{D_t^R - \delta_t, Y_t\})$, respectively. The rival can only shrink the gap between them: if $Y_t \geq D_t^R$ the gap is 0; if $D_t^R - \delta_t \leq Y_t < D_t^R$ it is $\lambda(D_t^R) - \lambda(Y_t) \leq \lambda(D_t^R) - \lambda(D_t^R - \delta_t)$; and if $Y_t < D_t^R - \delta_t$ it equals $\lambda(D_t^R) - \lambda(D_t^R - \delta_t)$. Since also $\delta_t \leq \delta_s$ for $t \leq s$ and λ is nondecreasing, $\lambda_{L+t}^R - \lambda_{L+t} \leq \lambda(D_t^R) - \lambda(D_t^R - \delta_s), \qquad 0 \leq t \leq s$. Integrating and changing variables with $D_t^R = D_0^R + (1-g)t$, $\int_L^{L+s} (\lambda_u^R - \lambda_u)\, \mathrm{d}u \leq \frac{1}{1-g}\left[\int_{D_s^R - \delta_s}^{D_s^R} \lambda(x)\, \mathrm{d}x - \int_{D_0^R - \delta_s}^{D_0^R} \lambda(x)\, \mathrm{d}x\right] \leq \frac{\bar{\lambda}}{1-g}\,\delta_s$. Thus slowing down by δ_s saves at most $\bar{\lambda}\delta_s/(1-g)$ of cumulative hazard.

Now compare integrands date by date. For a path A_i, the normalized payoff has integrand $f(s; A_i) = \exp\{-rs + \alpha(A_{is} - A_{i0}) - \beta(A_{js} - A_{j0}) - (\Lambda_{L+s} - \Lambda_L)\}$. The rival's path is the same for both plans, so the β term cancels. Racing raises $\alpha(A_{is} - A_{i0})$ by $\alpha\delta_s$, and by the inequality above it raises cumulative hazard by at most $\bar{\lambda}\delta_s/(1-g)$. Hence $\dfrac{f(s; \text{race})}{f(s; A_i)} \geq \exp\left\{\left(\alpha - \dfrac{\bar{\lambda}}{1-g}\right)\delta_s\right\} \geq 1 \quad \text{for every } s \geq 0.$

Integrating over s proves ($\star$). The same comparison holds in expectation under any fixed distribution over rival continuation paths, including when the firm's own plan conditions on information as it arrives.

Now we show that joint pacing is a Nash equilibrium outcome if and only if $\widehat{B}(L) \leq P(0)$. Let $q_L(s) := g\min\{s, L\} + (s - L)_+$. This is the rival's capability increase when it paces through date L and then races. For any absolutely continuous u with $u(0) = 0$ and $\dot{u} \in [0, 1]$ almost everywhere, define

$$J_L(u) := \int_0^\infty \exp\left\{-rs + \alpha u(s) - \beta q_L(s) - \int_0^s \lambda\big(\max\{u(v), q_L(v)\} - gv\big)\, \mathrm{d}v\right\} \mathrm{d}s$$

where $\widehat{B}(L) := \sup_u J_L(u)$. As above, s measures development time, whose consequences are deployed L units later.

To show necessity, suppose an equilibrium outcome jointly paces. Fix any feasible devel-

opment path u for firm i and let it deviate to that path. Before date L, firm j observes no deviation from the initial history, so there is an admissible continuation on which it paces throughout this interval. Before date L, replacing firm i's unobserved development leaves firm j's private history unchanged. The nearby histories supporting firm j's pacing speed can therefore be transferred to histories with the deviating firm's hidden development path, so pacing remains admissible for firm j throughout this interval under the deviation. Thereafter its speed is at most one, so its capability increase is at most $q_L(s)$ at every date s. Holding firm i's path fixed, a lower rival capability raises firm i's flow profit and weakly lowers the frontier hazard. The deviation therefore gives firm i at least $J_L(u)$. Joint pacing gives $P(0)$, so equilibrium requires $J_L(u) \leq P(0)$ for every u, or $\widehat{B}(L) \leq P(0)$.

For sufficiency, suppose $\widehat{B}(L) \leq P(0)$. Prescribe speed g while a firm has always paced and all its observations of its rival are consistent with pacing; otherwise prescribe speed one. Joint pacing is admissible. Consider any admissible path under this profile or a unilateral deviation, and let t be its first departure from joint pacing. At date t the firms are still tied on the prescribed internal-capability path. From then on, the nondeviating firm develops at speed at least g, and at speed one after it has observed the deviation, no later than $t + L$. Its capability increase from date t is therefore at least $q_L(s)$ at every elapsed date s. Holding the other firm's development path fixed, this higher rival capability weakly lowers its flow profit and raises the frontier hazard. Its normalized payoff from the deviation onward is thus at most $\widehat{B}(L) \leq P(0)$.

Before t, the path agrees with joint pacing. Adding the payoff on that initial interval shows that the whole path gives the firm at most its joint-pacing payoff. The same argument applies to either firm and to every unilateral deviation. Joint pacing attains both assigned payoffs on one path, so the proposed profile is a Nash equilibrium.

Differentiating the displayed formula for $B(L)$ gives $B'(L) = \beta(1-g)e^{\beta(1-g)L}\int_L^\infty e^{-(r-\alpha+\beta)s-\Lambda_s^R}\,\mathrm{d}s > 0$. The integrals are finite for every finite lag because $r > \alpha - \beta$. At zero lag, $B(0) = R(0) = \frac{1}{r-\alpha+\beta+\lambda^{\text{eff}}}$.By definition, $\widehat{B}(L) \geq B(L)$. Under the hazard bound $\bar{\lambda} < (1-g)\alpha$, claim ($\star$) shows that racing maximizes the payoff against the rival path q_L, so $\widehat{B}(L) = B(L)$. In part (a), for every $L > 0$, $\widehat{B}(L) \geq B(L) > B(0) = R(0) > P(0)$, which rules out pacing.

For the long-lag limit, write $B(L) = \int_0^\infty e^{-(r-\alpha+\beta g)s-\Lambda_s^R}e^{-\beta(1-g)\max\{s-L,0\}}\,\mathrm{d}s$. The last factor increases to one with L so monotone convergence gives $B(\infty) := \lim_{L\to\infty} B(L) = \int_0^\infty e^{-(r-\alpha+\beta g)s-\Lambda_s^R}\,\mathrm{d}s$. Under the hazard bound in part (b), $\Lambda_s^R \leq \bar{\lambda}s$ gives $B(\infty) \geq \int_0^\infty e^{-(r-\alpha+\beta g+\bar{\lambda})s}\,\mathrm{d}s > P(0)$. If $r - \alpha + \beta g + \bar{\lambda} > 0$, the integral exceeds $P(0)$ because $\bar{\lambda} < (1-g)\alpha$; otherwise it is infinite. When $\lambda^{\text{eff}} > (1-g)(\alpha-\beta)$, we have $B(0) < P(0)$, so continuity and strict monotonicity give the unique finite root $L^* > 0$. At equality,

$B(0) = P(0)$ and $L^* = 0$. The condition $\widehat{B}(L) \leq P(0)$ and the equality $\widehat{B}(L) = B(L)$ therefore give the cutoff in part (b).

To prove part (c), suppose $\lambda^{\text{eff}} \geq (1-g)\alpha$. By the definition of effective hazard, $\int_0^\infty e^{-(r-\alpha+\beta)s}[e^{-\Lambda_s^R} - e^{-(1-g)\alpha s}]\,ds \leq 0$. Because Λ_s^R/s is nondecreasing, the term in brackets changes sign at most once, from positive to negative. Multiplying by the increasing factor $e^{\beta(1-g)s}$ therefore preserves the weakly negative integral, so $B(\infty) \leq \dfrac{1}{r-\alpha+\beta g+(1-g)\alpha} = P(0)$.

Now prescribe speed g at every private history. Against a pacing rival, firm i faces the auxiliary problem in Lemma 6 with $\eta = \alpha$ and $\rho = r + \beta g$. Its racing and pacing values from safety are $B(\infty)$ and $P(0)$, respectively. Since $B(\infty) \leq P(0)$, the lemma bounds every deviation payoff by $P(0)$. Joint pacing attains this payoff, proving part (c). □

Proof of result on Endogenous Safety.

Proof of Proposition 5. Write $c = \alpha - \beta$, $b = 1 - g$, $q = g + 2\eta$, and $k = 1 + 2\eta$, so $v = q/k$. The path and verification lemmas apply with the new safety equation: the state velocity is bounded and affine in the speed pair, and restriction, concatenation, and state translations preserve feasibility.

Log convexity of R on $[0,\infty)$, the equation $bR' = (r - c + \lambda)R - 1$, and the formula $R(D) = (1-z)/(r-c) + zR(0)$ for $D < 0$, where $z = e^{(r-c)D/b}$, imply

$$\sup_{D\in\mathbb{R}} \frac{-R'(D)}{R(D)} = \frac{-R'(0)}{R(0)} = \frac{\lambda^{\text{eff}}}{b} \leq \frac{\kappa}{r-c}.$$

The last inequality follows because $-R'(0)/R(0)$ is κ times the mean of a density proportional to $e^{-(r-c)t-\kappa b t^2/2}$; multiplying an exponential density by the decreasing factor $e^{-\kappa b t^2/2}$ weakly lowers its mean.

Against a racing opponent, the candidate payoff is $\pi_i R(D)$. Its own-speed coefficient, divided by current profit, is $\alpha R + (1+\eta)R'$ when the firm leads and $\alpha R + \eta R'$ when it follows. When $\lambda^{\text{eff}} \leq \alpha b/(1+\eta)$, both coefficients are nonnegative at every state, so racing has zero payoff residual and every slower speed has a nonpositive residual. At almost every date spent tied, the deviator also has speed one. Integrating therefore bounds every deviation by $\pi_i R(D)$, which racing attains. If the condition fails, a strict leader on safety gains by pausing for a sufficiently short interval and then racing. This proves (i).

To prove sufficiency, we first obtain a bound at ties. Let $A_L = \max\{A_1, A_2\}$, $A_F =$

$\min\{A_1, A_2\}$, and $G = A_L - A_F$. Put $Y = [(1+\eta)A_L + \eta A_F]/k$ and $Z = [X + \eta(A_1 + A_2)]/k$. Then $\dot{Z} = v$, $D = k(Y - Z)$, and $\dot{Y} \in [0,1]$ almost everywhere. Against an opponent that races whenever strictly behind, a firm starting tied remains weakly behind. Its log profit is $cY - [\alpha(1+\eta) + \beta\eta]G/k \leq cY$. Lemma 6, applied to Y with safety speed v and hazard $\kappa k(Y - Z)_+$, bounds its payoff by the optimal tied payoff.[20] For $D \geq 0$, this is $\max\{P(D), R(D)\}$ per unit of initial profit; below safety it is the value of racing to zero and taking the optimal continuation. A strict leader facing a racing follower satisfies the same bound: before catch-up its change in log profit is $c(Y_t - Y_0) + [\alpha\eta + \beta(1+\eta)](G_t - G_0)/k \leq c(Y_t - Y_0)$, and after catch-up the tied bound applies. Tied firms attain the bound with the prescribed tied policy.

If $P(0) = R(0)$, the scalar bound is R everywhere. Prescribe racing except at safety ties, where both firms choose v. Since $cb/k < \alpha b/(1+\eta)$, part (i) bounds deviations until the first tie, and the scalar bound applies thereafter. Racing from other states and pacing from a safety tie attain these bounds, proving sufficiency at equality.

Henceforth assume (ES) and $P(0) > R(0)$. The Lemma 6and its proof give a unique crossing $\bar{D} > 0$ and, on $[0, \bar{D}]$, $(r + \kappa D)P + qP' = 1$ and $kP' + cP \leq 0$, with strict inequality for $0 < D \leq \bar{D}$. Write $p = g + \eta$, $h = p/(1+\eta)$, $w = 1 - h = b/(1+\eta)$, and $B = r - \alpha h + \beta$.

For $0 \leq D \leq \bar{D}$, let $C(D, G)$ be the leader's payoff from immediate accommodation, per unit of its current profit, when the follower races. With $E_D(t) = e^{-(r+\beta)t - \kappa(Dt - pt^2/2)}$, it is

$$C(0,G) = \frac{1}{B} + \left(P(0) - \frac{1}{B}\right) e^{-BG/w},$$

$$C(D,G) = \begin{cases} \displaystyle\int_0^G E_D(t)\, \mathrm{d}t + E_D(G)P(D - pG), & G \leq D/p, \\ \displaystyle\int_0^{D/p} E_D(t)\, \mathrm{d}t + E_D(D/p)C(0, G - D/p), & G \geq D/p. \end{cases}$$

The first derivatives agree at $G = D/p$. Differentiation gives $C_D \leq 0$ and $C_G < 0$; in the first case, $C_G = E_D(G)[\eta P'(D - pG) - \beta P(D - pG)]$.

Set $M(D) = e^{[(r-c)D + \kappa D^2/2]/b}$ and $H = (C - R)/M$. The leader's value among paths that race and then accommodate, including perpetual racing, is

$$U(D,G) = R(D) + M(D) \max\left\{0, \max_{D \leq d \leq \bar{D}} H(d,G)\right\}, \qquad 0 \leq D \leq \bar{D}.$$

[20] Here and subsequently $x_+ := \max\{x, 0\}$.

Here U is measured per unit of current profit. Among targets attaining the maximum, select the smallest finite one if any, and otherwise select ∞. Above $\bar{D}$, set $U = R$; below safety, set $U(D,G) = (1-z)/(r-c) + zU(0,G)$ and use the target selected at $(0,G)$. The function U is locally Lipschitz and extends continuously to $G = 0$, where it equals the scalar tied value. For $G > 0$, immediate accommodation is selected on the closed set $\{(D,G) : 0 \le D \le \bar{D},\ U(D,G) = C(D,G)\}$. The scalar bound gives $U \le \max\{P,R\}$ above safety. In particular, $C(\bar{D},G) < R(\bar{D})$ for $G > 0$: the strict leader's gap falls over a positive interval during accommodation, making the scalar profit comparison strict.

The waiting equation is $1 - pC_D - C_G - (r+\beta+\kappa D)C = 0$. Define $Q_C = (1+\eta)C_D + C_G + \alpha C = [1 - wC_G - (B+\kappa D)C]/h$, so $H_D = Q_C/(bM)$. For $G > D/p$, $(Q_C)_G = -\kappa D C_G/h \ge 0$. For $G < D/p$, put $d = D - pG$ and $a = h(r-c+\kappa d) + \kappa pG$. Differentiation and $qP'' = -\kappa P - (r+\kappa d)P'$ give

$$(Q_C)_G = \frac{E_D(G)}{h}\left\{\left(\beta a - \frac{wp\eta\kappa}{q}\right)P(d) - \left(\eta a + \frac{wp\eta(r+\kappa d)}{q} + wp\beta\right)P'(d)\right\} \ge 0.$$

Indeed, the second bound in (ES) gives $\beta h(r-c) \ge wp\eta\kappa/q$, and $P' < 0$. Thus $H_G < 0$ and H has increasing differences. Assigning gain zero to target infinity preserves increasing differences, so the selected target is nondecreasing in G. It is also nondecreasing in D, since increasing D only removes smaller targets.

Accommodation remains optimal once begun. During a pause, $(\dot{D},\dot{G}) = (-p,-1)$ and $\dot{Q}_C = (r+\beta+\kappa D)Q_C + (1+\eta)\kappa C - \alpha$. Since C increases along this path, the derivative of $e^{-\int_0^t (r+\beta+\kappa D_s)\,\mathrm{d}s} Q_C(D_t,G_t)$ changes sign at most once, from negative to positive. Its initial value is nonpositive by stopping optimality. Its terminal value is zero at safety or $kP' + cP \le 0$ at a tie. Hence $Q_C \le 0$ throughout the pause.

Suppose a pause starts at (D_0,G_0) and reaches (D_1,G_1). For $d \in [D_1,D_0]$, the gap when the path passed d was $G_1 + (d-D_1)/p \ge G_1$. Monotonicity of Q_C in G therefore gives $Q_C(d,G_1) \le 0$, and hence $H(D_1,G_1) \ge H(d,G_1)$. For $d \ge D_0$, increasing differences and initial optimality give $H(d,G_1) \le H(D_0,G_1)$, while $H(D_0,G_1) \ge H(D_0,G_0) \ge 0$. Thus immediate accommodation remains optimal. On safety, $J(G) = U(0,G) - C(0,G)$ is nonnegative, nondecreasing, and locally Lipschitz, because it is the maximum of 0, $R(0) - C(0,G)$, and $H(d,G) - H(0,G)$ over $d \in [0,\bar{D}]$. Its safety residual is $1 - wU_G - BU = -wJ' - BJ \le 0$ almost everywhere. When $J = 0$, pacing reduces G and keeps J zero.

We show that no leader deviation yields more than U per unit of current profit. The

waiting and racing residuals are

$$\mathcal{W} = 1 - pU_D - U_G - (r + \beta + \kappa D)U,$$
$$\mathcal{R} = 1 + bU_D - (r - c + \kappa D)U.$$

Their difference is $Q = (1+\eta)U_D + U_G + \alpha U$. In the stopping region, $\mathcal{W} = 0$ and $Q = Q_C \leq 0$. For $D \geq \bar{D}$, racing attains the scalar upper bound. For $0 < D < \bar{D}$ where $U = R$, the limit $C(D, \infty)$ is the waiting value in the auxiliary scalar problem with safety speed h, distance $D/(1+\eta)$, growth coefficient α, and discount rate $r + \beta$. Since $C(D, \infty) \leq C(D, G) \leq R(D)$, the Lemma 6gives $\alpha R + (1+\eta)R' \geq 0$, so $\mathcal{W} \leq 0$.

For each rational $a \in [0, \bar{D}]$, let $m_a(G) = \max_{a \leq d \leq \bar{D}} H(d, G)$. These functions are locally Lipschitz. Fix a gap G where every m_a and J is differentiable, excluding a null set. On each racing interval (d_-, d_*) where $U > \max\{C, R\}$, write $U = R + Mm(G)$, with m independent of D. Choosing a rational a in this interval identifies m locally with m_a. The envelope identity gives $H_G(d_*, G) = m'(G)$ and the same derivative for every maximizing target. The right endpoint satisfies $d_* < \bar{D}$ and $H_D(d_*, G) = 0$. Thus the limits from within the racing interval satisfy $Q(d_*, G) = 0$ and $U_D(d_*, G) = C_D(d_*, G) \leq 0$.

Backward integration of $bU_{DD} = (r - c + \kappa D)U_D + \kappa U$ gives $U_D \leq 0$ throughout the interval. Moreover, $b\partial_D(Q/M) = [(1+\eta)\kappa U - \alpha]/M$. Its numerator is nonincreasing, so Q/M is minimized at an endpoint. If $d_- > 0$, define the fixed-target value $U_*(D, G) = R(D) + M(D)H(d_*, G)$. Persistence of accommodation gives, for sufficiently small $t \geq 0$,

$$U_*(d_- - pt, G - t) \leq C(d_- - pt, G - t), \qquad U_*(d_-, G) = C(d_-, G).$$

Right differentiation at zero yields $-p(U_*)_D - (U_*)_G \leq -pC_D - C_G$. The waiting equation for C therefore gives a nonpositive waiting residual for U_* at (d_-, G). Its racing residual is zero. By the envelope identity, its derivatives are precisely the limiting derivatives of U from the racing interval, so $Q(d_-, G) \geq 0$. If $d_- = 0$, the corresponding limiting coefficient satisfies $bQ(0+, G) = -(1+\eta)(1 - wU_G(0, G) - BU(0, G)) \geq 0$ by the safety residual. Hence $Q \geq 0$ throughout the interval. Fubini's theorem gives both residual inequalities almost everywhere on $D > 0, G > 0$. Below safety, $Q(D, G) = \alpha(1-z)/(r - c) + zQ(0+, G) \geq 0$.

These inequalities verify the leader's incentive constraints. Approximate U by smooth local averages U_ε. On compact subsets with $D \neq 0$ and $G > 0$, both payoff residuals evaluated at U_ε are bounded above by a constant times ε. Integrate the smoothed inequalities along a control path and pass to the limit. At almost every date on safety the leader's

speed is h, so the safety residual applies; at a tie the scalar bound applies. The bound $U \leq 1/(r-c)$ makes the discounted terminal term vanish before catch-up, since the follower races and the leader's profit grows at most at rate c. Persistence of accommodation shows that the selected stopping path attains U.

It remains to verify the follower's incentive constraints. Normalize its payoff by e^{cA_L}, and let $F_C(D,G)$ denote its value under immediate accommodation and racing until catch-up. On safety,

$$F_C(0,G) = e^{-\alpha G}\int_0^{G/w} e^{-(r-\alpha+\beta h)t}\,\mathrm{d}t + e^{-(r-ch)G/w}P(0).$$

Above safety, $e^{-\alpha G} - p(F_C)_D - (F_C)_G - (r+\kappa D)F_C = 0$ and $F_C(D,0) = P(D)$. The first derivatives agree where safety and catch-up occur simultaneously. Every continuation from (D,G) is bounded by $e^{-\alpha G}/(r-\alpha)$, since own capability grows at most at speed one and rival capability cannot fall.

Define $H_F = \eta(F_C)_D - (F_C)_G$ and $K_F = (1+\eta)(F_C)_D + (F_C)_G + cF_C$. Along accommodation,

$$\dot{H}_F = (r+\kappa D)H_F + \eta\kappa F_C - \alpha e^{-\alpha G},$$
$$\dot{K}_F = (r+\kappa D)K_F + (1+\eta)\kappa F_C + \beta e^{-\alpha G}.$$

The first bound in (ES) makes the forcing term in the equation for H_F nonpositive. At a tie, $(F_C)_G = \eta P'$, so $H_F = 0$ and $K_F = kP' + cP \leq 0$. On safety, $K_F = 0$ and $(1+\eta)H_F = [\alpha + \eta(\alpha+\beta)]e^{-\alpha G}\int_0^{G/w} e^{-(r-\alpha+\beta h)t}\,\mathrm{d}t \geq 0$. Backward integration yields $H_F \geq 0$ and $K_F \leq 0$ throughout the accommodation region.

For a fixed finite target $d \in [\max\{D,0\}, \bar{D}]$, define

$$F^d(D,G) = e^{-\alpha G}R(D) + \frac{M(D)}{M(d)}\left[F_C(d,G) - e^{-\alpha G}R(d)\right],$$

using $M(D) = e^{(r-c)D/b}$ below safety. The waiting equation implies that $\partial_d F^d$ has the sign of $K_F(d,G)$, so $\partial_d F^d \leq 0$. Immediate accommodation also weakly dominates perpetual racing. Until catch-up, the follower has the same capability, a weakly lower rival capability, and weakly lower exposure; at catch-up, $P \geq R$ and its distance from safety is weakly lower than under joint racing. Hence $F_C(D,G) \geq e^{-\alpha G}R(D)$. With $F^\infty = e^{-\alpha G}R$, the follower's value is therefore nonincreasing in its target, including infinity.

For a fixed finite target, put $L^d = \eta F^d_D - F^d_G$. At that target, $L^d(d-,G) = H_F(d,G) - \eta K_F(d,G)/b \geq 0$. Before the target, $bL^d_D = (r-c+\kappa D)L^d + \eta\kappa F^d - \alpha e^{-\alpha G}$. The payoff bound and the first restriction in (ES) make the forcing term nonpositive; backward

integration gives $L^d \geq 0$. Below safety the same equation holds without the term $\eta\kappa F^d$. For target infinity, $L^\infty = e^{-\alpha G}(\alpha R + \eta R') \geq 0$ by the bound on $-R'/R$ and (ES). Thus, while the leader races, the residual at follower speed a_F is $(a_F - 1)L^d \leq 0$ for every fixed target.

Let $d(D, G)$ be the leader's selected target and let $W(D, G)$ be the resulting normalized follower payoff: it equals F_C at stopping states and $F^{d(D,G)}$ elsewhere. During any racing segment of a follower deviation, $\dot{G} = 1 - a_F \geq 0$ and $\dot{D} = b - \eta(1 - a_F) \geq 0$, using $\eta \leq b$. Thus $d(D_t, G_t)$ is nondecreasing. The bounded-variation chain rule applied to $F^{d(D_t,G_t)}(D_t, G_t)$ gives a nonpositive contribution from continuous target changes and from target jumps, in addition to the nonpositive fixed-target residual. For this purpose, infinity can be represented by an isolated point above $\bar{D}$; crossing to it changes the value to F^∞ and has a nonpositive jump.

At a stopping state the follower's value is F_C. Its residual at any admissible pair (a_L, a_F) is $a_L K_F + (a_F - 1)H_F \leq 0$, which also covers regularized mixtures at the boundary. A transition from stopping to racing cannot raise the value, because $F_C \geq F^d$ for every available target. A transition from racing to stopping is continuous: monotonicity of the target in both state coordinates implies $d(D_s, G_s) \leq d(D_t, G_t) = D_t$ for racing dates $s < t$ preceding the stopping date t, while $d(D_s, G_s) \geq D_s$. Thus the target converges to D_t and $F^{d(D_s,G_s)}(D_s, G_s)$ converges to $F_C(D_t, G_t)$. On safety, $K_F = 0$ and $H_F \geq 0$ give the same residual inequality for every speed pair that keeps $D = 0$.

To include infinitely many switches, put $W_t = W(D_t, G_t)$. On a compact time interval before catch-up, the state derivatives of F_C, F^∞, and all finite-target values are uniformly bounded. Since target changes on a racing interval only lower the value, $W_t - W_s \leq L(t - s)$ there for some common L. The same bound holds between any two stopping dates, where W agrees with F_C. For arbitrary $s < t$, join these comparisons at the first and last intervening stopping dates; the boundary comparisons introduce no positive jumps. Thus $W_t - Lt$ is nonincreasing, so W has bounded variation and its singular derivative is nonpositive. Its absolutely continuous derivative satisfies the residual inequality almost everywhere, including on the stopping set where $W = F_C$. Writing $\Gamma_t = e^{-rt-\Lambda_t+c(A_{Lt}-A_{L0})}$, integration gives $\Gamma_T W_T + \int_0^T \Gamma_t e^{-\alpha G_t}\,\mathrm{d}t \leq W_0$.

The follower value extends continuously to ties. Indeed, $H(d, 0) = [P(d) - R(d)]/M(d)$ is strictly decreasing on $[0, \bar{D}]$, because $H_D(d, 0) = (kP'(d) + cP(d))/(bM(d)) < 0$ for $d > 0$. For $D < \bar{D}$, the gain at the smallest feasible target $\max\{D, 0\}$ is positive and uniquely maximal when $G = 0$. Continuity and compactness therefore imply that selected targets converge to $\max\{D, 0\}$ as $(D', G) \to (D, 0)$ with $G > 0$. Substitution in F^d gives

$\max\{P(D), R(D)\}$ when $D \geq 0$, and the value of racing to safety and then pacing when $D < 0$. At $D = \bar{D}$, finite targets converge to $\bar{D}$, whereas an infinite target gives $R(\bar{D})$; the limits agree because $P(\bar{D}) = R(\bar{D})$. For $D > \bar{D}$, the target is infinity in a neighborhood. Thus $W(D, G)$ converges to the scalar tied value at every tie. Letting T approach catch-up in the payoff inequality and applying the scalar bound verifies the continuation after catch-up. If catch-up never occurs, the terminal term vanishes, since it is bounded by the discounted current follower profit divided by $r - \alpha$, and $r > \alpha$.

Thus every follower deviation is bounded by racing until catch-up and then following the tied policy. The prescribed path attains both firms' bounds from every state, so Lemma 4 establishes the pacing MPE. Its outcome from a safety tie is perpetual joint pacing, proving the the sufficiency of the condition in (ii).

For necessity, a firm that races from a safety tie remains a weak leader, and safety advances at speed at least g, so its distance above safety is at most $(1 - g)t$ and its flow is at least $e^{c(A_0+t)}$. It therefore guarantees $R(0)$, so pacing requires $P(0) \geq R(0)$. □

ar