Discussion Paper Series – CRC TR 224

Discussion Paper No. 648
Project B 02

Strategic Concealment in Innovation Races

Yonggyun Kim[1]
Francisco Poggi[2]

September 2026
(First Version : February 2025)

[1]Florida State University, ykim22@fsu.edu
[2]University of Mannheim, poggi@uni-mannheim.de

Support by the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation) through CRC TR 224 is gratefully acknowledged.

Strategic Concealment in Innovation Races*

Yonggyun Kim† **Francisco Poggi**‡

July 7, 2026

Abstract

Firms racing to innovate often make interim breakthroughs that speed up, but are not necessary for, the final innovation. When such a breakthrough is privately acquired, a firm can disclose it by filing a patent, or conceal it. This paper studies how this trade-off is shaped by the race structure and the intellectual property system. We develop a dynamic model where firms allocate resources between developing with an existing technology and researching a faster one. We show that firms strategically conceal their breakthroughs when the prize for winning is large and the prior-use defense is strong, impeding knowledge spillovers and slowing the pace of innovation.

JEL Classification: C73, D21, O30

Keywords: Interim Technology, Patents, License, Prior-use Defense, Direction of Innovation

*We are indebted to Curtis Taylor for his constant encouragement and support. Kim thanks him, Arjada Bardhi and Fei Li for their excellent guidance. We are also grateful to Attila Ambrus, Jim Anton, Raphael Boleslavsky, Jaden Chen, Xiaoyu Cheng, Jeffrey Ely, Felix Zhiyu Feng, Ilwoo Hwang, Alan Jaske, Dongyoung Damien Kim, Seung Joo Lee, Jorge Lemus, David McAdams, Philipp Sadowski, Todd Sarver, Johannes Schneider, Ludvig Sinander, Yangbo Song, Konrad Stahl, Can Tian, Zichang Wang, Huseyin Yildirim, Mofei Zhao, Dihan Zou and seminar participants at CUHK Business School, CUHK Shenzhen, Duke University, FSU, KDI, National Taiwan University, University of Copenhagen, UNC Chapel Hill, University of Mannheim, University of Queensland, the 33rd Stony Brook International Conference on Game Theory, SEA 92nd Annual Meeting, 2023 Midwest Theory Conference, and KES-Micro Seminar for comments and suggestions. All remaining errors are our own. Support by the German Research Foundation (DFG) through CRC TR 224 (Project B02) is gratefully acknowledged.

†Department of Economics, Florida State University. Email: ykim22@fsu.edu

‡Department of Economics, University of Mannheim. Email: poggi@uni-mannheim.de

1 Introduction

In many innovation races, the path to the finish line is not unique: often, firms can rely on established technology, or invest instead in research that, if successful, delivers a faster new technology. *Interim technological breakthroughs*—partial advances that do not themselves complete the innovation but sharply accelerate its remaining development—are central to how these races unfold. The race to develop COVID-19 vaccines illustrates the two paths: Oxford–AstraZeneca and Janssen (Johnson & Johnson) relied on the established viral-vector technology, whereas Moderna and Pfizer–BioNTech invested on messenger RNA (mRNA), a faster but, at the outset, unavailable technology.[1] Similar dynamics appear in other industries, such as battery chemistry reshaping electric-vehicle development, and new software frameworks speeding up the delivery of finished products.

Two features make these races distinctive. First, firms allocate scarce resources across *multiple pathways*, trading off the slow but reliable progress of an existing technology against the faster development that an interim breakthrough would unlock. Second, a firm that *privately* achieves such a breakthrough need not reveal it: it can **patent** the discovery—disclosing it and opening the door to licensing—or **conceal** it. Together, these features create a margin that is absent from standard innovation races: whether and when the frontier technology *diffuses* to a rival, and hence whether the industry squanders effort by either duplicating research or developing the final product with an inferior technology.

This margin motivates the central question of our paper: what institutional arrangements lead firms to share their interim discoveries, generating the knowledge spillovers that curb duplicated research and inefficient technology adoption and thereby accelerate innovation? Because a privately held breakthrough does not diffuse on its own, the answer turns on how the intellectual property system—patents, licensing, and the prior-use defense—shapes the private incentive to disclose, and on when that incentive aligns with the social value of diffusion.

[1] Viral-vector vaccines had already been deployed against earlier outbreaks, including the recombinant vesicular stomatitis virus (rVSV-ZEBOV) vaccine used during the 2014–2016 West African Ebola epidemic (Henao-Restrepo et al., 2017). By contrast, mRNA technology was not yet in practical use before COVID-19: its application had long been restricted by the instability and inefficient *in vivo* delivery of mRNA, hurdles that had to be overcome—for instance, protecting the mRNA from degradation during delivery—before a vaccine could be developed with it (Pardi et al., 2018).

High-level Intuition To fix ideas, consider two firms, A and B, competing to innovate. Firm A has privately obtained a new technology that accelerates development, while Firm B remains uncertain whether A possesses it. Firm A must decide whether to **patent** the technology ($a = P$) or **conceal** it ($a = C$). Let V_a^i denote Firm i's expected payoff given A's choice a, and let $W_a = V_a^A + V_a^B$ denote the total expected payoff. If A patents, it may offer to **license** the technology to B; let W_L denote the total payoff under licensing. Assuming A holds full bargaining power, its payoff after licensing is $W_L - V_P^B$. Figure 1 illustrates this decision.

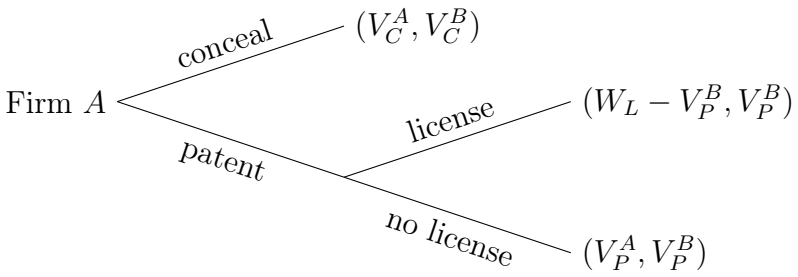


Figure 1: Patenting Decision and Payoffs

Sharing the technology can raise overall welfare by enabling faster innovation through knowledge spillover, so that $W_L > \max\{W_P, W_C\}$. Patenting, however, also reveals to Firm B that A holds the new technology, potentially potentially prompting B to adjust its own R&D, in doing so, raise its outside option to $V_P^B > V_C^B$, which weakens A's bargaining position. Firm A patents only if the net gain from patenting and licensing, $W_L - V_P^B$, exceeds the value of concealment, $V_C^A = W_C - V_C^B$, and this comparison is not clear-cut. A **hold-up problem** therefore arises: even though patenting and licensing would maximize social welfare, Firm A may withhold the discovery because disclosure strengthens its rival and erodes its own payoff.[2] The remainder of the paper formalizes this tension within a dynamic race and asks when the patent system overcomes it.

[2]This mechanism parallels the *feedback effects* studied in financial markets, where revealing information induces a real response that in turn feeds back on the value of revealing it. In Edmans et al. (2015), trading on private information leads a manager to re-optimize a real decision, which can render acting on the information self-defeating and endogenously limit arbitrage; here, patenting reveals the breakthrough to the rival, who reallocates its R&D and erodes the very licensing rent the patent was meant to capture. Boleslavsky et al. (2017) study a related setting with a policy dimension close to ours, in which an authority conditions a costly intervention on a market signal that informed agents strategically distort in anticipation. We thank an anonymous referee for pointing out this connection.

Race Setup We study a continuous-time model in which two firms race to develop a new product, such as a vaccine, and the first to succeed earns a fixed reward Π (e.g., temporary monopoly profits). Each firm continuously divides its limited resources between two pathways. It can pursue *direct development* with the currently *available technology*—developing the product at a baseline rate λ_L. Alternatively, it can devote resources to *research*, discovering a faster new technology at rate μ; once discovered, the firm develops the product at the higher rate $\lambda_H > \lambda_L$. A firm that discovers the new technology then decides whether to patent it or keep it secret. This dual-path structure embeds the adoption margin at the heart of our analysis: catching up through research duplicates effort, while developing with the old technology means producing at an inferior frontier.

Benchmark: No Technological Spillovers We begin by abstracting from the patent system and shutting down technological spillovers, as would be the case if the new technology were non-patentable and discoveries could not be licensed. This benchmark isolates the resource-allocation trade-off at the heart of the race: shifting resources toward research slows near-term development but raises the chance of acquiring the superior technology, and with it the expected future development rate.

How this trade-off plays out depends on what firms observe. When research progress is *public*, a firm that knows it's lagging behind faces a choice. It can aim to *catch up* by continuing to research the same technology, in which case there is duplication of efforts, or it can *fall back* to direct development with the old technology, developing at the inferior rate. Either response is inefficient relative to the first best, which would have the breakthrough spillover instantly (Proposition 1). When research progress is *private*, the lagging firm cannot condition on its rival's discovery, so it loses the option to fall back at the right moment. Firms keep researching instead, which deepens duplication and slows the aggregate pace of innovation relative to the public benchmark (Proposition 2). In both cases the absence of spillovers leaves the industry strictly short of the first best, and the question is whether the patent system can close the gap.

Patents and Licensing We then introduce patents. Patents address both inefficiencies at once. By assigning an exclusive right that can be licensed, a patent enables *controlled spillovers*: the discovering firm can transfer the frontier technology to its rival for a fee. And because patents are *public*, filing one transmits the very information that the public-progress benchmark relied on, allowing a rival to fall back appropriately. Patenting thus restores both the diffusion channel and the information channel that private concealment had shut down.

We model three institutional features. A patent grants its holder an exclusive right, but a rival who has independently discovered the same technology can mount a *prior-use defense*, which succeeds with probability β; if it succeeds, both firms may use the technology (cf. 35 U.S.C. §273).[3] Licensing takes a take-it-or-leave-it form with all bargaining power assigned to the patent holder. We deliberately strip away other frictions—there is no information leakage beyond the patent itself and no bargaining inefficiency—so that any failure to diffuse is driven purely by the strategic hold-up identified above.

Within this framework, we first show that the act of disclosure is never the source of inefficiency. Conditional on a patent being filed, the property right is well defined and the parties always reach a licensing agreement that diffuses the technology, with the fee extracting the rival's surplus—a Coasean outcome in which post-disclosure bargaining is efficient (Lemma 1). Inefficiency can therefore reside only in the *decision* to patent.

Our first main result characterizes when efficient disclosure is sustainable. Immediate patenting upon discovery—which implements the first best—is an equilibrium whenever research progress is public, and, when it is private, if and only if the value of licensing is large relative to the value of protecting one's lead (Theorem 1). Translating this condition into primitives, efficiency obtains when the normalized stakes of the race are not too high, or when the prior-use defense β is not too strong. Intuitively, high stakes make a protected lead more valuable than licensing rents, and a strong prior-use defense erodes the security of a patent; both sharpen the hold-up problem and tempt firms toward secrecy.

When efficiency fails—private progress, a high prize, and a strong prior-use defense—we fully characterize the unique symmetric equilibrium (Theorem 2). As the relative value of

[3] A rival who has not filed cannot claim exclusive rights even if it discovered the technology earlier, since trade-secret protection does not extend to independent discoveries by competitors.

licensing falls, equilibrium disclosure passes through three regimes: *immediate patenting*, in which firms disclose at once; *delayed mixing*, in which a firm conceals its breakthrough for a while and then patents gradually, balancing the option value of future licensing against the rising risk of preemption; and *full concealment*, in which firms never patent and keep the rival on the slower research path. The characterization shows that the strategic incentive to conceal emerges precisely when the prize is large and prior-use protection is strong, and that the inefficiency takes the concrete form of suppressed spillovers and prolonged, duplicative racing.

Empirical and Policy Implications These results deliver sharp comparative statics. When the prior-use defense is strong, higher stakes monotonically erode the incentive to disclose: at low-to-moderate stakes firms patent their breakthroughs immediately, sustaining perfect spillovers and the fastest rate of innovation, whereas at high stakes they retreat into concealment, eliminating spillovers and slowing cumulative progress. When the prior-use defense is weak, this force vanishes and immediate disclosure remains the unique equilibrium however high the stakes. Because the de facto enforceability of prior-use defenses varies across industries—readily documented in manufacturing but ambiguous in software—these predictions are testable by comparing disclosure timing and innovation speed across sectors with different protections and market stakes.

On the policy side, the analysis cautions that levers meant to spur innovation along one margin can suppress it along another. A larger product-market prize encourages initial R&D, yet beyond a point it tips firms from disclosure into concealment, choking off the spillovers that accelerate innovation; a stronger prior-use defense has the same chilling effect on disclosure. The two levers, however, are not equally available to a policymaker. The prior-use defense embodies commitments that lie outside our model—protecting independent inventors, deterring patent trolls, and curbing wasteful defensive patenting—so its strength is better treated as fixed by these broader objectives than tuned to encourage disclosure. The stakes of the race, by contrast, policy can shape: by compressing the gap between winning and losing—taxing the winner and partly compensating the laggard—a regulator can restore immediate disclosure and efficient diffusion even under a strong prior-use regime, provided

the prize remains large enough to sustain participation.

Roadmap Section 2 sets up the model and the first-best benchmark. Section 3 analyzes the no-spillover benchmark under public and private progress. Section 4 introduces patents and licensing and characterizes when they restore efficiency. Section 5 characterizes equilibrium concealment and disclosure when efficiency fails, and reinterprets the results in terms of the stakes and the prior-use defense. Section 6 discusses empirical and policy implications and concludes. The Related Literature is discussed below.

1.1 Related Literature

This paper contributes to the economics of innovation and intellectual property, and in particular to the literature on firms' choice between patenting and secrecy. Our central contribution is to make the cost of disclosing an interim breakthrough *endogenous* to the patent system and to the structure of the race: the decision to patent, and the knowledge spillovers it governs, respond to interpretable institutional primitives—the strength of the prior-use defense and the size of the product-market prize—yielding predictions that bear directly on innovation-policy debates.

Patenting versus Secrecy We contribute most directly to the literature studying the trade-off between patenting and secrecy, to which we add a novel motive for concealment: preventing a rival's *strategic response*.[4] Much of this work emphasizes the limitations of patent protection. The seminal article of Horstmann et al. (1985) argues that patent coverage may fail to exclude profitable imitation, so that firms may prefer secrecy to avoid being copied.[5] A second concern is that patents are time-limited: Denicolò and Franzoni (2004) model a setting in which a patent confers temporary monopoly power, while secrecy offers potentially indefinite protection that may nonetheless be lost to leakage or duplication. Our mechanism is distinct from both. A firm conceals not to avoid imitation or to prolong a

[4]For a comprehensive review of the literature on patents versus secrecy, see the survey by Hall et al. (2014).

[5]Many subsequent papers study the imitation threat and patent infringement, e.g., Gallini (1992); Takalo (1998); Anton and Yao (2004); Kultti et al. (2007); Kwon (2012); Zhang (2012); Krasteva (2014); Krasteva et al. (2020).

monopoly term, but to withhold the information its rival would need to re-optimize—keeping the rival's outside option low and its own licensing rents high. Because this motive operates through the rival's equilibrium R&D rather than through a patent's coverage or duration, its strength is shaped by the market and institutional environment, which is precisely what makes it responsive to policy.

Disclosure in Innovation Races Our analysis relates to a growing literature on information disclosure in innovation races (Hopenhayn and Squintani, 2016; Bobtcheff et al., 2017; Curello, 2023).[6] In these models, the value of an innovation grows over time, and firms trade off disclosing early to secure priority against waiting to enhance value. Our model instead holds the value of the final innovation fixed; disclosure matters only because it informs the rival's R&D, so a firm may delay it purely to confound the rival's allocation decisions. Closest to our paper is Chatterjee et al. (2026), who likewise study the decision to disclose intermediate breakthroughs and find, as we do, that large rewards to the final innovation can lead firms to withhold interim findings. The two papers differ in three respects. First, in their setting a disclosed discovery can be freely used by the rival. In ours, disclosure takes the form of a patent application: the patent holder can exclude the rival—subject to the prior-use defense—and license on its own terms. Second, in their model the intermediate breakthrough is the only route to the final innovation. We add a second, always-available path: direct development with the old technology. Third, in their model disclosure carries an exogenous intrinsic value (prestige, priority, a prize). In ours the value of disclosing is purely instrumental and endogenous: the expected licensing fee, determined within the race and by institutional features of the patent system. It is this endogeneity that allows our framework to deliver comparative statics in policy parameters.

Patent Institutions and Licensing Because disclosure in our model operates through patenting and licensing, our results speak to the design of patent institutions. The strength of the prior-use defense and the choice between first-to-file and first-to-invent regimes are central policy levers in our analysis; in particular, we offer a rationale for first-to-file as a disclosure-

[6] See also Lichtman et al. (2000); Baker and Mezzetti (2005); Gill (2008); Baker et al. (2011); Ponce (2011), reviewed in Section 3.3 of Hall et al. (2014).

promoting institution, complementing work comparing these regimes (Miyagiwa, 2015). The licensing channel connects us to the literature on licensing interim technologies (Bhattacharya et al., 1992; d'Aspremont et al., 2000; Bhattacharya and Guriev, 2006; Spiegel, 2008), which typically takes the identity of the technology holder—and hence of the licensor—as given. By contrast, in our model both the decision to license and the identity of the licensor are equilibrium outcomes, shaped by what each firm does and knows about its rival's progress.

Methodological Lineage Methodologically, we build on models of multi-stage innovation with Poisson breakthroughs (Scotchmer and Green, 1990; Denicolò, 2000; Green and Taylor, 2016; Song and Zhao, 2021; Kocourek and Kováč, 2023) and on races in which firms pursue multiple R&D pathways (Das and Klein, 2024; Akcigit and Liu, 2016; Bryan and Lemus, 2017). For instance, Kocourek and Kováč (2023) study a two-stage race in which firms choose arrival intensities under public or private information about an intermediate breakthrough; relative to such models, we add technological spillovers and a slower direct path that bypasses the breakthrough entirely. Our dual-path structure also connects to Carnehl and Schneider (2023) and Kim (2022), where agents choose between sequential and direct routes; unlike those single-agent or principal–agent settings, the competitive race itself induces firms to adopt the established technology, and we impose no exogenous deadline. Across these literatures, our point of departure is that interim discoveries are privately observed and that their licensing is an equilibrium object.

2 Model

2.1 Setup

We model a continuous-time ($t \in [0, \infty)$) innovation race between two risk-neutral, non-discounting firms, A and B. They compete to develop a new product; the first to successfully develop the product wins a prize Π.[7] Throughout the race, each firm incurs a constant flow cost $c > 0$.

[7] This winner-take-all payoff structure has been commonly used in the innovation race literature, e.g., Loury (1979); Lee and Wilde (1980); Denicolò and Franzoni (2010).

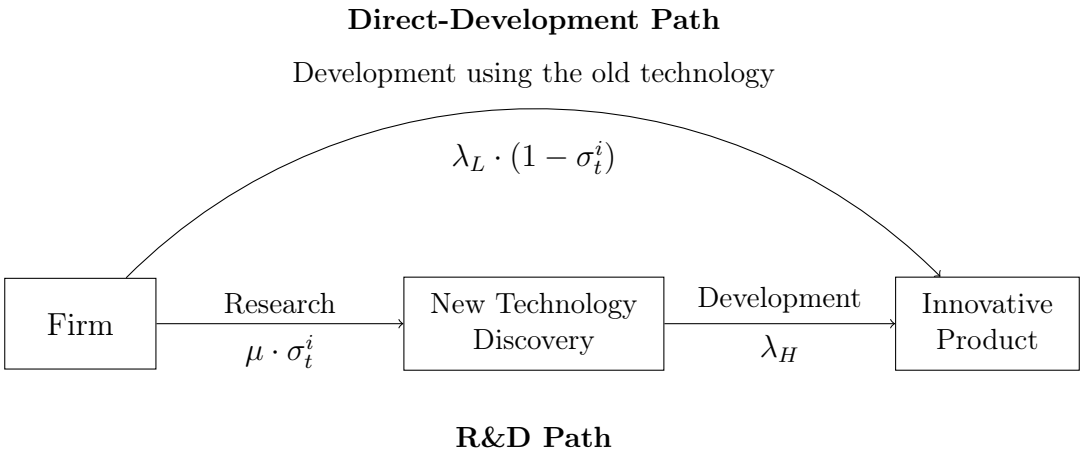


Figure 2: Development Paths and Resource Allocations

Technology and Resource Allocation Initially, firms possess an **old technology** that allows product development at a maximum rate λ_L. They can also attempt to discover a **new technology**, which enables a faster development rate $\lambda_H > \lambda_L$.

A firm that has not yet discovered the new technology must allocate a single unit of resources between two activities:

1. **R&D Path:** Researching the new technology. If Firm i allocates a fraction $\sigma_t^i \in [0, 1]$ of its resources to research, it discovers the new technology at a Poisson rate $\mu\sigma_t^i$.

2. **Direct-Development Path:** Developing the product using the old technology. The remaining fraction, $1 - \sigma_t^i$, yields product development according to an independent Poisson process with rate $\lambda_L(1 - \sigma_t^i)$.

Once a firm acquires the new technology, it ceases research and dedicates all resources to product development ($\sigma_t^i = 0$), developing at the maximum rate λ_H. Figure 2 illustrates the development paths and resource allocation framework.

Benchmarks: No Spillover For our benchmark analysis, we focus on the case with no technology spillovers across firms: the only way for a firm to access the new technology is to discover it independently. Our primary goal is to characterize the socially optimal resource allocation without spillovers, and show that it is implementable as an equilibrium of the race game.

We consider two informational setups regarding the observability of the rival's discovery of the new technology. Formally, let $N_t^i \in \{0, 1\}$ denote the state of Firm i (with $N_0^i = 0$), where $N_t^i = 1$ indicates the firm has discovered the new technology. While a firm always observes its own state, its knowledge of the rival's state depends on the information setting. We analyze two cases: *public progress*, where N_t^j is observable, and *private progress*, where it is not. These cases are analyzed in Section 3.

Main Analysis: Spillover through Licensing In our main analysis, we introduce the possibility of technology spillover through patenting and licensing. A firm that discovers the new technology can choose to file for a patent.

We assume that if the rival has already independently discovered the technology (and has not yet patented it), it can challenge the patent application. This challenge succeeds with probability $\beta \in [0, 1]$, a policy parameter capturing the strength of intellectual property right protection via a **prior-use defense**.

If a firm successfully secures an exclusive patent, it obtains monopoly control over the new technology. The patent-holding firm can license the technology to its rival via a take-it-or-leave-it offer.[8] The formal game, including the strategic timing of patent filing and the ensuing subgame, is detailed in Section 4.

2.2 Parametric Assumptions

We impose two assumptions on the model parameters for the subsequent analysis.

Assumption 1: Viability of Direct Development We assume that the prize for winning is large enough to justify the development costs, even when firms rely exclusively on the old technology:

$$\Pi - \frac{c}{\lambda_L} > 0 \tag{2.1}$$

This condition ensures that firms have no incentive to drop out of the race: staying active for an additional dt using the old technology generates an expected profit of $\lambda_L \Pi dt$ at a cost of

[8] The firm can effectively exclude the rival by demanding a prohibitively high licensing fee, though we show later that this never occurs in equilibrium.

cdt, which yields a net positive expected payoff. If this assumption fails, the old technology is never utilized, and the model collapses to a single, sequential innovation path.

Equivalently, defining the *normalized stakes* of the race $\pi := \frac{\lambda_L \Pi}{c}$—the ratio of the expected flow benefit $\lambda_L \Pi$ to the flow cost c of remaining in the race via direct development—Assumption 1 reads $\pi > 1$. We use π as the central primitive in the equilibrium characterization and policy analysis later in the paper (Section 5 and 6).

Assumption 2: Superiority of the R&D path We assume the following inequality:

$$\underbrace{\frac{1}{2\mu} + \frac{1}{\lambda_H}}_{\text{LHS}} \; < \; \underbrace{\frac{1}{2\lambda_L}}_{\text{RHS}}. \tag{2.2}$$

The two sides of the inequality represent the expected development times under two extreme scenarios: (i) LHS is the expected development time in a world where only the R&D path is available, and where only the first firm to discover the new technology is permitted to use it; and (ii) RHS is the expected development time in a world where only the direct-development path is available to both firms.

This assumption implies that the new technology is so efficient that racing to discover it—even if the winner finishes alone—is faster than both firms relying on the old technology. Note that this is a *stronger* condition than assuming the R&D path beats direct development for a single, isolated firm, i.e., $\frac{1}{\mu} + \frac{1}{\lambda_H} < \frac{1}{\lambda_L}$.

2.3 The First-Best Outcome

As a benchmark for efficiency, consider a social planner who dictates the firms' actions to maximize their joint expected payoff. Because the identity of the race winner is irrelevant for total surplus, maximizing the joint payoff is equivalent to minimizing the expected time until the product is developed. The planner's problem therefore has two parts: how to allocate resources before the new technology is discovered, and what to do once it is.

Once a firm discovers the new technology, efficiency requires an immediate **technology spillover**: sharing the discovery at once lets both firms develop at the faster rate λ_H, with

no further duplicative research or reliance on the old technology. Before any discovery, Assumption (2.2) makes the R&D path socially optimal: the expected development time is strictly decreasing in the common research share, since the assumption implies $\frac{1}{\mu} + \frac{1}{\lambda_H} < \frac{1}{\lambda_L}$. The planner thus directs both firms to devote all resources to research and, under this policy of perfect, instantaneous spillovers, the first-best expected development time is $\frac{1}{2\mu} + \frac{1}{2\lambda_H}$.

3 Benchmark: No Technological Spillovers

In this section, we shut down the technological spillover channel to analyze equilibrium behavior under public and private research progress. The absence of spillovers naturally creates an efficiency gap relative to the first best, primarily by forcing the lagging firm to either duplicate research or rely on sub-frontier technology.

Beyond this direct technological friction, shutting down spillovers introduces a secondary, strategic source of inefficiency: the potential for firms to misallocate their resources depending on what they can observe about their rival's progress.

3.1 Public Research Progress

When firms observe their rival's research progress, they can condition resource allocation on the *public history* $H_t := \left\{(N_s^A, N_s^B)\right\}_{s\in[0,t]}$. To simplify the exposition, we focus on Markov strategies, $\mathbf{s}_i : \{0,1\} \times \{0,1\} \to [0,1]$, where $\mathbf{s}_i(N^A, N^B)$ denotes Firm i's resource allocation toward research given the current state of progress (N^A, N^B). Also note that $\mathbf{s}_i = 0$ whenever $N^i = 1$—Firm i develops the product with the new technology when it discovers it.

We define two natural benchmark strategies:

1. **Research strategy ($\mathbf{s}^R$):** A firm researches until it successfully obtains the new technology. Formally, $\mathbf{s}_i^R = 1$ when $N^i = 0$, and $\mathbf{s}_i^R = 0$ when $N^i = 1$.

2. **Fall-back strategy ($\mathbf{s}^F$):** A firm researches only until *either* firm obtains the new technology. Formally, $\mathbf{s}_i^F = 1$ when $N^A = N^B = 0$, and $\mathbf{s}_i^F = 0$ otherwise.

To understand the trade-offs driving these strategies, suppose one firm discovers the new technology while the other has not. In the first-best scenario, both firms immediately develop

with the new technology, yielding an expected remaining time of $\frac{1}{2\lambda_H}$. Because there are no spillovers, the lagging firm's response necessarily generates one of two inefficiencies relative to the first best:

- **Duplication of effort:** If the lagging firm continues researching to catch up ($\mathbf{s}^R$), the expected remaining time is $T = \frac{1}{\mu+\lambda_H} + \frac{\mu}{\mu+\lambda_H} \cdot \frac{1}{2\lambda_H} = \frac{1}{2\lambda_H} + \frac{1}{2(\lambda_H+\mu)}$.
- **Sub-frontier development:** If the lagging firm develops with the old technology ($\mathbf{s}^F$), the expected remaining time is $T = \frac{1}{\lambda_H+\lambda_L} = \frac{1}{2\lambda_H} + \frac{\lambda_H-\lambda_L}{2\lambda_H(\lambda_H+\lambda_L)}$.

That is, the research strategy extends the duration by $\frac{1}{2(\lambda_H+\mu)}$ due to duplicative research effort, while the fall-back strategy extends the duration by $\frac{\lambda_H-\lambda_L}{2\lambda_H(\lambda_H+\lambda_L)}$ due to sub-frontier development.

Definition 1. We say that the *fall-back condition* holds if the fall-back strategy extends the expected duration less than the research strategy; that is, if:

$$\frac{\lambda_H - \lambda_L}{2\lambda_H(\lambda_H + \lambda_L)} < \frac{1}{2(\lambda_H + \mu)}. \tag{3.1}$$

Note that this condition is equivalent to:

$$\frac{1}{\mu} > \frac{1}{2\lambda_L} - \frac{1}{2\lambda_H} \quad \Longleftrightarrow \quad \lambda_H > \lambda_\star := \mu\lambda_H \left(\frac{1}{\lambda_L} - \frac{1}{\mu} - \frac{1}{\lambda_H}\right). \tag{3.2}$$

In Online Appendix OA.1, we analyze the case of a single firm facing an exogenous termination rate λ. We show that such a firm optimally chooses the R&D path when $\lambda < \lambda_\star$, and the direct development path when $\lambda > \lambda_\star$.[9] Consequently, once a firm discovers the new technology and develops at rate λ_H, its rival faces a termination rate of λ_H; when this exceeds the threshold $\lambda_\star$—that is, when the fall-back condition (3.2) holds—the rival's best response is to develop directly rather than to keep researching.

The following proposition formalizes this observation and shows that the strategy profile minimizing the expected duration constitutes a Markov Perfect equilibrium (MPE).

[9] Given this definition, Assumption (2.2) is equivalent to $\mu < \lambda_\star$.

Proposition 1. *Suppose research progress is public, there are no technological spillovers, and the assumptions in Section 2.2 hold.*

(a) ***Efficiency:*** *The expected duration of the race is minimized when both firms adopt the fall-back strategy ($\mathbf{s}^F$) if (3.1) holds, and the research strategy ($\mathbf{s}^R$) otherwise.*

(b) ***Equilibrium:*** *This duration-minimizing strategy profile constitutes a MPE.*

This proposition identifies the mechanical inefficiency in the absence of technological spillovers: the lagging firm must either duplicate the research effort or fall back to the inferior technology. Beyond this mechanical gap, there is no strategic misallocation of resources. Intuitively, the winner-take-all structure is such that each firm acts to maximize its own payoff that depends on the probability of winning but, given the technological constraints it faces, this leads it to exactly the allocation a duration-minimizing planner would choose.

3.2 Private Research Progress

We now consider the case where firms cannot observe their rivals' research progress. Therefore, a firm can only condition its resource allocation on its own progress. Since Firm i will not allocate its resources to research once the new technology is discovered, its strategy comes down to choosing the pre-discovery allocation $\boldsymbol{\sigma}_i \in [0,1]$.[10]

We let $\boldsymbol{\sigma}_i^R := 1$ denote the benchmark pre-discovery allocation. With a slight abuse of terminology, we refer to this as the *research strategy*, as it shares the core feature of its counterpart in the public progress case: a firm researches as long as it has not independently discovered the new technology.

The following proposition shows that the research strategy not only minimizes the expected duration but also constitutes a Nash equilibrium.

Proposition 2. *Suppose research progress is private, there are no technological spillovers, and the assumptions in Section 2.2 hold.*

[10]For expositional clarity, the main text restricts attention to stationary (time-independent) allocation strategies. In Online Appendix OA.2, we show this restriction is without loss of generality: the constant research profile remains globally optimal among the unrestricted space of all measurable, time-dependent strategies $\boldsymbol{\sigma}_i : \mathbb{R}_+ \to [0,1]$.

(a) ***Efficiency:*** *The expected duration of the race is minimized when both firms adopt the research strategy.*

(b) ***Equilibrium:*** *The research strategy profile constitutes a Nash equilibrium.*

As in the public-progress case, the equilibrium allocation is efficient: it minimizes expected duration of the race given the technological and informational constraints firms face. The gap relative to the first-best outcome, however, has two distinct sources. The first, already present under public progress, is the absence of spillovers—a discovery cannot be transferred to the rival. The second is new: firms cannot condition their strategies on rivals' progress, which prevents them from re-optimizing after a rival's discovery. This second channel motivates the analysis in the following section: even if spillovers were possible through licensing, firms would have an incentive to conceal discoveries precisely because disclosure allows the rival to re-optimize.

3.3 Comparison of Expected Durations

We conclude our benchmark analysis by visually summarizing the expected duration of the innovation race. Figure 3 plots the expected duration under the two benchmark strategies—the research strategy (red dashed curve) and the fall-back strategy (blue solid curve)—against $\frac{1}{\mu}$ (the expected time required for a firm to independently discover the new technology). These are plotted alongside the first-best benchmark, which achieves a shorter expected duration than both strategies due to the presence of technological spillovers.

The intersection of the two strategy curves perfectly illustrates the fall-back condition. When research is sufficiently fast (i.e., $\frac{1}{\mu} \leq \frac{1}{2\lambda_L} - \frac{1}{2\lambda_H}$), the fall-back condition fails, as the research strategy yields a shorter expected duration. In this region, attempting to catch up via the R&D path is more efficient. Because firms seek to minimize expected duration, they adopt the research strategy under both public and private progress. Consequently, the expected durations for both informational settings coincide exactly on the red dashed curve.

The strategic friction emerges when research is relatively slow, satisfying the fall-back condition. Here, the fall-back strategy becomes the efficient response. Under public progress, a firm observes its rival's breakthrough and immediately switches to the fall-back strategy.

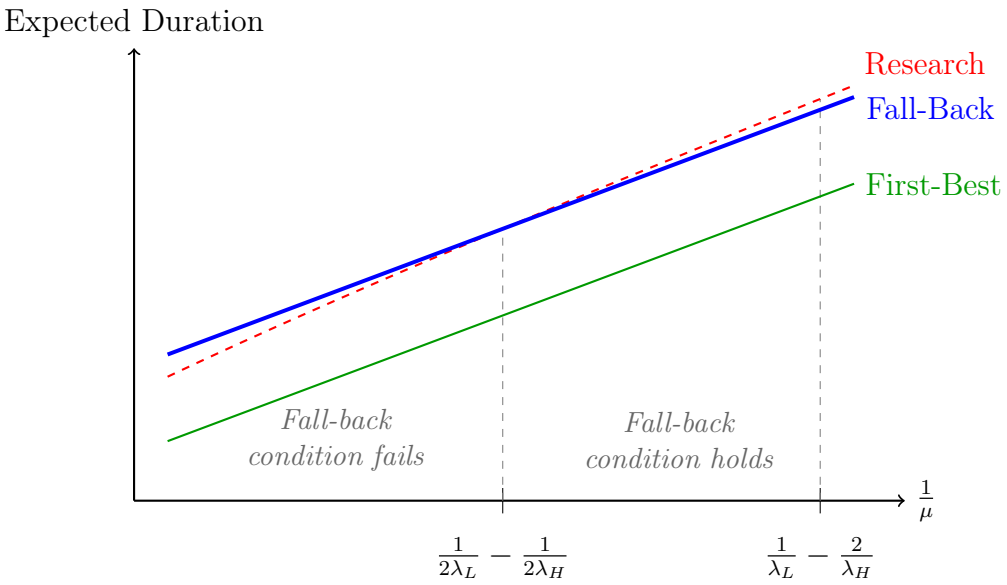


Figure 3: Expected race duration across informational settings.

Thus, the expected duration under public progress tracks the lower envelope of the two curves, pivoting onto the solid blue trajectory.

Under private progress, however, firms remain ignorant of their rival's state and are forced to blindly continue duplicating research efforts. This inability to re-optimize locks the private progress race onto the research strategy (the red dashed curve). Unable to track the lower envelope, the expected duration under private progress diverges upward, widening the efficiency gap.

4 Restoring Efficiency through Patents and Licensing

The preceding analysis shows that—relative to the first-best—there are two sources of inefficiencies: (a) the absence of spillovers, and (b) under private research progress, the inability of firms to observe a rival's breakthrough and re-optimize their R&D strategy accordingly. Patents offer a natural institutional fix that targets both sources. Filing a patent is publicly observable, so patents address (b) as the rival learns of the discovery and may adjust its strategy. Furthermore, patents address (a) by facilitating spillovers through licensing.

In this section, we incorporate patents into the innovation race and explore when the first-best outcome can be achieved via patenting and licensing. A firm that discovers the

new technology may file for a patent; if the patent is granted, the new technology can be transferred to the rival in exchange for a fee through a licensing agreement.

4.1 Patenting Game

We model the patenting decision as a continuous-time stopping game. At any instant $t \geq 0$, an active firm that has independently discovered the new technology, but has not yet filed, must choose whether to file a patent application. Filing is an irreversible and publicly observable action. If a firm files, the game transitions to the patent resolution and licensing subgame. If neither firm files, the race proceeds continuously.

Patent Filing Strategy A firm that discovers the new technology at time τ adopts a patenting strategy characterized by a CDF G_τ over $[\tau, \infty]$.[11] For any finite $t \geq \tau$, $G_\tau(t)$ denotes the probability of having filed by time t. As usual, the CDF G_τ is assumed to be right-continuous. A discontinuity in the CDF represents an atom—a discrete probability mass of filing a patent at a specific instant. A firm's complete patenting strategy profile is therefore given by the collection $\{G_\tau\}_{\tau \geq 0}$.

Subgame after Patent Filing Once Firm i applies for a patent, the outcome of the patent application depends on the status of the rival (Firm j):

- **Uncontested Patent:** If the rival has not yet discovered the technology ($N_t^j = 0$), the patent is granted immediately to Firm i.
- **Prior-Use Defense:** If the rival has independently discovered the technology ($N_t^j = 1$), it can challenge the patent (e.g., by claiming prior use). The challenge succeeds with probability $\beta \in [0, 1]$: (i) if successful, the patent is invalidated or deemed non-exclusive; both firms may use the technology freely; and (ii) if the challenge fails, Firm i is granted an exclusive patent.

[11] The closed interval at infinity is used, following Chatterjee et al. (2026), to formally include full concealment as a strategy. This corresponds to a measure that places the entire probability mass at $t = \infty$, such that the CDF is $G_\tau(t) = 0$ for all finite $t \geq \tau$.

The parameter β captures the strength of intellectual property defenses for independent inventors: $\beta = 0$ corresponds to a pure *first-to-file* system, while $\beta = 1$ represents a perfect *prior-use* defense.

An exclusive patent holder effectively controls the technology. It can deny access or license the technology to the rival via a take-it-or-leave-it offer, thereby capturing the full surplus generated by the rival's use of the superior technology.[12] Figure 4 summarizes the subgame.

Equilibrium Concept The appropriate equilibrium concept for our dynamic game depends on the informational environment. Under public progress, firms perfectly observe each other's research state. Thus, the game is one of complete information, and we focus on Markov Perfect Equilibrium (MPE), where firms' strategies depend only on the current payoff-relevant state rather than the full history of the game. Under private progress, however, a firm cannot observe its rival's discovery unless a patent is explicitly filed. Because firms must make decisions under incomplete information regarding the rival's hidden progress, we employ Perfect Bayesian Equilibrium (PBE). A PBE requires that each firm's strategy is sequentially rational given its beliefs about the rival's state, and that these beliefs are updated according to Bayes' rule whenever possible along the equilibrium path.

4.2 Equilibrium Licensing

A granted patent triggers a licensing subgame: the patent holder makes a take-it-or-leave-it offer of a license fee $l \geq 0$ to the rival, who then accepts or rejects.

If the rival accepts, both firms develop with the new technology, and the rival pays l to the patent holder. The continuation payoffs of the firms are therefore $V_{11} + l$ for the patent holder and $V_{11} - l$ for the rival, where $V_{11} := \frac{\lambda_H \Pi - c}{2\lambda_H}$ represents the expected continuation profit from the race when both firms use the new technology

If the rival rejects the offer, the patent holder develops with the new technology while the rival continues with the old one. In this case, the rival's continuation payoff is $V_{01} := \frac{\lambda_L \Pi - c}{\lambda_H + \lambda_L}$,

[12] An equivalent interpretation is that if Firm j infringes on Firm i's patent, the court orders Firm j to pay damages equal to the full licensing value. This threat forces Firm j to accept a take-it-or-leave-it offer that extracts all surplus generated by the superior technology.

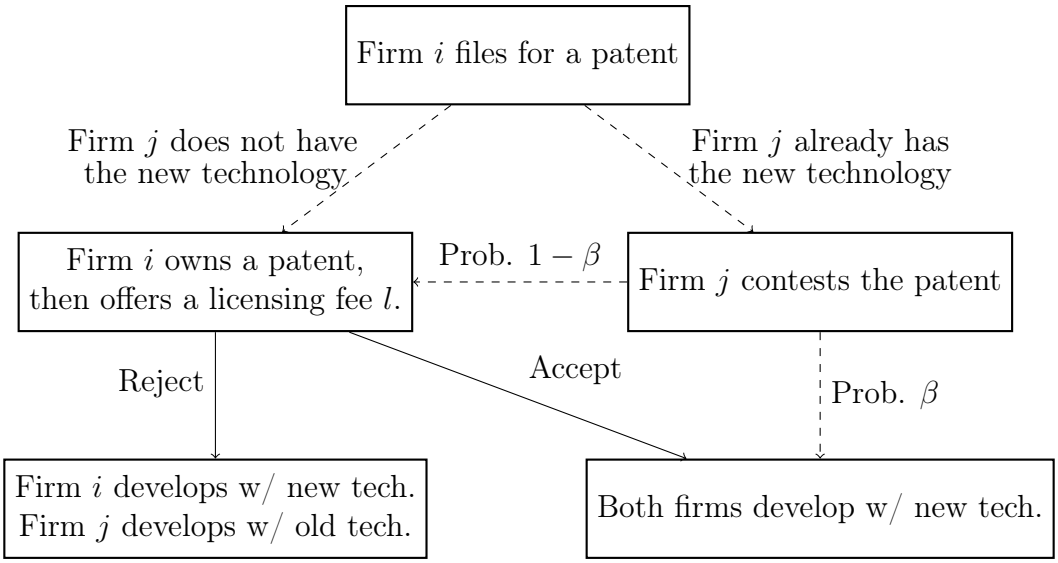


Figure 4: Game Tree after patent filing

which is the expected continuation profit from racing with the old technology against a patent holder possessing the new technology.

In any equilibrium, the license fee must leave the rival indifferent between accepting and rejecting, and the rival accepts the offer. This is formalized in the following lemma.

Lemma 1. *In any subgame perfect Nash equilibrium (SPNE), following the grant of a patent for the new technology, the patent holder offers a license fee* $l^* = V_{11} - V_{01}$ *to the rival, who accepts. The expected continuation payoffs are* $U_P = 2V_{11} - V_{01}$ *for the patent holder and* $U_R = V_{01}$ *for the rival.*[13]

Intuitively, a licensing agreement reduces the expected development time and is therefore efficient. By a Coasian logic, the efficiency gains from the spillover ensure that a mutually agreeable license fee exists. Because the patent holder has all the bargaining power, the equilibrium fee l^* is the maximum it can charge—just enough to leave the licensee indifferent between accepting and rejecting. The patent holder captures, in this way, the entire surplus from the spillover.

In the subsequent subsections, we show that firms' patenting decisions crucially depend on the relative benefit of patenting, defined by $\rho := l^*/V_{11}$.

[13]Because the post-patenting phase constitutes a proper subgame of complete information, this SPNE result applies directly to both the public (MPE) and private (PBE) progress cases.

4.3 Efficient Patenting

The possibility of patenting the new technology can address both sources of inefficiencies identified above. First, filing for a patent is a public event, meaning patents transmit information about research progress to the rival. Second, as established in Lemma 1, a granted patent leads to a licensing agreement, restoring the spillover that is otherwise absent.

A first-best efficient equilibrium is one in which firms exclusively conduct research prior to discovery, patent immediately upon discovering the new technology, and license the technology with probability one, ensuring perfect spillovers.

Definition. We say that patents *restore efficiency* if the patent game admits a first-best efficient equilibrium.

The following result provides the necessary and sufficient conditions under which patents restore efficiency.

Theorem 1. *Suppose that the assumptions in Section 2.2 hold. Patents restore efficiency if and only if either of the following holds:*

1. *Research progress is public information.*
2. *Research progress is private information and* $\rho \geq \rho_M := \frac{\lambda_H}{\lambda_H + \mu(2-\beta)}$, *where* $\rho := \frac{l^*}{V_{11}}$.

When deciding whether to patent a discovery or conceal it, a firm faces a fundamental economic trade-off. Patenting grants the firm exclusivity and the right to extract licensing rents. However, filing also publicly reveals the discovery, which allows the rival to re-optimize its resource allocation.

Under public progress, the rival already observes discoveries instantly. Because filing conveys no new information, it carries no strategic downside. Firms therefore always patent, easily sustaining the efficient equilibrium.

Under private progress, however, filing is the sole mechanism that reveals the breakthrough. To understand when the efficient equilibrium is sustainable in this environment, we compare a firm's payoff from immediate patenting against the alternative of full concealment. In Lemma B.1 in the Appendix, we show that the payoff from any delayed patenting

strategy is simply a convex combination of these two extremes, making this binary comparison sufficient.

Suppose Firm A makes a discovery. Because no prior patent has been filed, it infers that the rival has not yet succeeded. If Firm A patents immediately, it secures the patent with certainty, yielding a continuation payoff of $V_{11} + l^*$.

If instead Firm A fully conceals the discovery, it quietly continues development and wins the race at rate λ_H. Simultaneously, the rival independently discovers the technology at rate μ. If the rival makes the discovery and patents it, Firm A must mount a prior-use defense: it succeeds with probability β (yielding V_{11}) and fails with probability $1-\beta$ (forcing Firm A to pay the license fee, yielding $V_{11} - l^*$). This leaves Firm A with an expected continuation payoff of $V_{11} - (1-\beta)l^*$ in the event of a rival patent. Thus, Firm A's expected continuation payoff from full concealment is:

$$V_C = \frac{\lambda_H \Pi + \mu(V_{11} - (1-\beta)l^*) - c}{\lambda_H + \mu}. \tag{4.1}$$

Immediate patenting dominates full concealment if and only if $V_{11} + l^* \geq V_C$. Substituting the identity $\lambda_H \Pi - c = 2\lambda_H V_{11}$ and rearranging yields the exact condition $\rho \geq \rho_M$.

4.4 Hold-up and the Prior-Use Defense

Theorem 1 condenses the efficiency question into a single inequality, $\rho \geq \rho_M$. To draw out the economics behind it, we isolate two regimes in which the inequality holds automatically, so that efficiency is restored regardless of the stakes: when the fall-back condition fails, and when the prior-use defense is absent.

Proposition 3. *Suppose research progress is private and the assumptions in Section 2.2 hold. If either (i) the fall-back condition (3.2) fails, or (ii) $\beta = 0$, then $\rho > \rho_M$; hence patents restore efficiency by Theorem 1.*

Proof of Proposition 3. Since $\frac{\lambda_L \Pi - c}{\lambda_H \Pi - c} < \frac{\lambda_L}{\lambda_H}$,

$$\rho = 1 - \frac{2\lambda_H}{\lambda_H + \lambda_L} \cdot \frac{\lambda_L \Pi - c}{\lambda_H \Pi - c} \; > \; 1 - \frac{2\lambda_H}{\lambda_H + \lambda_L} \cdot \frac{\lambda_L}{\lambda_H} = \frac{\lambda_H - \lambda_L}{\lambda_H + \lambda_L}.$$

It therefore suffices to show $\frac{\lambda_H-\lambda_L}{\lambda_H+\lambda_L} \geq \rho_M = \frac{\lambda_H}{\lambda_H+\mu(2-\beta)}$ in each case.

(i) The failure of (3.2) reads $\mu(\lambda_H - \lambda_L) \geq 2\lambda_H\lambda_L$, equivalently $\frac{\lambda_H-\lambda_L}{\lambda_H+\lambda_L} \geq \frac{\lambda_H}{\lambda_H+\mu}$. Since $\beta \leq 1$ gives $\rho_M \leq \frac{\lambda_H}{\lambda_H+\mu}$, we obtain $\frac{\lambda_H-\lambda_L}{\lambda_H+\lambda_L} \geq \frac{\lambda_H}{\lambda_H+\mu} \geq \rho_M$.

(ii) With $\beta = 0$, $\rho_M = \frac{\lambda_H}{\lambda_H+2\mu}$. Assumption (2.2) implies $\mu(\lambda_H - \lambda_L) > \lambda_H\lambda_L$, equivalently $\frac{\lambda_H-\lambda_L}{\lambda_H+\lambda_L} > \frac{\lambda_H}{\lambda_H+2\mu} = \rho_M$. □

The two parts of Proposition 3 are not coincidental: inefficient concealment requires two forces at once, and each part switches one off. The first is the *hold-up* that arises when disclosure helps the rival, governed by the fall-back condition; the second is the *safety of concealment*, governed by the prior-use defense.

Hold-up: does disclosure help the rival? The hold-up problem from the introduction tempts a discoverer to conceal only when patenting *raises* the rival's value, $V_P^B > V_C^B$, and whether it does is decided by the fall-back condition. When the condition fails—part (i)—the comparison runs the other way. Suppose Firm A has discovered the technology. If A conceals, Firm B keeps researching, which is its preferred course (Proposition 1). If A instead patents, the technology becomes its exclusive property: researching it no longer pays B, since an independent discovery would now infringe, and B is left to fall back to the slower direct-development path. This is worse for B than the research it would freely pursue under concealment, so patenting *lowers* its value, $V_P^B < V_C^B$. Disclosure weakens the rival rather than strengthening it; no hold-up arises, and A patents immediately—with $\rho > \rho_M$ to spare. When the fall-back condition holds, the comparison reverses: a firm that learns its rival holds the technology would rather abandon research and fall back, so patenting lets B do exactly that and raises its outside option to $V_P^B > V_C^B$. The hold-up problem now bites.

Prior use: is concealment safe? Hold-up alone, however, does not generate inefficiency; concealment must also be safe enough to be worth attempting, and this is the role of the prior-use defense. A firm that conceals risks preemption—its rival may discover the technology independently and patent first—and the defense determines how much it salvages when that happens. When protection is absent—part (ii), the pure first-to-file limit $\beta = 0$—preemption is too costly, and the firm patents immediately whatever the prize, so immediate patenting

constitutes an equilibrium. Only a sufficiently strong defense makes concealment worthwhile: a stronger defense raises the threshold $\rho_M = \lambda_H/[\lambda_H + \mu(2-\beta)]$ while leaving the relative value of licensing ρ unchanged, enlarging the region $\{\rho < \rho_M\}$ in which hold-up bites. Inefficiency therefore requires both forces: the hold-up created by the fall-back condition and a prior-use defense strong enough to shelter concealment.

5 Equilibrium Concealment and Disclosure

The previous section showed that patents restore efficiency unless research progress is private, the fall-back condition holds, and the relative value of licensing is low ($\rho < \rho_M$)—the regime in which the hold-up and prior-use forces jointly bind and firms inefficiently delay disclosure. This section characterizes equilibrium patenting behavior throughout that regime, maintaining the baseline assumptions of Section 2.2 and the fall-back condition. Proofs are relegated to Online Appendix OA.3.

5.1 Equilibrium Concept

To cleanly isolate the strategic timing of patent disclosure, we focus on a certain class of equilibria, namely the **symmetric research equilibrium**. We structure this equilibrium concept across three strategic layers:

(1) prior to any patent application, both firms continuously allocate all resources to research ($\sigma = \mathbf{1}$);

(2) once a patent is filed, firms play the subgame perfect Nash equilibrium (SPNE) of the licensing subgame, yielding the optimal fee l^* (Lemma 1); and

(3) both firms commit to the same patenting strategy $\{G_\tau\}_{\tau \geq 0}$ that maximizes their expected payoff given the first two conditions.

By fixing the pre-patenting resource allocation and analytically resolving the post-patenting bargaining subgame, this equilibrium concept effectively collapses the complex dynamic race. The overarching game simplifies into a single strategic dimension: the optimal timing of

patent disclosure. Therefore, in the analysis that follows, we treat the race strictly as a game of choosing $\{G_\tau\}_{\tau\geq 0}$. We consider a class of patenting strategies characterized by three distinct phases over time.

Definition 2. A patenting strategy $\{G_\tau\}_{\tau\geq 0}$ is called a **gradual patenting strategy** if there exist thresholds T_1, T_2 with $0 \leq T_1 \leq T_2$ such that:[14]

(i) **Silent Phase**: the firm does not patent before T_1: $G_\tau(t) = 0$ for all $0 \leq \tau \leq t < T_1$;

(ii) **Mixing Phase**: the firm engages in partial patenting between T_1 and T_2: $0 < G_\tau(t) < 1$ for all $\tau < t$ and $T_1 \leq t < T_2$, and by the end of this phase, the firm patents with certainty: $G_\tau(T_2) = 1$ for all $\tau < T_2$;

(iii) **Immediate Phase**: the firm patents immediately upon discovery after T_2: $G_t(t) = 1$ for all $t \geq T_2$.

We first show that in any symmetric research equilibrium, the equilibrium patenting strategy takes the form of a gradual patenting strategy.

Proposition 4. *In any symmetric research equilibrium, the equilibrium patenting strategy must be a gradual patenting strategy.*

The intuition for this result is driven by the evolution of firms' beliefs. When a firm discovers the new technology early in the race, it infers that its rival is highly unlikely to have made the discovery yet. Because the threat of the rival securing the patent first or mounting a successful prior-use defense is low, the discovering firm prefers to delay patenting (the silent phase). This concealment prevents the rival from observing the breakthrough, keeping the rival on the R&D path rather than shifting to the direct-development path.

However, as the race drags on without a winner, the conditional probability that the rival has *also* discovered the technology quietly grows. This escalating belief increases the risk that the rival will file a patent first or successfully challenge a late filing. To balance the benefit of concealing against the growing risk of losing the intellectual property, the firm begins to randomize its filing timing (the mixing phase). Eventually, if the race continues

[14] We allow T_1 and T_2 to take ∞ to encompass the boundary cases to be discussed later.

long enough, the belief that the rival possesses the technology becomes so high that the risk of losing the patent strictly outweighs any informational advantage, compelling the firm to patent the moment it makes a discovery (the immediate phase).

5.2 Equilibrium Characterization

Note that even within the class of gradual patenting strategies, the patenting behavior can vary significantly depending on the thresholds T_1 and T_2. We consider three benchmark patenting strategies that represent extreme cases of patenting behavior where at least one phase is degenerate:

1. **Immediate patenting** ($T_1 = T_2 = 0$): The firm patents at the moment of discovery.
2. **Full concealment** ($T_1 = T_2 = \infty$): The firm never patents the discovery.
3. **Delayed mixing** ($T_1 > 0,\ T_2 = \infty$): The firm does not patent before T_1, but continuously and indefinitely patents after T_1.

We show that every symmetric research equilibrium takes one of these three benchmark forms; intermediate configurations—a finite mixing phase, or a silent phase that ends in immediate patenting at a positive time—cannot be sustained in equilibrium.

Theorem 2. *Suppose that firms' research progress is private information and the assumptions in Section 2.2 and the fall-back condition hold. Then, the patenting strategies in the symmetric research equilibria are characterized as follows:*

(a) The unique equilibrium patenting strategy is the ***immediate patenting strategy*** *if* $\rho > \rho_M$.

(b) Every equilibrium takes the form of a ***delayed mixing strategy*** *if*

$$\rho_M > \rho > \rho_D := \frac{\lambda_H(\lambda_H - \mu)}{(\lambda_H + \mu)(\lambda_H - \beta\mu)}. \tag{5.1}$$

(c) The unique equilibrium patenting strategy is the ***full concealment strategy*** *if* $\rho < \rho_D$.

When ρ is high, the potential licensing revenue is highly lucrative relative to the baseline value of the race. In this regime, the firm's priority is rent extraction, driving it to secure the patent immediately upon discovery, even at the cost of accelerating the rival's progress (Immediate Patenting). Conversely, when ρ is low, the baseline payoff of winning the race dominates. The firm's primary objective shifts to protecting its competitive lead, making the public disclosure of a patent too costly. As a result, the firm prefers to hide its breakthrough (Full Concealment) to keep the rival on the slower, less-efficient R&D path. The intermediate regime (Delayed Mixing) emerges as a smooth transition where the firm dynamically balances the option value of future licensing against the escalating risk of rival preemption.

Crucially, Theorem 2 confirms the absolute robustness of the efficiency condition established in Theorem 1. It demonstrates that whenever the relative value of licensing crosses the threshold $\rho > \rho_M$, immediate patenting is not merely *an* equilibrium, but the *unique* symmetric research equilibrium. In this region, the economic incentives to extract licensing rents are sufficiently powerful to completely overcome the strategic fear of expropriation.

5.3 The Impact of Stakes and Prior-Use Defense

The thresholds ρ_M and ρ_D characterize the equilibrium purely in terms of the relative value of licensing. To translate these findings into actionable economic insights, we now reinterpret our results using two fundamental parameters of the innovation environment: the normalized stakes of the race π (recall $\pi > 1$) and the strength of the prior-use defense (β).

Because the relative value of licensing (ρ) is strictly decreasing in π, the threshold inequalities reverse. To find the cutoffs, we solve the equations $\rho = \rho_M$ and $\rho = \rho_D$. First, $\rho = \rho_M$ is equivalent to:

$$\pi_M(\beta) := \frac{\lambda_L}{\lambda_H}\left[\frac{2(\lambda_H + \mu(2-\beta))}{\mu(\beta - \beta_M)} + 1\right] \quad \text{where} \quad \beta_M := \frac{2\lambda_\star}{\lambda_H + \lambda_\star}. \tag{5.2}$$

Note that $\beta_M < 1$ from (3.2), π_M is defined over $(\beta_M, 1]$, and decreasing in β. For all $\beta \leq \beta_M$, we define $\pi_M(\beta) = \infty$.

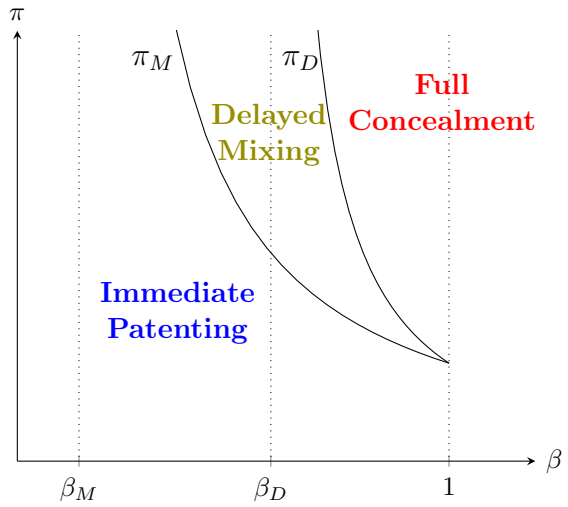


Figure 5: Equilibrium Characterization depending on π and β

Next, $\rho = \rho_D$ is equivalent to:

$$\pi_D(\beta) := \frac{\lambda_L}{\lambda_H}\left[\frac{2(\lambda_H - \mu\beta)}{\mu(\beta - \beta_D)} + 1\right] \quad \text{where} \quad \beta_D := \frac{2\lambda_H(\mu + \lambda_\star)}{(\lambda_H + \mu)(\lambda_H + \lambda_\star)}. \tag{5.3}$$

Note that $\beta_D > \beta_M$ from (3.2), π_D is defined over $(\beta_D, 1]$, and decreasing in β. For all $\beta \leq \beta_D$, we define $\pi_D(\beta) = \infty$.

Corollary 1. *When the fall-back condition holds, the patenting strategies in the symmetric research equilibrium are characterized by the stakes π and the strength of the prior-use defense β as follows:*

1. *if $\pi > \pi_D(\beta)$, the unique equilibrium is the **full concealment** strategy;*
2. *if $\pi \in (\pi_M(\beta), \pi_D(\beta))$, every equilibrium is a **delayed mixing** strategy;*
3. *if $\pi < \pi_M(\beta)$, the unique equilibrium is the **immediate patenting** strategy.*

The corollary and Figure 5 translate ρ into the model's fundamental parameters: the stakes of the race (π) and the strength of the prior-use defense (β).

First, observe that the relative value of licensing is given by:

$$\rho = \frac{l^*}{V_{11}} = 1 - \frac{V_{01}}{V_{11}} = 1 - \frac{\frac{\lambda_L \Pi - c}{\lambda_L + \lambda_H}}{\frac{\lambda_H \Pi - c}{2\lambda_H}} = \left(\frac{\lambda_H - \lambda_L}{\lambda_H + \lambda_L}\right)\left(1 + \frac{2\lambda_L}{\lambda_H \pi - \lambda_L}\right).$$

Because $\lambda_H > \lambda_L$ and active races require $\pi > 1$, the right-hand side is strictly decreasing in π. This proves that as the stakes grow, the relative value of licensing strictly declines.

Economically, when the stakes of winning the race are massive, securing the win vastly overshadows the secondary benefit of extracting licensing rents. Consequently, high-stakes races ($\pi > \pi_D(\beta)$) strongly disincentivize disclosure, driving firms into full concealment to protect their technological lead. Conversely, immediate patenting occurs when the race is of relatively low stakes ($\pi < \pi_M(\beta)$), making the licensing revenue a critical component of the firm's overall profitability. In the intermediate region where $\pi \in (\pi_M(\beta), \pi_D(\beta))$, firms initially conceal their discovery but eventually begin to partially patent as the risk of preemption grows, making delayed mixing the unique equilibrium.

Next, consider the strength of the prior-use defense, β. The thresholds ρ_M and ρ_D crucially depend on β. Note that the inequality $\rho_M \geq \rho_D$ always holds, and the two thresholds perfectly converge when the prior-use defense is absolute ($\beta = 1$). As illustrated in Figure 5, this convergence implies that both π cutoffs collapse to a single point: $\bar{\pi} \equiv \pi_M(1) = \pi_D(1)$.

This convergence reveals that while strengthening the prior-use defense is universally detrimental to public disclosure, its structural impact depends on the stakes of the race:

1. **High Stakes ($\pi > \bar{\pi}$):** In this regime, the strength of the prior-use defense dictates the equilibrium structure. Because ρ_M and ρ_D are strictly increasing in β, the required licensing thresholds become harder to satisfy as prior-use rights strengthen. Visually, when we fix π and increase β, the thresholds π_M and π_D both decrease, pulling the environment out of immediate patenting and into the delayed mixing or full concealment regions. Here, stronger protections merely exacerbate the fear of expropriation, worsening inefficiency by actively promoting secrecy.

2. **Low Stakes ($\pi \leq \bar{\pi}$):** When the stakes are sufficiently low, the relative value of licensing is so substantial that it easily clears the threshold for immediate disclosure, even under a perfect prior-use defense ($\beta = 1$). Because $\pi_M(\beta)$ is strictly decreasing in β, any environment with $\pi \leq \bar{\pi}$ trivially satisfies $\pi < \pi_M(\beta)$ for all $\beta \in [0, 1]$. In this regime, the first-best outcome of immediate patenting is uniquely implemented regardless of the prior-use defense level.

6 Discussion

6.1 Empirical Implications

Our characterization of equilibrium disclosure in Section 5 yields empirical implications for how the normalized stakes of winning the race (π) and the strength of prior-use defense protections (β) affect the social speed of innovation.

With respect to π, the model predicts that its relationship with the aggregate speed of innovation depends crucially on the baseline strength of the prior-use defense (β). Under a strong prior-use defense ($\beta > \beta_D$), higher stakes monotonically weaken the incentive to disclose.[15] When the stakes are low to intermediate ($\pi < \pi_M(\beta)$), firms research and immediately patent their interim breakthroughs; this early disclosure ensures perfect knowledge spillovers, avoids duplicative effort, and maximizes the overall rate of innovation. As the stakes rise into the intermediate region ($\pi_M(\beta) < \pi < \pi_D(\beta)$), firms begin to delay disclosure, partially withholding breakthroughs and slowing diffusion. When the stakes become massively lucrative ($\pi > \pi_D(\beta)$), the baseline payoff of winning overshadows the licensing revenue, and firms switch to full concealment; this breakdown in disclosure eliminates spillovers, forces redundant research, and slows cumulative innovation.

Conversely, under a weak prior-use defense ($\beta \leq \beta_M$), this non-monotonicity disappears. Because the threat of a rival successfully mounting an independent prior-use defense is sufficiently low, the efficient outcome becomes remarkably robust. Regardless of how high the stakes (π) become, immediate patenting remains the unique equilibrium. In this regime, the first-best outcome is always implemented, meaning high stakes consistently sustain perfect knowledge spillovers without triggering a retreat into secrecy.

Regarding β, the empirical impact is strongly negative but ultimately bounded by the stakes of the race. In high-stakes environments ($\pi > \overline{\pi}$), stronger prior-use rights increase the likelihood that a rival can successfully contest a patent, meaning a higher β strictly discourages early patenting. However, in low-stakes competitions ($\pi < \overline{\pi}$), this dampening

[15] This comparative static also implies that *lowering* the stakes π accelerates innovation. The effect is bounded below: our analysis maintains the assumptions of Section 2.2, equivalently $\pi > 1$. Were π pushed below this threshold, Assumption 1 (viability of direct development) would fail and firms might not initiate the race at all. Extended beyond the region we model, the relationship between π and the speed of innovation is therefore non-monotonic.

effect is neutralized. When the baseline prize is sufficiently small, the relative value of licensing revenue is so dominant that the first-best outcome of immediate disclosure is uniquely implemented, completely regardless of the strength of the prior-use defense (β).

These structural variations provide a rich, testable matrix for empirical investigation. Because the practical enforceability of prior-use defenses varies significantly across industries—easily documented in manufacturing via physical prototypes, but highly ambiguous in software—researchers can compare disclosure timing and aggregate innovation speeds across sectors with varying de facto intellectual property protections and observable market stakes.

6.2 Policy Implications

While significant expected rewards are essential to incentivize initial R&D participation, our analysis reveals that excessively high stakes (π) in the product market can inadvertently stifle knowledge diffusion. When the payoff gap between winning the race and competing as a laggard is massive, firms are driven to conceal their breakthroughs to preserve their competitive lead, resulting in redundant research and a decelerated aggregate pace of innovation. To mitigate this inefficiency, policymakers might consider redistributive interventions in the final product market. For instance, imposing a tax on the monopolistic winner and redistributing a portion of this surplus to the lagging firm directly compresses the profit differential. This intervention effectively lowers the normalized stakes of the race. Provided the baseline rewards remain sufficient to satisfy the firms' R&D participation constraints, shifting the industry from a high-stakes concealment regime into a low-stakes immediate patenting regime can restore knowledge spillovers and ultimately accelerate overall technological progress.[16]

Our results also speak to the strength of the prior-use defense (β). Within the model, a strong prior-use defense reduces the expected security of a patent, exacerbating the fear of expropriation and pushing firms toward concealment; conversely, limiting independent prior-use claims—as a first-to-file regime does—raises the security of an early patent and

[16]This lower bound matters: the participation constraint is Assumption 1, equivalently $\pi > 1$, which ensures that developing with the old technology is worthwhile and firms enter the race. A tax that drove π below this threshold would violate the assumption and could deter firms from initiating the race altogether. Extended beyond this participation region, the relationship between the stakes and the speed of innovation is therefore non-monotonic.

encourages immediate disclosure. Read narrowly, this would seem to recommend a weak prior-use defense. Such a reading, however, ignores the functions that prior-use protection serves outside our model: it shields independent inventors who rely on trade secrecy from being held up by a later patentee, deters litigation by non-practicing patent holders, and spares firms the cost of purely defensive patenting. A policymaker who values these objectives cannot treat β as a free dial to be lowered whenever disclosure lags.

This asymmetry between the two levers is the practical upshot of our analysis. Because the strength of the prior-use defense is anchored by legal and normative commitments that have little to do with the diffusion of any single breakthrough, β is best regarded as fixed from the standpoint of innovation policy. The stakes of the race, π, are far more readily shaped—through taxation, procurement design, or other redistribution in the product market. Taking the prior-use regime as given, then, a policymaker can restore efficient disclosure by lowering the stakes of the race, provided the prize remains large enough to keep firms in it. Efficiency need not come at the expense of institutions society values for other reasons; it requires calibrating the one margin—competition for the prize—that policy can actually control.

6.3 Future Research

Our framework isolates the strategic timing of disclosure in a dual-path innovation race, but it naturally opens several avenues for future research. First, while we assume resource allocation is constant over time, allowing firms to endogenously dynamically adjust their R&D effort could yield richer preemption dynamics. Second, extending the model beyond two exogenous paths to encompass multi-stage or cumulative breakthroughs would allow for the study of partial technological disclosures. Finally, shifting from a positive analysis to a normative mechanism design perspective—such as exploring optimal contest architectures that dynamically adjust the winner-takes-all prize structure to mitigate strategic concealment—remains a promising direction for future work.

Appendix

A Proofs for Section 3

A.1 Public Research Progress

Proof of Proposition 1. Since $\mathbf{s}_A(1,\cdot) = 0$ is known, Firm A's strategy, $\mathbf{s}_A$, can be characterized by (a_0, a_1) where $a_{N_B} = \mathbf{s}_A(0, N_B)$. Likewise, Firm B's strategy can be characterized by (b_0, b_1) where $b_{N_A} = \mathbf{s}_B(N_A, 0)$.

(a) Let D_{ij} denote the expected duration when the state is $(N^A, N^B) = (i, j)$. First, note that $D_{11} = \frac{1}{2\lambda_H}$, as both firms develop the product with the new technology at rate λ_H.

Next, consider the state $(0, 1)$—Firm B has discovered the new technology, but Firm A has not. Then, the expected duration given Firm A's allocation a_1 is:

$$D_{01}(a_1) := \frac{1}{\lambda_H + \mu a_1 + \lambda_L(1 - a_1)} \left(1 + \mu a_1 \cdot D_{11}\right).$$

By taking the derivative with respect to a_1, we have:

$$D'_{01}(a_1) = \frac{\mu\lambda_L}{(\lambda_H + \mu a_1 + \lambda_L(1 - a_1))^2} \left\{ \frac{1}{\mu} - \left(\frac{1}{2\lambda_L} - \frac{1}{2\lambda_H} \right) \right\}.$$

This implies that $D'_{01}(a_1) > 0$ if and only if (3.1) holds. Thus, $a_1 = 0$ minimizes D_{01} if the fall-back condition holds; and $a_1 = 1$ minimizes D_{01} if the fall-back condition fails to hold. This result symmetrically holds for $D_{10}(b_1)$. We can then define the minimized expected duration when exactly one firm has discovered the new technology:

$$D_1 := \min\left\{D_{01}(0), D_{01}(1)\right\} = \min\left\{ \frac{1}{\lambda_H + \lambda_L}, \ \frac{2\lambda_H + \mu}{2\lambda_H(\lambda_H + \mu)} \right\}.$$

Last, consider the state $(0, 0)$—neither firm has discovered the new technology yet. Let $S = a_0 + b_0$ denote the aggregate research share before any breakthrough has been made. Then,

$$D_{00}(S) := \frac{1 + \mu S \cdot D_1}{\lambda_L(2 - S) + \mu S}.$$

By taking the derivative with respect to S, we have:

$$D'_{00}(S) = \frac{2\lambda_L\mu}{(\lambda_L(2-S)+\mu S)^2}\left\{D_1 - \frac{1}{2}\left(\frac{1}{\lambda_L}-\frac{1}{\mu}\right)\right\}.$$

Note that by definition,

$$D_1 - \frac{1}{2}\left(\frac{1}{\lambda_L}-\frac{1}{\mu}\right) \leq \frac{1}{\lambda_H+\lambda_L} - \frac{1}{2}\left(\frac{1}{\lambda_L}-\frac{1}{\mu}\right) < \frac{1}{\lambda_H} - \frac{1}{2}\left(\frac{1}{\lambda_L}-\frac{1}{\mu}\right) < 0,$$

where the final strict inequality holds from Assumption (2.2). Therefore, D_{00} is decreasing in S, thus, it is minimized at $S = 2$, i.e., $a_0 = b_0 = 1$.

(b) At state $(1,1)$, both firms' expected payoffs are $V_{11} := \frac{\lambda_H\Pi - c}{2\lambda_H}$.

At state $(0,1)$, only Firm A makes a strategic choice, and its problem is to maximize:

$$V_{01}(a_1) := \frac{\mu a_1 V_{11} + \lambda_L(1-a_1)\Pi - c}{\mu a_1 + \lambda_L(1-a_1) + \lambda_H}.$$

By taking the derivative with respect to a_1, we have:

$$V'_{01}(a_1) = -\frac{\lambda_L\mu(\lambda_H\Pi + c)}{(\lambda_H + \mu a_1 + \lambda_L(1-a_1))^2}\left\{\frac{1}{\mu} - \left(\frac{1}{2\lambda_L} - \frac{1}{2\lambda_H}\right)\right\}.$$

Therefore, $a_1 = 0$ maximizes V_{01} if (3.1) holds; whereas $a_1 = 1$ maximizes V_{01} if (3.1) fails to hold. Given this result, Firm B's continuation payoff at state $(0,1)$ can also be computed. Then, Firm A's continuation payoff is V_1^-, while Firm B's continuation payoff is V_1^+, which are defined as:

$$V_1^+ := \begin{cases} \dfrac{\lambda_H\Pi - c}{\lambda_H+\lambda_L} & \text{if (3.1) holds;} \\ \dfrac{\lambda_H\Pi + \mu V_{11} - c}{\lambda_H+\mu} & \text{if (3.1) fails,} \end{cases} \quad \text{and} \quad V_1^- := \begin{cases} \dfrac{\lambda_L\Pi - c}{\lambda_H+\lambda_L} & \text{if (3.1) holds;} \\ \dfrac{\mu V_{11} - c}{\lambda_H+\mu} & \text{if (3.1) fails.} \end{cases}$$

This result symmetrically holds for the state $(1,0)$.

Now, consider the state $(0,0)$. Firm A's expected payoff is

$$V_A(a_0,b_0) := \frac{\lambda_L(1-a_0)\Pi + \mu a_0 V_1^+ + \mu b_0 V_1^- - c}{\mu(a_0+b_0) + \lambda_L(2-a_0-b_0)}.$$

Let $\Delta_0 := V_A(1,0) - V_A(0,0)$ and $\Delta_1 := V_A(1,1) - V_A(0,1)$. With some algebra, we can derive that

$$\frac{\partial V_A}{\partial a_0} = \frac{2(\lambda_L + \mu)}{\{\mu(a_0 + b_0) + \lambda_L(2 - a_0 - b_0)\}^2} \cdot \{\lambda_L(1 - b_0) \cdot \Delta_0 + \mu b_0 \cdot \Delta_1\}.$$

Therefore, if $\Delta_0, \Delta_1 \geq 0$, then $a_0 = 1$ is a dominant strategy. Since Firm B's problem is symmetric to Firm A's, it is sufficient to show that $\Delta_0, \Delta_1 \geq 0$.

When (3.1) holds, using the definition of V_1^+, V_1^- and V_A, we can derive that:

$$\Delta_0 = \frac{(\lambda_L \Pi + c)(\lambda_\star - \lambda_L)}{2(\lambda_L + \mu)(\lambda_L + \lambda_H)}, \qquad \Delta_1 = \frac{\lambda_L(\mu\Pi + c)(\lambda_\star - \lambda_L)}{2\mu(\lambda_L + \mu)(\lambda_L + \lambda_H)}.$$

From Footnote 9, we have $\lambda_\star > \mu > \lambda_L$, thus, $\Delta_0, \Delta_1 > 0$.

Next, when (3.1) fails, or equivalently $\lambda_\star \geq \lambda_H$, using the definition of V_1^+, V_1^- and V_A, we can derive that:

$$\Delta_0 = \frac{\lambda_H \cdot \lambda_\star \cdot (\lambda_L \Pi + c) + \mu \cdot (\lambda_\star - \lambda_H) \cdot c}{2\lambda_H(\lambda_H + \mu)(\lambda_L + \mu)},$$

$$\Delta_1 = \frac{\lambda_L \cdot \{\lambda_H \cdot \lambda_\star \cdot (\mu\Pi + c) + \mu \cdot (\lambda_\star - \lambda_H) \cdot c\}}{2\mu\lambda_H(\lambda_H + \mu)(\lambda_L + \mu)}.$$

From $\lambda_\star \geq \lambda_H$, we have $\Delta_0, \Delta_1 > 0$. Therefore, at state $(0,0)$, both firms choose $a_0 = b_0 = 1$. □

A.2 Private Research Progress

Proof of Proposition 2. For simplicity, let $a = \boldsymbol{\sigma}_A$ and $b = \boldsymbol{\sigma}_B$.

(a) Let $\mathcal{D}_1(x)$ denote the expected remaining duration when one firm has discovered the new technology and the rival allocates a share x of resources to research:

$$\mathcal{D}_1(x) := \frac{1}{\lambda_H + \mu x + \lambda_L(1 - x)} \left\{1 + \frac{\mu x}{2\lambda_H}\right\}.$$

Given a strategy profile (a, b), the expected duration from the beginning of the race is:

$$\mathcal{D}_0(a,b) := \frac{1}{\mu(a+b) + \lambda_L(2-a-b)} \left\{1 + \mu a \mathcal{D}_1(b) + \mu b \mathcal{D}_1(a)\right\}.$$

Let $S := a + b$ and $M := ab$. Note that $\max\{0, S-1\} \leq M \leq (S/2)^2$. Using these variables, $\mathcal{D}_0(a,b)$ can be rewritten as follows:

$$\tilde{\mathcal{D}}_0(S,M) := \frac{1}{2\lambda_L + (\mu - \lambda_L)S} \left[1 + \frac{\mu}{2\lambda_H} \frac{NUM_0(S,M)}{DEN_0(S,M)}\right]$$

where

$$\begin{aligned} NUM_0(S,M) &:= (2\lambda_L(\lambda_H - \lambda_\star) + \mu(\mu - \lambda_L)S)M + 2\lambda_H(\lambda_H + \lambda_L + (\mu - \lambda_L)S)S, \\ DEN_0(S,M) &:= (\lambda_H + \lambda_L)^2 + (\mu - \lambda_L)(\lambda_H + \lambda_L)S + (\mu - \lambda_L)^2 M. \end{aligned}$$

Taking the partial derivative with respect to M, we have:

$$\frac{\partial \tilde{\mathcal{D}}_0}{\partial M} = \frac{\lambda_L \mu \{(\mu - \lambda_L)S + 2(\lambda_H + \lambda_L)\}(\lambda_H + \lambda_L + (\mu - \lambda_L)S)(\lambda_H - \lambda_\star)}{2\lambda_H(2\lambda_L + (\mu - \lambda_L)S)\{(\lambda_H + \lambda_L)^2 + (\mu - \lambda_L)(\lambda_H + \lambda_L)S + (\mu - \lambda_L)^2 M\}^2}.$$

The sign of this derivative depends entirely on the term $(\lambda_H - \lambda_\star)$. This shows that (i) if $\lambda_H > \lambda_\star$, the expected duration is increasing in M, meaning it is minimized at the lower bound $M = \max\{0, S-1\}$; and (ii) if $\lambda_H < \lambda_\star$, the expected duration is decreasing in M, meaning it is minimized at the upper bound $M = (S/2)^2$.

First, consider the case where $\lambda_H > \lambda_\star$. When $S \leq 1$, $\tilde{\mathcal{D}}_0$ is minimized at $M = 0$. Evaluating the derivative with respect to S at this boundary yields:

$$\frac{\partial \tilde{\mathcal{D}}_0}{\partial S}(0,S) = -\frac{\lambda_L(\lambda_\star - \lambda_L)}{(\lambda_H + \lambda_L)(2\lambda_L + (\mu - \lambda_L)S)^2} < 0,$$

since $\lambda_\star > \mu > \lambda_L$ from Footnote 9.

When $S > 1$, $\tilde{\mathcal{D}}_0$ is minimized at $M = S - 1$. Note that $\tilde{\mathcal{D}}_0(S-1, S) = \mathcal{D}_0(a, 1)$ where

$a = S - 1$. Differentiating with respect to a gives $\frac{\partial \mathcal{D}_0}{\partial a}(a,1) = -\lambda_L \cdot \frac{NUM_1(a)}{DEN_1(a)}$ where

$$\begin{aligned} NUM_1(a) =& \{\lambda_H(\lambda_\star - \mu) + (\lambda_H + \mu)\lambda_\star\}(\lambda_L + \mu + (\mu - \lambda_L)a)^2 \\ &+ 2\lambda_H^2(2\lambda_L + \lambda_H + \mu + 2(\mu - \lambda_L)a)(\lambda_\star - \mu), \\ DEN_1(a) =& 2\lambda_H(\lambda_H + \mu)(\lambda_L + \mu + (\mu - \lambda_L)a)^2(\lambda_L + \lambda_H + (\mu - \lambda_L)a)^2. \end{aligned}$$

Since $\lambda_\star > \mu > \lambda_L$, all terms in $NUM_1(a), DEN_1(a)$ are strictly positive, implying $\frac{\partial \mathcal{D}_0}{\partial a}(a,1) < 0$. Thus, $\mathcal{D}_0(a,1)$ is minimized at $a = 1$. Therefore, overall, $a = b = 1$ minimizes $\mathcal{D}_0$.

Next, consider the case where $\lambda_H < \lambda_\star$. Here, $\tilde{\mathcal{D}}_0(M,S)$ is minimized at $M = (S/2)^2$, which implies the symmetric allocation $a = b = S/2$. Evaluating the derivative along this symmetric path yields $\frac{\partial \mathcal{D}_0}{\partial a}(a,a) = -\lambda_L \frac{NUM_2(a)}{DEN_2(a)}$, where

$$\begin{aligned} NUM_2(a) =& \{(\mu - \lambda_L)(2(\lambda_H + \mu a)(\lambda_\star - \lambda_H) + \lambda_H\lambda_L a) + \lambda_H\lambda_L(2\lambda_\star + \lambda_L)\}(\lambda_L + (\mu - \lambda_L)a) \\ &+ \lambda_H\lambda_L\lambda_\star^2, \\ DEN_2(a) =& 2(\mu - \lambda_L)\lambda_H\{\lambda_L + (\mu - \lambda_L)a\}^2\{\lambda_H + \lambda_L + (\mu - \lambda_L)a\}^2 \end{aligned}$$

Since $\lambda_\star > \lambda_H$, we have $NUM_2(a), DEN_2(a) > 0$, implying that $\mathcal{D}_0(a,a)$ is minimized at $a = 1$. Therefore, $a = b = 1$ also minimizes $\mathcal{D}_0$ in this case as well.

(b) Recall that $V_{11} = \frac{\lambda_H \Pi - c}{2\lambda_H}$ is a firm's expected payoff when both firms have discovered the new technology. When Firm A has discovered the new technology and Firm B uses the strategy b, Firm A's expected payoff is:

$$\mathcal{V}^+(b) := \frac{\lambda_H \Pi + \mu b V_{11} - c}{\lambda_H + \mu b + \lambda_L(1-b)}.$$

Conversely, when Firm B has discovered the new technology and Firm A uses strategy a, Firm A's expected payoff is:

$$\mathcal{V}^-(a) := \frac{\mu a V_{11} + \lambda_L(1-a)\Pi - c}{\lambda_H + \mu a + \lambda_L(1-a)}.$$

Prior to any discovery, Firm A's total expected payoff when using allocation a while Firm

B uses b can be written as:

$$\mathcal{V}_0(a,b) := \frac{\mu a \mathcal{V}^+(b) + \mu b \mathcal{V}^-(a) + \lambda_L(1-a)\Pi - c}{\mu(a+b) + \lambda_L(2-a-b)}.$$

To show that $(a,b) = (1,1)$ constitutes a Nash equilibrium, it suffices to show that $\frac{\partial}{\partial a}\mathcal{V}_0(a,1) > 0$ for all $a \in [0,1]$. Through algebraic expansion, we can derive the following equality:

$$\begin{aligned}
&\{\mu + \lambda_L + (\mu - \lambda_L)a\} \cdot \{\lambda_H + \lambda_L + (\mu - \lambda_L)a\} \cdot \frac{\partial \mathcal{V}_0}{\partial a}(a,1) \\
=&\mu(\lambda_\star - \mu)\left[\frac{2\lambda_H\lambda_L}{(\lambda_H+\mu)(\mu+\lambda_L+(\mu-\lambda_L)a)} + \frac{\lambda_L}{\lambda_H+\lambda_L+(\mu-\lambda_L)a}\right] \cdot \frac{\lambda_H\Pi + c}{2\lambda_H} \\
&+ \frac{\mu(\mu-\lambda_L)(\lambda_H^2 + \mu\lambda_L + \mu(\lambda_H - \lambda_L)a)}{(\lambda_H+\mu)(\lambda_H+\lambda_L+(\mu-\lambda_L)a)} \cdot \frac{\lambda_H\Pi + c}{2\lambda_H} \\
&+ \frac{\lambda_L(\lambda_\star - \mu)(\lambda_H + \lambda_L + (\mu-\lambda_L)a)}{(\lambda_H+\mu)(\mu+\lambda_L+(\mu-\lambda_L)a)} \cdot c.
\end{aligned}$$

Given $\lambda_\star > \mu > \lambda_L$ and $\lambda_H > \lambda_L$, every additive term on the right-hand side is strictly positive. Consequently, $\mathcal{V}_0(a,1)$ is increasing in a. This implies that Firm A's best response to Firm B playing $b = 1$ is to set $a = 1$. By symmetry, $(a,b) = (1,1)$ constitutes a Nash equilibrium. □

B Proofs for Section 4

B.1 Licensing Fee: Proof of Lemma 1

Proof of Lemma 1. When the offer is rejected, Firm j's expected payoff is $\frac{\lambda_L\Pi - c}{\lambda_H + \lambda_L}$. Note that V_{11} is the expected payoff when both firms race with the new technology. Thus, when the license offer with the fee l is accepted, Firm j's expected payoff is $V_{11} - l$. Then, Firm i's optimal offer is $l^* = V_{11} - (\lambda_L\Pi - c)/(\lambda_H + \lambda_L)$ with simple algebra. Then, once the offer is accepted, Firm i's expected payoff is $V_{11} + l^*$ and Firm j's expected payoff is $V_{11} - l^*$. □

B.2 Proof of Theorem 1

In this section, we provide the proof of Theorem 1. In the first subsection, we show that efficiency is restored in the public progress case. In the subsequent subsection, we focus on the private progress case and characterize the exact conditions under which the efficient equilibrium can be supported.

B.2.1 Public Progress

In this subsection, we show that an efficient profile—where both firms research exclusively and patent immediately upon discovery ($G_\tau(\tau) = 1$ for all $\tau \geq 0$)—constitute a subgame perfect equilibrium.

Subgame after Discovery We consider a subgame where Firm A discovers the new technology at time τ. We show that the immediate patenting strategy ($G_\tau(\tau) = 1$) is the unique equilibrium strategy.

First, consider the off-path case where Firm B discovered the new technology earlier ($\tau' < \tau$) but has not yet patented (on the equilibrium path, Firm B patents immediately). Since both firms have discovered the new technology, they both develop the product at rate λ_H, yielding an expected payoff of V_{11} if neither patents. If Firm A patents before Firm B, it secures an exclusive patent with probability $(1-\beta)$ (since Firm B's prior-use challenge fails with probability $1-\beta$), guaranteeing Firm A an additional expected payoff of $(1-\beta)l^*$. Conversely, delaying patenting allows Firm B to patent first, which would not only forfeit Firm A's licensing revenue but also force Firm A to pay an expected fee of $(1-\beta)l^*$ to Firm B. Thus, Firm A's strict best response is to patent immediately.

Next, consider the case where Firm B has not yet discovered the new technology. Suppose Firm A uses a patenting strategy G_τ such that $G_\tau(\tau) < 1$, while Firm B uses some resource allocation strategy $\mathbf{s}$ and a subsequent patenting strategy. Let U^A and U^B denote the expected equilibrium payoffs for Firm A and B in this subgame. Note that $U^A + U^B < 2V_{11}$, because the joint expected payoff is strictly maximized when both firms use the new technology from time τ onward, which is not achieved if $G_\tau(\tau) < 1$.

If Firm B allocates no resources to research, its expected payoff is:

$$G_\tau(\tau)\cdot(V_{11}-l^*)+\int_\tau^\infty(\lambda_L\Pi-c)\cdot e^{-(\lambda_H+\lambda_L)(t-\tau)}(1-G_\tau(t))dt \\ +\int_\tau^\infty(V_{11}-l^*)\cdot e^{-(\lambda_H+\lambda_L)(t-\tau)}dG_\tau(t).$$

Using $V_{11}-l^*=V_{01}=\frac{\lambda_L\Pi-c}{\lambda_H+\lambda_L}$ and integration by parts, the sum of the last two terms simplifies exactly to $(1-G_\tau(\tau))\cdot(V_{11}-l^*)$. This shows that Firm B's guaranteed expected payoff is $V_{11}-l^*$, thus, $U^B\geq V_{11}-l^*$. Combining this lower bound with $U^A+U^B<2V_{11}$, we obtain $U^A<V_{11}+l^*$. However, $V_{11}+l^*$ is exactly the expected payoff Firm A secures by using the immediate patenting strategy ($G_\tau(\tau)=1$). Therefore, any strategy other than immediate patenting strategy yields a strictly lower payoff and cannot be part of equilibrium. By symmetry, the same holds for Firm B.

Resource allocation We now consider the state where neither firm has discovered the new technology yet. We must show that when Firm B employs the constant research strategy ($b=1$), Firm A's best response is also the constant research strategy ($a=1$).

Suppose Firm B uses the constant strategy $b=1$, and Firm A chooses a potentially time-dependent, measurable allocation policy $\sigma:\mathbb{R}_+\to[0,1]$. Anticipating that both firms will immediately patent upon any discovery in the continuation subgame, let $U(x)$ denote Firm A's expected payoff if Firm A were to simply use a constant policy $x\in[0,1]$:

$$U(x):=\frac{\mu(V_{11}-l^*)+\mu x(V_{11}+l^*)+\lambda_L(1-x)\Pi-c}{\mu+\mu x+\lambda_L(1-x)}.$$

Taking the derivative with respect to x yields:

$$U'(x)=\frac{\lambda_L\{\lambda_H\mu(\lambda_\star-\lambda_L)\Pi+(\lambda_H\mu+(\lambda_H+\lambda_L+\mu)\lambda_\star)c\}}{\lambda_H(\lambda_H+\lambda_L)(\mu x+\lambda_L(1-x)+\mu)^2},$$

which is strictly positive because $\lambda_\star>\mu>\lambda_L$. Therefore, $U(x)$ is strictly increasing in x, implying that $x=1$ is the unique point-wise maximizer of the stationary payoff.

Because the game environment is stationary, this point-wise maximization translates directly to dynamic optimality. As we rigorously establish in Online Appendix OA.2.3 using

optimal control theory (specifically, Arrow's sufficiency condition), because $x = 1$ maximizes the stationary objective $U(x)$, the constant policy $\sigma^*(t) = 1$ for all $t \geq 0$ is Firm A's global best response among all measurable time-dependent strategies. By symmetry, the profile where both firms play the constant research strategy constitutes a Nash equilibrium. □

B.2.2 Private Progress

In this subsection, we assume that research progress is private and characterize conditions under which the efficient profile can be supported as an equilibrium. It suffices to show that Firm A's best response to Firm B's researching and immediate patenting is also researching and immediate patenting.

Equilibrium Patenting Strategy Suppose Firm A discovered the new technology at time τ and Firm B has not patented yet. Since Firm B uses the immediate patenting strategy, Firm B has not discovered the new technology yet. We begin with a lemma specifying Firm A's payoff when it employs the patenting strategy G_τ.

Lemma B.1. *Suppose that Firm A discovered the new technology and Firm B has not, and Firm B uses the research and immediate patenting strategy. Firm A's expected payoff is:*

$$V_C + G_\tau(\tau) \cdot \{V_{11} + l^* - V_C\} + \int_\tau^\infty (V_{11} + l^* - V_C) e^{-(\lambda_H + \mu)(t-\tau)} dG_\tau(t). \tag{B.1}$$

Proof of Lemma B.1. First, observe that Firm A's expected payoff of employing G_τ is:

$$\begin{aligned} &G_\tau(\tau) \cdot (V_{11} + l^*) + \int_\tau^\infty (V_{11} + l^*) \cdot e^{-(\lambda_H + \mu)(t-\tau)} \, dG_\tau(t) \\ &+ \int_\tau^\infty (\lambda_H \Pi - c) \cdot e^{-(\lambda_H + \mu)(t-\tau)} \cdot (1 - G_\tau(t)) \, dt \\ &+ \int_\tau^\infty (V_{11} - (1-\beta) \cdot l^*) \cdot e^{-(\lambda_H + \mu)(t-\tau)} \cdot (1 - G_\tau(t)) \mu \, dt. \end{aligned}$$

The first two terms correspond to the payoffs when Firm A patents (either immediately at τ or at some $t > \tau$). The third term corresponds to the case where Firm A develops the product first, net of cost. The last term is the expected payoff when Firm B independently discovers and patents before Firm A files.

Observe that by the definition of $V_C = \frac{\lambda_H \Pi - c + \mu(V_{11} - (1-\beta)l^*)}{\lambda_H + \mu}$, the sum of the last two terms can be factored as $\int_\tau^\infty (1 - G_\tau(t)) \cdot V_C \cdot (\lambda_H + \mu) \cdot e^{-(\lambda_H+\mu)(t-\tau)} dt$. By expressing $1 - G_\tau(t) = \int_t^\infty dG_\tau(s)$ and reversing the order of integration (or equivalently, applying integration by parts), we can directly evaluate this integral as:

$$V_C(1 - G_\tau(\tau)) - \int_\tau^\infty V_C \cdot e^{-(\lambda_H+\mu)(t-\tau)} \, dG_\tau(t).$$

Substituting this evaluated expression back into the original payoff equation yields (B.1). □

This lemma implies that the immediate patenting strategy ($G_\tau(\tau) = 1$) is the maximizer of Firm A's expected payoff if and only if $V_{11} + l^* \geq V_C$, which is equivalent to $\rho \geq \rho_M$.

Resource allocation We now consider the pre-discovery resource allocation stage under private progress. We must show that if Firm B plays the efficient strategy profile—allocating all resources to research and patenting immediately upon discovery—Firm A's optimal response is to also allocate all resources to research.

Because Firm B files a patent without delay, its private discovery is instantly made public. Consequently, as long as no patent has been filed, Firm A can infer with certainty that Firm B has not yet discovered the new technology. This implies that, from Firm A's perspective, the evolution of the state and the information sets prior to a breakthrough are identical to those under public progress.

Furthermore, because both firms are playing the immediate patenting strategy in the continuation subgame, the expected payoffs upon discovery are perfectly preserved: Firm A receives $V_{11} + l^*$ upon its own discovery and $V_{11} - l^*$ upon Firm B's discovery. Because both the state transition probabilities and the continuation payoffs exactly match the public progress environment, Firm A's expected payoff functional for any pre-discovery allocation strategy is mathematically identical to the objective $U(x)$ defined in the public progress case. Therefore, by the identical optimal control logic established in Appendix OA.2.3, Firm A's unique global best response is the constant research strategy ($x = 1$). By symmetry, this sustains the efficient resource allocation.

Online Appendix for "*Strategic Concealment in Innovation Races*"

OA.1 Single-Firm Benchmark

Proposition OA.1.1. *Suppose that there is only one firm in the race, and it faces an exogenous termination rate λ. The optimal allocation policy is*

(a) to follow the R&D path if $\lambda < \lambda_\star$;

(b) to develop with the old technology if $\lambda > \lambda_\star$.

Proof of Proposition OA.1.1. Let $V_\lambda(N^i)$ denote the value function of Firm i in state N^i. If $N^i = 1$, the firm develops at rate λ_H until termination λ. The expected duration is $1/(\lambda_H + \lambda)$, and the success probability is $\lambda_H/(\lambda_H + \lambda)$. Thus, the continuation value is

$$V_\lambda(1) := \frac{\lambda_H}{\lambda_H + \lambda} \cdot \Pi - \frac{1}{\lambda_H + \lambda} \cdot c = \frac{\lambda_H \Pi - c}{\lambda_H + \lambda}. \tag{OA.1.1}$$

If $N^i = 0$, the firm chooses allocation x to maximize the expected payoff. The corresponding Hamilton-Jacobi-Bellman (HJB) equation is:

$$0 = \max_{x \in [0,1]} \Big\{ \lambda_L (1-x)(\Pi - V_\lambda(0)) + \mu x (V_\lambda(1) - V_\lambda(0)) + \lambda(0 - V_\lambda(0)) - c \Big\}. \tag{OA.1.2}$$

Solving for $V_\lambda(0)$ yields:

$$V_\lambda(0) = \max_{x \in [0,1]} \frac{\lambda_L (1-x)\Pi + \mu x V_\lambda(1) - c}{\lambda_L (1-x) + \mu x + \lambda}.$$

Differentiating the objective function with respect to x:

$$\frac{\lambda_L (\lambda \Pi + c)(\lambda_\star - \lambda)}{(\lambda + \lambda_H)[\lambda + (1-x)\lambda_L + x\mu]^2}. \tag{OA.1.3}$$

Therefore, the objective function is monotone in x and the sign of the first derivative is

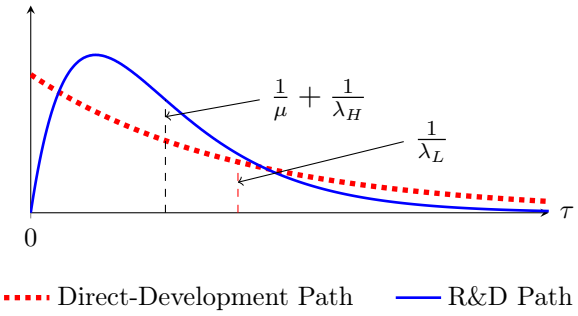


Figure OA.1: Probability distribution functions of a firm's development time

determined solely by the sign of $\lambda_\star - \lambda$. Thus, the optimal policy is to research ($\sigma_t = 1$) when $\lambda < \lambda_\star$, and to directly develop when $\lambda > \lambda_\star$. □

To illustrate the intuition behind this proposition, Figure OA.1 shows the probability distributions of development times with the direct-development path (red dotted curve) and the R&D path (blue solid curve). The direct-development path is more likely to produce a successful product in a short time frame, as it requires only one breakthrough. In contrast, the R&D path produces a successful product after two breakthroughs and, although it has a shorter expected development time, is less likely to lead to quick development. Therefore, when facing a high termination rate, the firm prefers direct development to increase the probability of early success. Under a low termination rate, the firm prefers the R&D path to lower the expected development time.

OA.2 Non-stationary Extension: Private Progress

In the proof of Proposition 2, we restricted the strategy space to stationary (time-independent) allocations. In this section, we extend our analysis to the case of non-stationary allocations. We formulate the duration-minimization problem (Efficiency) and Firm A's payoff-maximization problem (Equilibrium) over the unrestricted space of all measurable, time-dependent strategies $\sigma : \mathbb{R}_+ \to [0, 1]$. We then show that the same conclusion as in Proposition 2 holds.

OA.2.1 Efficiency (Social Planner's Problem)

We formulate the social planner's duration-minimization problem over all measurable allocations $\sigma_A, \sigma_B : \mathbb{R}_+ \to [0,1]$. The planner's joint control is captured through the distribution of race phases, where $P_{jk}(t)$ denotes the probability that $N_t^A = j$ and $N_t^B = k$. Given the allocation pair $(\sigma_A(t), \sigma_B(t))$, the state dynamics are:

$$\dot{P}_{00} = -\big[2\lambda_L + (\mu - \lambda_L)(\sigma_A + \sigma_B)\big]P_{00}, \tag{OA.2.1}$$

$$\dot{P}_{10} = \mu\sigma_A P_{00} - \big[\lambda_H + \lambda_L + (\mu - \lambda_L)\sigma_B\big]P_{10}, \tag{OA.2.2}$$

$$\dot{P}_{01} = \mu\sigma_B P_{00} - \big[\lambda_L + (\mu - \lambda_L)\sigma_A + \lambda_H\big]P_{01}, \tag{OA.2.3}$$

$$\dot{P}_{11} = \mu\sigma_A P_{01} + \mu\sigma_B P_{10} - 2\lambda_H P_{11}. \tag{OA.2.4}$$

Since these phases are mutually exclusive, the race is ongoing at time t with probability $\sum_{jk} P_{jk}(t)$, and the planner minimizes the expected duration $J(\boldsymbol{\sigma}) = \int_0^\infty \sum_{jk} P_{jk}(t)\, dt$.

Proof of Proposition 2(a). We show via a verification argument that the constant policy $\sigma_A^* \equiv \sigma_B^* \equiv 1$ is the unique global minimizer of J.[17]

Let W_{jk}^* denote the stationary expected remaining durations from phase jk under $\boldsymbol{\sigma}^*$:

$$W_{11}^* = \frac{1}{2\lambda_H}, \qquad W_{10}^* = W_{01}^* = \frac{1 + \mu W_{11}^*}{\lambda_H + \mu}, \qquad W_{00}^* = \frac{1 + 2\mu W_{10}^*}{2\mu}, \tag{OA.2.5}$$

and define the switching functions $\Phi_A(\mathbf{P}) = c_1 P_{00} + c_2 P_{01}$ and $\Phi_B(\mathbf{P}) = c_1 P_{00} + c_2 P_{10}$, with

$$c_1 := (\mu - \lambda_L)W_{00}^* - \mu W_{10}^*, \qquad c_2 := (\mu - \lambda_L)W_{01}^* - \mu W_{11}^*.$$

Step 1: A global identity. Consider the linear function $V(\mathbf{P}) := \sum_{jk} W_{jk}^* P_{jk}$. Substituting (OA.2.1)–(OA.2.4), the derivative $\dot{V} = \sum_{jk} W_{jk}^* \dot{P}_{jk}$ is affine in (σ_A, σ_B); the recursions (OA.2.5) are exactly the conditions under which $\dot{V} = -\sum_{jk} P_{jk}$ at $\sigma_A = \sigma_B = 1$, and

[17] The standard sufficiency route via Arrow's theorem does not apply here: the maximized Hamiltonian is a pointwise maximum of functions affine in $\mathbf{P}$, hence convex, and whenever the fall-back condition holds the switching function changes sign on $\mathbb{R}^4_+$, rendering it strictly non-concave. The identity (OA.2.6) circumvents the need for concavity.

collecting the remaining terms gives

$$\dot{V}(\mathbf{P}(t)) = -\sum_{jk} P_{jk}(t) + (1-\sigma_A(t))\Phi_A(\mathbf{P}(t)) + (1-\sigma_B(t))\Phi_B(\mathbf{P}(t)).$$

Because the minimum exit rate $2\lambda_L > 0$, every admissible trajectory satisfies $\mathbf{P}(t) \to \mathbf{0}$, so $V(\mathbf{P}(t)) \to 0$. Integrating from $\mathbf{P}(0) = (1,0,0,0)$, where $V(\mathbf{P}(0)) = W_{00}^*$, yields

$$J(\boldsymbol{\sigma}) = W_{00}^* + \int_0^\infty \big[(1-\sigma_A(t))\Phi_A(\mathbf{P}(t)) + (1-\sigma_B(t))\Phi_B(\mathbf{P}(t))\big]\, dt. \tag{OA.2.6}$$

Step 2: Nonnegativity of the switching functions. Write $\Phi_A(\mathbf{P}(t)) = P_{00}(t)\big[c_1 + c_2 r(t)\big]$ with $r(t) := P_{01}(t)/P_{00}(t) \geq 0$. We first show $c_1 > 0$ unconditionally. Expanding,

$$c_1 = \frac{\mu\lambda_L\lambda_\star + \lambda_L\lambda_H(\lambda_\star - \mu)}{2\mu\lambda_H(\lambda_H + \mu)};$$

the numerator is positive from $\lambda_\star > \mu$. Hence $c_1 > 0$ for all admissible parameters.

The coefficient $c_2 = \frac{\lambda_L(\lambda_\star - \lambda_H)}{2\lambda_H(\lambda_H+\mu)}$ has the sign of the fall-back condition: $c_2 < 0$ if and only if $\lambda_\star < \lambda_H$.

If the fall-back condition fails, then $c_2 \geq 0$ and $\Phi_A(\mathbf{P}) = c_1 P_{00} + c_2 P_{01} \geq 0$ immediately.

If the fall-back condition holds, then $\lambda_H > \lambda_\star > \mu$ by Assumption (2.2), so $\bar{r} := \mu/(\lambda_H - \mu) > 0$ is well defined. Along any admissible trajectory, $r(0) = 0$ and $\dot{r} = \mu\sigma_B - r\big[(\lambda_H - \lambda_L) - (\mu - \lambda_L)\sigma_B\big]$; at $r = \bar{r}$, $\dot{r} = \frac{\mu(\lambda_H - \lambda_L)}{\lambda_H - \mu}(\sigma_B - 1) \leq 0$ for every $\sigma_B \leq 1$, so $r(t) \leq \bar{r}$ throughout. Since $c_2 < 0$, the affine map $r \mapsto c_1 + c_2 r$ attains its minimum on $[0, \bar{r}]$ at $\bar{r}$, where

$$c_1 + c_2\bar{r} = \frac{\lambda_H\lambda_L(\lambda_\star - \mu)}{2\mu(\lambda_H^2 - \mu^2)} > 0$$

by Assumption (2.2). Hence $\Phi_A(\mathbf{P}(t)) \geq 0$. By symmetry, $\Phi_B(\mathbf{P}(t)) \geq 0$ in both cases.

Conclusion. Since $1 - \sigma_i(t) \geq 0$ and $\Phi_i(\mathbf{P}(t)) \geq 0$ for $i \in \{A, B\}$, the integrand in (OA.2.6) is nonnegative, so $J(\boldsymbol{\sigma}) \geq W_{00}^*$ for every admissible allocation. As $P_{00}(t) > 0$ for all finite t, the bracket $c_1 + c_2 r(t)$ is in fact strictly positive, so equality holds if and only if $\sigma_A(t) = \sigma_B(t) = 1$ for almost every t. Hence $\sigma_A^* \equiv \sigma_B^* \equiv 1$ is the unique global minimizer. □

OA.2.2 Equilibrium (Firm A's Best Response)

We now formulate Firm A's payoff-maximization problem over the unrestricted space of all measurable, time-dependent strategies $\sigma_A : \mathbb{R}_+ \to [0,1]$. Suppose Firm B plays the constant allocation strategy $\sigma_B^*(t) = 1$. From Firm A's perspective, Firm B's hidden state probabilities $P_{B0}(t)$ and $P_{B1}(t)$ are exogenous and evolve as:

$$P_{B0}^*(t) = e^{-\mu t}, \quad P_{B1}^*(t) = \frac{\mu}{\lambda_H - \mu}\big(e^{-\mu t} - e^{-\lambda_H t}\big). \tag{OA.2.7}$$

The probability that Firm B has not yet finished is $S_B^*(t) = P_{B0}^*(t) + P_{B1}^*(t)$.

If Firm A discovers the new technology at time t, its expected payoff relies on Firm B's state at that exact instant. Let V_{10}^* and V_{11}^* be the stationary expected continuation values of racing in development against a rival in research or development, respectively:

$$V_{11}^* = \frac{\lambda_H \Pi - c}{2\lambda_H}, \qquad V_{10}^* = \frac{\lambda_H \Pi + \mu V_{11}^* - c}{\lambda_H + \mu} = \frac{2\lambda_H + \mu}{\lambda_H + \mu} V_{11}^*.$$

Firm A's conditionally expected continuation payoff upon discovery at time t is $U_A^*(t) = P_{B0}^*(t)V_{10}^* + P_{B1}^*(t)V_{11}^*$.

Let the state variable $R_A(t)$ denote the probability that Firm A has not yet finished the race, which evolves as $\dot{R}_A(t) = -\big[\lambda_L(1-\sigma_A(t)) + \mu\sigma_A(t)\big]R_A(t)$. Firm A's objective is to maximize:

$$\max_{\sigma_A} \int_0^\infty R_A(t)\Big[\lambda_L(1-\sigma_A(t))S_B^*(t)\Pi + \mu\sigma_A(t)U_A^*(t) - S_B^*(t)c\Big]dt. \tag{OA.2.8}$$

Proof of Proposition 2 (b). We again apply Arrow's sufficiency condition. Let $\eta_A(t)$ be the co-state variable associated with $R_A(t)$. The Hamiltonian is:

$$\begin{aligned} H(\sigma_A(t), R_A(t), \eta_A(t)) = {} & R_A(t)\Big[\lambda_L(1-\sigma_A(t))S_B^*(t)\Pi + \mu\sigma_A(t)U_A^*(t) - S_B^*(t)c\Big] \\ & - \eta_A(t)R_A(t)\big[\lambda_L + (\mu - \lambda_L)\sigma_A(t)\big]. \end{aligned} \tag{OA.2.9}$$

We construct the candidate optimal paths $(\eta_A^*(t), R_A^*(t))$ generated by the constant policy

$\sigma_A^*(t) = 1$. The state trajectory is $R_A^*(t) = e^{-\mu t}$. Let the co-state variable be Firm A's conditionally expected continuation value under the constant profile $(1, 1)$, which is exactly $\eta_A^*(t) = P_{B0}^*(t)V_{00}^* + P_{B1}^*(t)V_{01}^*$, where

$$V_{01}^* = \frac{\mu V_{11}^* - c}{\lambda_H + \mu}, \qquad V_{00}^* = \frac{\mu V_{10}^* + \mu V_{01}^* - c}{2\mu} = \frac{1}{2}(V_{10}^* + V_{01}^*) - \frac{c}{2\mu}.$$

Arrow's sufficiency condition requires verifying four properties:

1. **Maximum Principle:** For all $t \geq 0$, the policy $\sigma_A^* = 1$ maximizes the Hamiltonian. Grouping the terms in the Hamiltonian containing $\sigma_A(t)$ and substituting $\eta_A^*(t)$, we obtain Firm A's switching function:

$$\begin{aligned}\Phi_A(t) =& -\lambda_L S_B^*(t)\Pi + \mu U_A^*(t) - (\mu - \lambda_L)\eta_A^*(t) \\ =& P_{B0}^*(t)\underbrace{\left[\mu V_{10}^* - \lambda_L \Pi - (\mu - \lambda_L)V_{00}^*\right]}_{\Delta_{00}^*} + P_{B1}^*(t)\underbrace{\left[\mu V_{11}^* - \lambda_L \Pi - (\mu - \lambda_L)V_{01}^*\right]}_{\Delta_{01}^*},\end{aligned} \tag{OA.2.10}$$

   where $\sigma_A^* = 1$ maximizes the Hamiltonian if and only if $\Phi_A(t) \geq 0$.

   By substituting (OA.2.7) into (OA.2.10), we can derive the exact algebraic form:

$$\Phi_A(t) = e^{-\mu t}\left[\Delta_{00}^* + \frac{\mu}{\lambda_H - \mu}\left(1 - e^{-(\lambda_H - \mu)t}\right)\Delta_{01}^*\right].$$

   With some algebra, we can derive that

$$\begin{aligned}\Phi_A(t) =& \frac{\lambda_L(\lambda_\star - \mu)}{2\mu}\left(\frac{\mu}{\lambda_H + \mu}\Pi + \frac{c}{\lambda_H}\right)e^{-\mu t} \\ &+ \frac{\lambda_L \mu}{2(\lambda_H + \mu)}\left[\frac{\lambda_\star - \mu}{\lambda_H - \mu}(1 - e^{-(\lambda_H - \mu)t} + e^{-(\lambda_H - \mu)t})\right]\left(\Pi + \frac{c}{\lambda_H}\right)e^{-\mu t}.\end{aligned}$$

   Since $\lambda_\star > \mu$ and $\lambda_H > \mu$, $\Phi_A(t) > 0$ for all $t \geq 0$. Therefore, $\sigma_A^* = 1$ maximizes the Hamiltonian point-wise.

2. **Evolution of the co-state:** The co-state must satisfy $\dot{\eta}_A^*(t) = -\frac{\partial H}{\partial R_A}$ evaluated at

$\sigma_A^* = 1$. This yields the differential equation:

$$\dot{\eta}_A^*(t) = \mu\eta_A^*(t) - \big[\mu U_A^*(t) - S_B^*(t)c\big].$$

Note that $\dot{P}_{B0}^*(t) = -\mu P_{B0}^*(t)$ and $\dot{P}_{B1}^*(t) = \mu P_{B0}^*(t) - \lambda_H P_{B1}^*(t)$. Using these evolutions and the definitions of V_{00}^* and V_{01}^*, we can verify that the above equation holds.

3. **Transversality Condition:** We must show $\lim_{t\to\infty} \eta_A^*(t)(R_A^*(t) - R_A(t)) \leq 0$ for any feasible trajectory $R_A(t)$. For any admissible policy, the probabilities $R_A(t)$ and $R_A^*(t)$ are bounded in $[0, 1]$ and decay to 0 as $t \to \infty$ due to the minimum baseline hazard rate $\lambda_L > 0$. Because $\eta_A^*(t)$ represents bounded expected continuation payoffs, the limit evaluates exactly to 0, satisfying the condition.

4. **Concavity:** The maximized Hamiltonian

$$\hat{H}(R_A(t), \eta_A(t)) = \max_{\sigma_A(t)\in[0,1]} H(\sigma_A(t), R_A(t), \eta_A(t))$$

must be concave in $R_A(t)$. As explicitly shown in (OA.2.9), the Hamiltonian is strictly linear in the state variable $R_A(t)$, thereby satisfying weak concavity.

Because all four conditions hold, Arrow's sufficiency theorem guarantees that the constant policy $\sigma_A^*(t) = 1$ is Firm A's global best response among all measurable, time-dependent allocation strategies. Symmetrically, Firm B's global best response to Firm A's research strategy is also the research strategy. Therefore, the profile where both firms research constitutes an equilibrium. □

OA.2.3 Optimality of Constant Resource Allocation Strategies

In the main text (Proof of Theorem 1), we restricted attention to constant resource allocation strategies. In this appendix, we use optimal control theory to rigorously prove that if a constant allocation strategy maximizes a firm's stationary expected payoff, it is a global best response among all measurable, time-dependent strategies.

Consider the state where neither firm has discovered the new technology. Suppose Firm B uses a constant allocation strategy $b \in [0,1]$. Firm A employs a potentially time-dependent, measurable allocation policy $\sigma : \mathbb{R}_+ \to [0,1]$.

Let V_A and V_B denote Firm A's continuation payoffs if Firm A or Firm B discovers the new technology first, respectively. Under the candidate equilibrium, both firms immediately patent upon discovery, locking in the licensing subgame payoffs: $V_A = V_{11} + l^*$ and $V_B = V_{11} - l^*$. The flow rate at which Firm A develops the old technology is $\lambda_L(1-\sigma_t)$, and the rates at which Firm A and Firm B discover the new technology are $\mu\sigma_t$ and μb, respectively.

Let $\Sigma_t = \int_0^t \sigma_s ds$ be Firm A's cumulative research effort. The probability that the race is still ongoing at time t (the survival function) is:

$$r_t = \exp\left(- \int_0^t \left\{ \mu(\sigma_s + b) + \lambda_L(2 - \sigma_s - b) \right\} ds \right). \tag{OA.2.11}$$

Equivalently, the state variable r_t evolves according to the differential equation:

$$\dot{r}_t = - \left\{ \mu(\sigma_t + b) + \lambda_L(2 - \sigma_t - b) \right\} \cdot r_t, \quad \text{with } r_0 = 1. \tag{OA.2.12}$$

Firm A's objective is to choose a policy σ_t to maximize its expected payoff from the race:

$$J(\sigma) = \int_0^\infty \left\{ \lambda_L(1-\sigma_t)\Pi + \mu\sigma_t V_A + \mu b V_B - c \right\} \cdot r_t \, dt. \tag{OA.2.13}$$

To link this dynamic problem to the stationary formulation in the main text, let $N(x,b)$ denote the flow payoff term inside the brackets of (OA.2.13) evaluated at a constant x, and let $D(x,b)$ denote the aggregate hazard rate term inside the brackets of (OA.2.12). If Firm A also plays a constant strategy $\sigma_t = x$, the expected payoff reduces exactly to the stationary form $U(x) = \frac{N(x,b)}{D(x,b)}$.

Lemma OA.2.1. *Suppose there exists $x^* \in [0,1]$ such that $x^* \in \arg\max_{x \in [0,1]} U(x)$. Then the constant policy $\sigma_t^* = x^*$ for all $t \geq 0$ is a global maximum of the optimal control problem defined by (OA.2.13) and (OA.2.12).*

Proof of Lemma OA.2.1. We apply Arrow's sufficiency condition for infinite-horizon optimal control problems (e.g., Seierstad and Sydsaeter, 1987, Theorem 3.14). Let η_t be the co-state

variable associated with r_t. The Hamiltonian is:

$$\begin{aligned} H(\sigma_t, r_t, \eta_t) &= N(\sigma_t, b) \cdot r_t - \eta_t \cdot D(\sigma_t, b) \cdot r_t \\ &= \{N(\sigma_t, b) - \eta_t D(\sigma_t, b)\} \cdot r_t. \end{aligned} \tag{OA.2.14}$$

We construct the candidate optimal paths (η_t^*, r_t^*) generated by the constant policy $\sigma_t^* = x^*$. Let the co-state variable be constant: $\eta_t^* = U(x^*)$ for all $t \geq 0$. The state trajectory under x^* is $r_t^* = e^{-D(x^*,b)t}$.

Arrow's sufficiency condition requires verifying four properties:

1. **Maximum Principle:** For all $t \geq 0$, $\sigma_t^* = x^*$ maximizes the Hamiltonian. Substituting $\eta_t^* = U(x^*)$, we have:

$$H(\sigma_t, r_t^*, \eta_t^*) = \{N(\sigma_t, b) - U(x^*)D(\sigma_t, b)\} \cdot r_t^*. \tag{OA.2.15}$$

   Because x^* maximizes $U(x) = N(x, b)/D(x, b)$, we know that for any $\sigma_t \in [0, 1]$, $U(x^*) \geq N(\sigma_t, b)/D(\sigma_t, b)$. Since $D(\sigma_t, b) > 0$, this implies $N(\sigma_t, b) - U(x^*)D(\sigma_t, b) \leq 0$. Thus, $H(\sigma_t, r_t^*, \eta_t^*) \leq 0$ for all admissible σ_t. Furthermore, evaluating at $\sigma_t = x^*$ yields exactly 0. Therefore, $\sigma_t^* = x^*$ maximizes the Hamiltonian point-wise.

2. **Evolution of the co-state:** The co-state must satisfy $\dot{\eta}_t^* = -\frac{\partial H}{\partial r_t}(x^*, r_t^*, \eta_t^*)$. Because $H(x^*, r_t^*, \eta_t^*) = 0$ for all r_t^*, its partial derivative with respect to r_t is 0. This is perfectly consistent with our construction that $\dot{\eta}_t^* = 0$.

3. **Transversality Condition:** We must show $\lim_{t\to\infty} \eta_t^*(r_t^* - r_t) \leq 0$ for any feasible state trajectory r_t. Note that for any admissible policy σ, $r_t \geq 0$ and $r_t^* \geq 0$. Furthermore, because $D(\sigma_t, b) \geq 2\lambda_L > 0$, both r_t and r_t^* converge strictly to 0 as $t \to \infty$. Since $\eta_t^* = U(x^*)$ is a finite constant, the limit evaluates exactly to 0, satisfying the condition.

4. **Concavity:** The maximized Hamiltonian $\hat{H}(r_t, \eta_t) = \max_{\sigma_t \in [0,1]} H(\sigma_t, r_t, \eta_t)$ must be concave in r_t. As seen in (OA.2.14), the Hamiltonian is strictly linear in r_t, satisfying weak concavity.

Because all four conditions hold, Arrow's sufficiency theorem guarantees that the constant policy $\sigma_t^* = x^*$ is a global best response among all measurable allocation strategies. □

OA.3 Proofs for Section 5

In this section, we provide the proofs for the results presented in Section 5. The exposition is organized as follows: Section OA.3.1 formalizes the evolution of firms' beliefs and develops a recursive formulation of the patenting optimization problem. Section OA.3.2 proves that any equilibrium patenting strategy must take the form of a gradual patenting strategy (Proposition 4). Building on this structure, Section OA.3.3 contains the proof of the main equilibrium characterization (Theorem 2). Proofs of all intermediate lemmas are relegated to Section OA.3.4.

OA.3.1 Belief Evolution and Recursive Formulation

Belief Evolution Suppose that a firm employs the patent filing policy $G = \{G_s\}_{s\geq 0}$. Let $\mathbf{p}_t$ denote the probability that the firm has the new technology at time t (conditional on the firm having neither developed the product nor patented the new technology). The following lemma characterizes the evolution of $\mathbf{p}$.

Lemma OA.3.1. *Suppose that $\boldsymbol{\sigma} = \mathbf{1}$ and the patent filing strategy is G. Then, $\mathbf{p}$ evolves as follows:*

$$d\,\mathbf{p}_t = (1-\mathbf{p}_t)\cdot\left[(\mu - \lambda_H\,\mathbf{p}_t)dt - d\mathcal{P}_t\right], \tag{OA.3.1}$$

where $d\mathcal{P}_t$ represents the patenting intensity—the instantaneous probability that the firm files a patent at time t—given by

$$d\mathcal{P}_t = \mu(1-\mathbf{p}_t)\left[G_t(t)dt + \int_0^t e^{-(\lambda_H-\mu)(t-s)}dG_s(t)ds\right].^{18} \tag{OA.3.2}$$

[18] We interpret $dG_s(t)$ in the Lebesgue-Stieltjes sense. If G_s is continuous at t, then $dG_s(t) = g_s(t)dt$, where $g_s(t)$ denotes the right-derivative of G_s. If G_s is discontinuous at t, then $dG_s(t)$ represents the probability mass of the atom, $G_s(t) - G_s(t^-)$.

Recursive Formulation Now suppose that Firm A discovers the new technology at time τ, chooses the patent filing policy $\tilde{G}_\tau$—a CDF over $[\tau, \infty]$—and Firm B uses the patent filing policy $G := \{G_s\}_{s\geq 0}$.

Given the resource allocation strategy $\boldsymbol{\sigma} = \mathbf{1}$, the probability that Firm B has neither filed the patent nor developed the product by time t is $e^{-\mu t}/(1-\mathbf{p}_t)$.[19] Define

$$K_t := \frac{e^{-(\lambda_H+\mu)t}}{1-\mathbf{p}_t}. \tag{OA.3.3}$$

Then, the term K_t/K_τ represents the probability that Firm A has not developed the product and Firm B has neither developed the product nor filed a patent by time t.

We first consider the full concealment benchmark: $\tilde{G}_\tau(t) = 0$ for all $t \geq \tau$. Let V_t^n denote the expected payoff in this case. The following lemma characterizes V_t^n.

Lemma OA.3.2. *Suppose Firm B uses the patent filing policy G. Then, Firm A's expected payoff from the full concealment policy is:*

$$V_\tau^n := \frac{1}{K_\tau}\int_\tau^\infty d\mathcal{F}_t \tag{OA.3.4}$$

where

$$d\mathcal{F}_t = (\lambda_H \Pi - c)\cdot K_t\ dt + (V_{11} - (1-\beta)\cdot l^*)\cdot K_t\ d\mathcal{P}_t. \tag{OA.3.5}$$

Additionally, for all $t \geq 0$, $|V_t^n| < \infty$.

Define $\Phi(t) := \{(V_{11} + (1-\beta\,\mathbf{p}_t)\cdot l^*) - V_t^n\}\cdot K_t$ as the *discounted excess payoff* of patenting over concealment—the difference in discounted expected payoffs between immediate patenting and full concealment. The following lemma shows that we can write Firm A's expected payoff (given Firm B's policy) using $\Phi(t)$.

Lemma OA.3.3. *Suppose Firm A discovers the new technology at time τ and employs the patent filing policy $\tilde{G}_\tau$, and Firm B uses the patent filing policy G. Then, Firm A's expected*

[19] Note that Pr(Firm B has not discovered the new technology) $= e^{-\mu t}$ and

$$1-\mathbf{p}_t = \frac{\Pr(\text{Firm B has not discovered the new technology})}{\Pr(\text{Firm B has neither filed a patent nor developed the product})}.$$

payoff at time τ can be written as follows:

$$V_\tau(\tilde{G}_\tau, G) = V_\tau^n + \frac{1}{K_\tau}\left[\tilde{G}_\tau(\tau)\cdot\Phi(\tau) + \int_\tau^\infty \Phi(t)\ d\tilde{G}_\tau(t)\right]. \tag{OA.3.6}$$

This lemma demonstrates a powerful property of the firm's optimization problem: the expected payoff from any arbitrary patenting strategy can be linearly decomposed. It is exactly equal to the baseline payoff of never patenting, V_τ^n, plus the expected incremental value of disclosure. This incremental value is captured by integrating the discounted excess payoff $\Phi(t)$ over the firm's chosen filing distribution $\tilde{G}_\tau$. The first term inside the brackets accounts for the immediate payoff jump if the firm files a discrete probability mass exactly at the moment of discovery, while the integral captures the expected excess benefit of any delayed filings.

The following lemma shows that the support of the optimal patenting policy should be concentrated on the times that maximize the discounted excess payoff.

Lemma OA.3.4. *Let G_τ^* be an optimal patenting policy for a firm that discovered the new technology at time τ: $G_\tau^* \in \arg\max_{\tilde{G}_\tau} V_\tau(\tilde{G}_\tau, G)$.*

(a) If $\sup_{t\geq\tau}\Phi(t) < 0$, then full concealment is optimal: $G_\tau^*(t) = 0$ for all $t \geq \tau$.

(b) If T is in the support of dG_τ^*, then T must maximize the discounted excess payoff:

$$T \in \arg\max_{t\geq\tau}\Phi(t) \qquad \textit{and} \qquad \Phi(T) \geq 0.$$

This result is a direct consequence of the linear payoff decomposition established in Lemma OA.3.3. Because the firm acts to maximize an expectation over the measure $d\tilde{G}_\tau$, its optimal strategy must allocate probability mass exclusively to the moments in time that yield the highest possible excess payoff. If $\Phi(t)$ remains strictly negative for all $t \geq \tau$, then any act of patenting strictly reduces the firm's expected payoff relative to the baseline, making full concealment the uniquely optimal choice. Conversely, if the firm does choose to patent, it will only do so at times T that globally maximize the excess payoff, provided that this maximum is at least zero. This dramatically simplifies our equilibrium analysis: identifying

the optimal timing of patenting is equivalent to tracking the trajectory of $\Phi(t)$ and locating its peaks.

OA.3.2 Gradual Patenting Strategy: Proof of Proposition 4

Now we show that every equilibrium patenting strategy takes the form of a gradual patenting strategy (Proposition 4). We first establish the differential evolution of the excess payoff.

Lemma OA.3.5. *In any symmetric research equilibrium, the discounted excess payoff* Φ *evolves as follows:*

$$\frac{d\Phi(t)}{K_t V_{11}} = (1-\mathbf{p}_t)\lambda_H dt - \rho \cdot \{(\lambda_H(1+\mathbf{p}_t - 2\beta\,\mathbf{p}_t) + \beta\mu(1-\mathbf{p}_t))\,dt + 2(1-\beta)d\mathcal{P}_t\}\,. \quad \text{(OA.3.7)}$$

This lemma decomposes the instantaneous change in the excess payoff of patenting. The evolution is driven by the natural progression of the race and the rival's patenting intensity, $d\mathcal{P}_t$. Crucially, the term scaled by ρ demonstrates that a higher relative value of licensing exerts stronger downward pressure on Φ—making immediate patenting more attractive than delay when ρ is high. Furthermore, the equation reveals that any active patenting by the rival ($d\mathcal{P}_t > 0$) also decreases a firm's own excess payoff—further pressing firms to patent earlier.

Lemma OA.3.6. *In any symmetric research equilibrium,* Φ *has no discrete jump.*

This lemma rules out mass points in the patenting timing. If firms were expected to file a patent at a specific, predictable instant, the simultaneous surge in disclosure would discontinuously destroy the value of exclusivity. Anticipating this sudden drop in value, a firm would strictly prefer to preempt the mass point by filing slightly earlier, thereby smoothing out the distribution of patent filings over time.

Lemma OA.3.7. *In any symmetric research equilibrium, if* $\Phi(T) = \Phi(T')$ *for some* $T < T'$*, then for any* $S \in (T, T')$*, we have* $\Phi(S) \geq \Phi(T) = \Phi(T')$*.*

This result implies that the excess payoff trajectory cannot have "valleys" or be U-shaped between optimal patenting times. During periods where no patenting occurs, a firm's belief $\mathbf{p}_t$ that the rival has the technology strictly increases.

Proof of Proposition 4. By Lemma OA.3.4, the support of any optimal patenting policy G_τ^* must be contained strictly within the set of times that maximize the discounted excess payoff $\Phi(t)$ over $t \geq \tau$.

Lemma OA.3.7 establishes that the set of global maximizers of $\Phi(t)$ on $\mathbb{R}_+$ cannot contain any gaps. Therefore, the set of times that maximize $\Phi(t)$ must form a single connected interval, which we denote as $[T_1, T_2]$ with $0 \leq T_1 \leq T_2 \leq \infty$. Additionally, Φ must increase for $t < T_1$ and decrease for $t > T_2$.

This connected interval perfectly dictates the three phases of the firm's patenting behavior:

(i) **Silent Phase:** For any $t < T_1$, $\Phi(t)$ has not yet reached its maximum. Thus, for any discovery at $\tau < T_1$, the optimal policy assigns zero probability mass to filing before T_1. Therefore, $G_\tau(t) = 0$ for all $t < T_1$.

(ii) **Mixing Phase:** For $t \in [T_1, T_2]$, $\Phi(t)$ is constant at its global maximum. If $T_1 < T_2$, the firm is indifferent regarding exactly when to patent within this interval, allowing for a continuous mixed strategy ($0 < G_\tau(t) < 1$). Lemma OA.3.6 ensures there are no discrete jumps (atoms) in this mixing distribution. Because $\Phi(t)$ strictly decreases after T_2, the optimal window closes, meaning the firm must have patented with certainty by the end of the interval: $G_\tau(T_2) = 1$.

(iii) **Immediate Phase:** For $t > T_2$, $\Phi(t)$ is strictly decreasing. If a firm discovers the new technology at a time $\tau \geq T_2$, the maximum of $\Phi(t)$ on the restricted domain $[\tau, \infty)$ occurs exactly at the boundary $t = \tau$. Thus, the firm's unique best response is to patent immediately upon discovery, yielding $G_\tau(\tau) = 1$ (and generally $G_t(t) = 1$ for all $t \geq T_2$).

Because any symmetric equilibrium patenting strategy must follow this exact sequential structure across the interval $[T_1, T_2]$, it is, by definition, a gradual patenting strategy. □

OA.3.3 Proof of Theorem 2

In this subsection, we present the formal proof of Theorem 2. We begin by establishing a foundational lemma that maps the structural phases of any gradual patenting strategy—defined by its phase transitions (T_1, T_2)—to precise parametric conditions. By evaluating the trajectory of the discounted excess payoff, this lemma provides the necessary conditions that systematically rule out impossible strategy profiles across the parameter space.

Lemma OA.3.8. *Suppose that a gradual patenting strategy with thresholds (T_1, T_2) constitutes a symmetric research equilibrium. Then, the following statements hold:*

(a) if $T_1 = \infty$, then $\rho \leq \rho_D$;

(b) if $T_1 < \infty$ and $T_2 = \infty$, then (i) $\rho \in (\rho_D, \rho_M]$, (ii) for all $t \geq T_1$, $\Phi(t) = 0$ and

$$\mathbf{p}_t = \mathbf{p}^* := 1 - \frac{2\mu(1-\beta)\rho}{\lambda_H(1-\rho) - \mu\beta\rho};$$

(c) if $T_2 < \infty$, then $\rho > \rho_M$.

Proof of Lemma OA.3.8. From Lemma OA.3.5, the differential evolution of the discounted excess payoff is given by:

$$d\Phi = K_t V_{11} \left[(1-\mathbf{p}_t)\lambda_H dt - \rho \cdot \left\{ \left(\lambda_H(1+\mathbf{p}_t - 2\beta\,\mathbf{p}_t) + \beta\mu(1-\mathbf{p}_t)\right) dt + 2(1-\beta) d\mathcal{P}_t \right\} \right].$$

By the definition of the gradual patenting strategy, $\Phi(t)$ attains its global maximum on the interval $[T_1, T_2]$.

(a) Case $T_1 = \infty$ (Full Concealment): Firms never patent, so $d\mathcal{P}_t = 0$ for all $t \geq 0$. For full concealment to be optimal, we must have $\Phi(t) \leq 0$ for all $t \geq 0$. Under $d\mathcal{P}_t = 0$, we have $M_t = e^{-\mu t}$ and $L_t = \frac{\mu}{\lambda_H - \mu}(e^{-\mu t} - e^{-\lambda_H t})$. The belief $\mathbf{p}_t = \frac{L_t}{M_t + L_t}$ simplifies to:

$$\mathbf{p}_t = \frac{\mu(e^{-\mu t} - e^{-\lambda_H t})}{\lambda_H e^{-\mu t} - \mu e^{-\lambda_H t}},$$

which is strictly increasing and converges to μ/λ_H as $t \to \infty$. Calculating the continuation value $V_t^n = \frac{1}{K_t}\int_t^\infty 2\lambda_H V_{11} K_s ds$ yields $V_t^n = V_{11}[1 + \frac{\lambda_H}{\lambda_H + \mu}(1 - \mathbf{p}_t)]$. Substituting this into the

condition $\Phi(t) \leq 0$ gives:

$$\Phi(t) = (1 - \beta\,\mathbf{p}_t)V_{11}K_t \left[\rho - \frac{\lambda_H}{\lambda_H + \mu} \cdot \frac{1 - \mathbf{p}_t}{1 - \beta\,\mathbf{p}_t}\right] \leq 0.$$

Since this must hold for all $t \geq 0$, and the right-hand bracket is strictly decreasing in $\mathbf{p}_t$, the inequality must hold at the limit $t \to \infty$ where $\mathbf{p}_t \to \mu/\lambda_H$. This limit evaluates exactly to $\rho \leq \rho_D$.

(b) Case $T_1 < \infty$ and $T_2 = \infty$ (Delayed Mixing): The mixing phase lasts indefinitely, meaning $\Phi(t)$ is constant at its maximum for all $t \geq T_1$. Thus, $d\Phi(t) = 0$ for $t \geq T_1$. Setting $d\Phi = 0$ in the evolution equation yields:

$$A := \frac{\lambda_H - (\lambda_H + \beta\mu)\rho}{2(1-\beta)\rho} = \frac{\lambda_H\,\mathbf{p}_t}{1 - \mathbf{p}_t} + \frac{d\mathcal{P}_t/dt}{1 - \mathbf{p}_t}. \tag{OA.3.8}$$

Substituting $d\mathcal{P}_t/dt = A(1 - \mathbf{p}_t) - \lambda_H\,\mathbf{p}_t$ into the belief evolution $d\,\mathbf{p}_t = (1 - \mathbf{p}_t)[\mu - \lambda_H\,\mathbf{p}_t - d\mathcal{P}_t/dt]dt$ gives $d\,\mathbf{p}_t = (1 - \mathbf{p}_t)[\mu - A(1 - \mathbf{p}_t)]dt$. Solving this differential equation yields:

$$\mathbf{p}_t = \frac{De^{\mu t} + (A - \mu)}{De^{\mu t} + A}. \tag{OA.3.9}$$

Plugging this back into (OA.3.8) yields the patenting intensity:

$$\frac{d\mathcal{P}_t}{dt} = \frac{(A + \lambda_H)\mu}{De^{\mu t} + A} - \lambda_H.$$

If $D \neq 0$, then $\lim_{t\to\infty} d\mathcal{P}_t/dt = -\lambda_H$, which contradicts the requirement that the patenting intensity must be non-negative ($d\mathcal{P}_t/dt \geq 0$). Thus, $D = 0$. This implies that $\mathbf{p}_t = (A - \mu)/A \equiv \mathbf{p}^*$ for all $t \geq T_1$.

For this belief to be valid ($\mathbf{p}^* \geq 0$), we require $A \geq \mu$, which algebraically simplifies to $\rho \leq \rho_M$. Additionally, because the silent phase lasts until T_1, firms do not patent before T_1. The belief $\mathbf{p}_t$ must therefore be strictly bounded below the no-patenting asymptote μ/λ_H. Thus, $\mathbf{p}^* < \mu/\lambda_H$, which simplifies to $\rho > \rho_D$. This proves (i). Finally, with $D = 0$, $d\mathcal{P}_t/dt = \mu - \lambda_H\,\mathbf{p}^*$. Calculating V_t^n with constant $\mathbf{p}_s = \mathbf{p}^*$ and constant patenting intensity confirms that $V_t^n = V_{11} + (1 - \beta\,\mathbf{p}^*)l^*$, yielding $\Phi(t) = 0$ for all $t \geq T_1$, proving (ii).

(c) Case $T_2 < \infty$ (Immediate Finish): After T_2, the immediate phase begins and $\Phi(t)$ must be strictly decreasing. Thus, $d\Phi(t) < 0$ for $t > T_2$. Because any firm that discovers the technology after T_2 patents it instantly, the stock of unpatented discoveries is zero: $\mathbf{p}_t = 0$. Consequently, the patenting intensity exactly equals the discovery rate: $d\mathcal{P}_t = \mu dt$. Evaluating the evolution equation for $t > T_2$ yields:

$$d\Phi(t) = K_t V_{11} \left[\lambda_H - \rho\big(\lambda_H + \beta\mu + 2(1-\beta)\mu\big)\right] dt < 0.$$

Rearranging this strict inequality yields $\lambda_H < \rho(\lambda_H + \mu(2-\beta))$, which simplifies to $\rho > \rho_M$. □

We are now ready to prove Theorem 2 using this lemma.

Proof of Theorem 2. We prove the theorem by combining the necessary conditions established in Lemma OA.3.8 with the constructive existence of the corresponding patenting strategies.

Part (a): $\rho > \rho_M$ By Lemma OA.3.8 (a), $T_1 = \infty$ requires $\rho \leq \rho_D$, which is violated here. Thus, $T_1 < \infty$. By Lemma OA.3.8 (b), $T_2 = \infty$ requires $\rho \in (\rho_D, \rho_M]$, which is also violated here. Thus, $T_2 < \infty$.

Now, consider any gradual strategy where $0 \leq T_1 < T_2 < \infty$. During the mixing phase $t \in (T_1, T_2)$, the indifference condition requires $d\Phi_t = 0$. Furthermore, because all accumulated discoveries must be patented by T_2, the unpatented discovery pool is fully depleted, meaning $\lim_{t \to T_2^-} \mathbf{p}_t = 0$. Using the belief evolution equation (OA.3.9) evaluated at T_1 and taking the limit as $t \to T_2^-$, we can solve for the integration constant D in two ways:

$$D = \left(\frac{\mu}{1 - \mathbf{p}_{T_1}} - A\right) e^{-\mu T_1} = (\mu - A)\, e^{-\mu T_2}.$$

Because $\mathbf{p}_{T_1} \geq 0$, we have $\frac{\mu}{1-\mathbf{p}_{T_1}} \geq \mu$, which implies $\frac{\mu}{1-\mathbf{p}_{T_1}} - A \geq \mu - A$. Additionally, the condition $\rho > \rho_M$ algebraically simplifies to $A < \mu$, meaning $\mu - A > 0$. Therefore, the right-hand side is strictly positive. However, because $T_2 > T_1$, we also have $e^{-\mu T_1} > e^{-\mu T_2}$.

Multiplying these strict inequalities yields:

$$\left(\frac{\mu}{1-\mathbf{p}_{T_1}}-A\right)e^{-\mu T_1} > (\mu - A)e^{-\mu T_2},$$

which contradicts the equality for D established above. Therefore, no equilibrium can feature $T_1 < T_2$, leaving only $T_1 = T_2$. It remains to exclude a silent-then-immediate switch at a positive time, $0 < T_1 = T_2 < \infty$. During the silent phase $[0, T_1)$ no patents are filed, so the pool of unpatented discoveries grows strictly from zero and $\lim_{t\to T_1^-} \mathbf{p}_t > 0$. Entering the immediate phase at T_1 would require this entire accumulated pool to be filed at the single instant T_1—an atom in the patenting distribution—which by Lemma OA.3.6 cannot occur in a symmetric research equilibrium. Hence $\lim_{t\to T_1^-} \mathbf{p}_t = 0$, forcing $T_1 = 0$. The only remaining possibility is thus $T_1 = T_2 = 0$, the immediate patenting strategy, which by Theorem 1 is mutually optimal when $\rho > \rho_M$; it is therefore the unique equilibrium.

Part (b): $\rho_M > \rho > \rho_D$ By Lemma OA.3.8 (c), $T_2 < \infty$ requires $\rho > \rho_M$, which is violated here. Thus, $T_2 = \infty$. By Lemma OA.3.8 (a), $T_1 = \infty$ requires $\rho \leq \rho_D$, which is also violated. Thus, any equilibrium must feature $0 \leq T_1 < \infty$ and $T_2 = \infty$ (Delayed Mixing). To prove existence and uniqueness of T_1, note that during the silent phase $[0, T_1]$, $d\mathcal{P}_t = 0$. The belief $\mathbf{p}_t$ strictly increases from 0 and asymptotically approaches μ/λ_H. By Lemma OA.3.8(b), the mixing phase requires the belief to be exactly $\mathbf{p}^* = 1 - \frac{2\mu(1-\beta)\rho}{\lambda_H(1-\rho)-\mu\beta\rho}$. Because $\rho \in (\rho_D, \rho_M)$, simple algebra confirms that $\mathbf{p}^* \in (0, \mu/\lambda_H)$. Because $\mathbf{p}_t$ is strictly monotonic, there exists a unique, finite time $T_1 > 0$ such that $\mathbf{p}_{T_1} = \mathbf{p}^*$. For $t \geq T_1$, the firm randomizes such that the patenting intensity exactly sustains $\mathbf{p}_t = \mathbf{p}^*$, holding $\Phi(t) = 0$. This explicitly constructs the unique delayed mixing equilibrium.

Part (c): $\rho < \rho_D$ By Lemma OA.3.8 (b) and (c), any equilibrium with $T_2 < \infty$ requires $\rho > \rho_M$, and any equilibrium with a mixing phase extending to infinity ($T_1 < \infty, T_2 = \infty$) requires $\rho > \rho_D$. Since $\rho < \rho_D < \rho_M$, neither is possible. Thus, the only possible equilibrium structure is full concealment ($T_1 = T_2 = \infty$). To verify existence, if both firms fully conceal, $d\mathcal{P}_t = 0$ for all t. As derived in the proof of Lemma OA.3.8 (a), under zero patenting, $\Phi(t) < 0$ globally if and only if $\rho \leq \rho_D$. Since $\Phi(t)$ remains strictly negative, full concealment is the

unique best response. □

OA.3.4 Proofs of Lemmas

Proof of Lemma OA.3.1. Define

$$M_t := e^{-\mu t},$$
$$L_t := \int_0^t \mu M_s e^{-\lambda_H (t-s)} (1 - G_s(t)) ds,$$

where M_t is the probability that the firm has not discovered the new technology. L_t is the probability that the new technology is discovered but the product has not yet been developed and the patent is not filed yet. Then, $\mathbf{p}_t$ is defined by $L_t/(M_t + L_t)$.

Observe that

$$dM_t = -\mu \cdot (1 - \mathbf{p}_t) dt \cdot (M_t + L_t), \tag{OA.3.10}$$

$$\begin{aligned} dL_t &= \mu M_t (1 - G_t(t)) dt - \lambda_H L_t dt - \int_0^t \mu M_s e^{-\lambda_H (t-s)} dG_s(t) ds \\ &= [\mu(1 - \mathbf{p}_t) dt - \lambda_H \, \mathbf{p}_t \, dt - d\mathcal{P}_t] \cdot (M_t + L_t). \end{aligned} \tag{OA.3.11}$$

By differentiating $\mathbf{p}_t / (1 - \mathbf{p}_t) = L_t / M_t$,

$$\frac{d\,\mathbf{p}_t}{(1 - \mathbf{p}_t)^2} = \frac{\mathbf{p}_t}{1 - \mathbf{p}_t} \left[\frac{dL_t}{L_t} - \frac{dM_t}{M_t} \right] = \frac{\mathbf{p}_t}{1 - \mathbf{p}_t} \left[\frac{\mu(1 - \mathbf{p}_t) dt - \lambda_H \, \mathbf{p}_t \, dt - d\mathcal{P}_t}{\mathbf{p}_t} + \mu dt \right].$$

By rearranging this equation, we obtain (OA.3.1). □

Proof of Lemma OA.3.2. By definition, V_t^n represents Firm A's expected continuation payoff from time t onward under the full concealment policy ($\tilde{G}_t(s) = 0$ for all $s \geq t$).

Under this policy, Firm A receives payoffs only through two events: either (i) Firm A develops the product first (net of cost), yielding $(\lambda_H \Pi - c) dt$; or (ii) Firm B files a patent, yielding $(V_{11} - (1 - \beta) \cdot l^*) d\mathcal{P}_t$. The term K_t / K_τ serves as the effective discount factor, representing the conditional probability that neither firm has ended the race by time t, given

that it was ongoing at time τ. Integrating these flow payoffs yields:

$$V_t^n = \int_\tau^\infty (\lambda_H \Pi - c) \cdot \frac{K_t}{K_\tau}\, dt + \int_\tau^\infty (V_{11} - (1-\beta)\cdot l^*) \cdot \frac{K_t}{K_\tau}\, d\mathcal{P}_t,$$

which is equivalent to (OA.3.4).

Note that

$$\int_t^\infty d\mathcal{F}_s = (\lambda_H \Pi - c) \cdot \int_t^\infty K_s ds + (V_{11} - (1-\beta)\cdot l^*) \cdot \int_t^\infty K_s d\mathcal{P}_s.$$

The first term is finite since $0 \le K_s \le e^{-\lambda_H s}$.

For the second term, observe that $K_s d\mathcal{P}_s$ represents the *discounted* probability density or mass of the firm patenting at time s. Thus, the integral is bounded by the total probability of patenting:

$$\int_t^\infty K_s d\mathcal{P}_s \le \int_t^\infty 1 \cdot d\Pr(\text{The firm patents at } s) \le 1.$$

Since both integrals are finite and $K_t > 0$, the value V_t^n is finite. □

Proof of Lemma OA.3.3. Note that Firm A's expected payoff can be written as follows:

$$\begin{aligned} V_\tau(\tilde{G}_\tau, G) :=& \tilde{G}_\tau(\tau) \cdot (V_{11} + (1-\beta\,\mathbf{p}_\tau)\cdot l^*) + \int_\tau^\infty (V_{11} + (1-\beta\,\mathbf{p}_t)\cdot l^*) \cdot \frac{K_t}{K_\tau}\, d\tilde{G}_\tau(t) \\ &+ \int_\tau^\infty (\lambda_H \Pi - c) \cdot \frac{K_t}{K_\tau} \cdot (1 - \tilde{G}_\tau(t))\, dt \\ &+ \int_\tau^\infty (V_{11} - (1-\beta)\cdot l^*) \cdot \frac{K_t}{K_\tau} \cdot (1 - \tilde{G}_\tau(t))\, d\mathcal{P}_t. \end{aligned}$$

The first two terms correspond to the payoffs when Firm A patents. The third term corresponds to the case where Firm A develops the product first, net of cost. The last term is the payoff when Firm B patents.

Observe that:

$$\begin{aligned}
V_\tau(\tilde{G}_\tau, G) =& \tilde{G}_\tau(\tau) \cdot (V_{11} + (1-\beta\,\mathbf{p}_\tau)\cdot l^*) + \frac{1}{K_\tau}\int_\tau^\infty (1-\tilde{G}_\tau(t))\; d\mathcal{F}_t \\
&+ \frac{1}{K_\tau}\int_\tau^\infty (V_{11} + (1-\beta\,\mathbf{p}_t)\cdot l^*)\cdot K_t\; d\tilde{G}_\tau(t) \\
=& \tilde{G}_\tau(\tau)\cdot (V_{11} + (1-\beta\,\mathbf{p}_\tau)\cdot l^*) + (1-\tilde{G}_\tau(\tau))\cdot V_\tau^n \\
&- \frac{1}{K_\tau}\int_\tau^\infty \int_\tau^t d\tilde{G}_\tau(s) d\mathcal{F}_t + \frac{1}{K_\tau}\int_\tau^\infty (V_{11} + (1-\beta\,\mathbf{p}_t)\cdot l^*)\cdot K_t\; d\tilde{G}_\tau(t) \\
=& V_\tau^n + \tilde{G}_\tau(\tau)\cdot\frac{\Phi(\tau)}{K_\tau} - \frac{1}{K_\tau}\int_\tau^\infty\int_t^\infty d\mathcal{F}_s d\tilde{G}_\tau(t) \\
&+ \frac{1}{K_\tau}\int_\tau^\infty (V_{11} + (1-\beta\,\mathbf{p}_t)\cdot l^*)\cdot K_t\; d\tilde{G}_\tau(t) \\
=& V_\tau^n + \frac{1}{K_\tau}\left[\tilde{G}_\tau(\tau)\cdot\Phi(\tau) + \int_\tau^\infty \Phi(t)\; d\tilde{G}_\tau(t)\right]. \qquad \text{(OA.3.12)}
\end{aligned}$$

Thus, (OA.3.6) holds. □

Proof of Lemma OA.3.4. If $\sup_{t\geq\tau}\Phi(t) < 0$, the above expression is maximized when $\tilde{G}_\tau(\tau) = 0$ and $d\tilde{G}_\tau(t) = 0$ for all $t \geq \tau$. Thus, full concealment is optimal.

Next, consider the support of the optimal policy G_τ^*. Note that the objective function is linear in the probability measure $d\tilde{G}_\tau$. Consequently, the expected payoff is maximized if and only if the probability measure is concentrated entirely on the set of times that maximize the integrand $\Phi(t)$. Thus, if T is in the support of dG_τ^*, it must be that $T \in \arg\max_{t\geq\tau}\Phi(t)$.

Furthermore, if $\Phi(T) < 0$ for such a T, then the maximum value of the integrand is negative. In this case, as shown above, the optimal policy is full concealment (zero measure everywhere), which implies the support is empty (or does not contain T). Hence, for T to be in the support, we must have $\Phi(T) \geq 0$. □

Proof of Lemma OA.3.5. Note that $\Phi(t) = (V_{11} + (1-\beta\,\mathbf{p}_t)\cdot l^*)\cdot K_t - \int_t^\infty d\mathcal{F}_s$. By taking a derivative, we have

$$d\Phi = -\beta l^* K_t d\,\mathbf{p}_t + (V_{11} + (1-\beta\,\mathbf{p}_t)\cdot l^*) dK_t + d\mathcal{F}_t.$$

Using (OA.3.10) and (OA.3.11), we can derive that

$$dK_t = -K_t \left(\lambda_H(1+\mathbf{p}_t)dt + d\mathcal{P}_t\right). \tag{OA.3.13}$$

Using this together with Lemma OA.3.1 and (OA.3.5), we can rewrite $d\Phi$ as follows:

$$\begin{aligned} d\Phi = &- \beta l^* K_t(1-\mathbf{p}_t)\left[(\mu - \lambda_H\,\mathbf{p}_t)dt - d\mathcal{P}_t\right] \\ &- (V_{11} + (1-\beta\,\mathbf{p}_t)l^*)K_t\left[\lambda_H(1+\mathbf{p}_t)dt + d\mathcal{P}_t\right] \\ &+ K_t\left[2\lambda_H V_{11}dt + (V_{11} - (1-\beta)l^*)d\mathcal{P}_t\right]. \end{aligned}$$

By rearranging this equation, we obtain (OA.3.7). □

Proof of Lemma OA.3.6. Suppose that there is a discrete downward jump of Φ at time T. This can happen only when firms patent with a mass at time T, which induces a discontinuous downward jump of Φ at time T:

$$\Delta\Phi(T) = -K_T[2(1-\beta)l^*]\Delta\mathcal{P}_T \leq 0.$$

Since Φ is right-continuous from the right-continuity of G, $\Phi(T)$ cannot be the maximum of Φ, which contradicts that there is a mass of patenting at time T from Lemma OA.3.4. □

Proof of Lemma OA.3.7. Suppose that there exists $S \in (T, T')$ such that $\Phi(S) < \Phi(T) = \Phi(T')$. By Lemma OA.3.6, Φ is continuous. Thus, there exists an interval $[S_1, S_2] \subset (T, T')$ such that $\Phi(t) < \Phi(T)$ for all $t \in [S_1, S_2]$, and $d\Phi(S_1) < 0$ and $d\Phi(S_2) > 0$.

Note that for any $t \in [S_1, S_2]$, since $\Phi(t) < \Phi(T)$, firms do not patent at time t from Lemma OA.3.4. Thus, $d\mathcal{P}(t) = 0$ for all $t \in [S_1, S_2]$ and $\mathbf{p}_{S_1} < \mathbf{p}_{S_2}$. Then, $d\Phi(S_1) < 0$ implies from (OA.3.7) that

$$\rho > \frac{\lambda_H}{\lambda_H + \beta\mu + 2\lambda_H(1-\beta)\frac{\mathbf{p}_{S_1}}{1-\mathbf{p}_{S_1}}} > \frac{\lambda_H}{\lambda_H + \beta\mu + 2\lambda_H(1-\beta)\frac{\mathbf{p}_{S_2}}{1-\mathbf{p}_{S_2}}}.$$

This implies that $d\Phi(S_2) < 0$ from (OA.3.7), which contradicts the assumption that $d\Phi(S_2) > 0$. □

References

Akcigit, U. and Liu, Q. (2016). The role of information in innovation and competition. *Journal of the European Economic Association*, 14(4):828–870.

Anton, J. J. and Yao, D. A. (2004). Little patents and big secrets: Managing intellectual property. *RAND Journal of Economics*, 35(1):1–22.

Baker, S., Lee, P. Y., and Mezzetti, C. (2011). Intellectual property disclosure as threat. *International Journal of Economic Theory*, 7(1):21–38.

Baker, S. and Mezzetti, C. (2005). Disclosure as a strategy in the patent race. *Journal of Law and Economics*, 48(1):173–194.

Bhattacharya, S., Glazer, J., and Sappington, D. E. (1992). Licensing and the sharing of knowledge in research joint ventures. *Journal of Economic Theory*, 56(1):43–69.

Bhattacharya, S. and Guriev, S. (2006). Patents vs. trade secrets: Knowledge licensing and spillover. *Journal of the European Economic Association*, 4(6):1112–1147.

Bobtcheff, C., Bolte, J., and Mariotti, T. (2017). Researcher's dilemma. *Review of Economic Studies*, 84(3):969–1014.

Boleslavsky, R., Kelly, D. L., and Taylor, C. R. (2017). Selloffs, bailouts, and feedback: Can asset markets inform policy? *Journal of Economic Theory*, 169:294–343.

Bryan, K. A. and Lemus, J. (2017). The direction of innovation. *Journal of Economic Theory*, 172:247–272.

Carnehl, C. and Schneider, J. (2023). On risk and time pressure: When to think and when to do. *Journal of the European Economic Association*, 21:1–47.

Chatterjee, K., Das, K., Dong, M., and Klein, N. (2026). Strategic disclosure in research races. *Available at https://ssrn.com/abstract=4803642*.

Curello, G. (2023). Incentives for collective innovation. Available at SSRN: https://arxiv.org/abs/2109.01885.

Das, K. and Klein, N. (2024). Do stronger patents lead to faster innovation? the effect of duplicative search. *International Economic Review*, 65(2):915–954.

d'Aspremont, C., Bhattacharya, S., and Gerard-Varet, L.-A. (2000). Bargaining and sharing innovative knowledge. *Review of Economic Studies*, 67(2):255–271.

Denicolò, V. (2000). Two-stage patent races and patent policy. *RAND Journal of Economics*, 31(3):488–501.

Denicolò, V. and Franzoni, L. A. (2004). Patents, secrets, and the First-Inventor Defense. *Journal of Economics & Management Strategy*, 13(3):517–538.

Denicolò, V. and Franzoni, L. A. (2010). On the winner-take-all principle in innovation races. *Journal of the European Economic Association*, 8(5):1133–1158.

Edmans, A., Goldstein, I., and Jiang, W. (2015). Feedback effects, asymmetric trading, and the limits to arbitrage. *American Economic Review*, 105(12):3766–3797.

Gallini, N. T. (1992). Patent policy and costly imitation. *RAND Journal of Economics*, 23(1):52–63.

Gill, D. (2008). Strategic disclosure of intermediate research results. *Journal of Economics & Management Strategy*, 17(3):733–758.

Green, B. and Taylor, C. R. (2016). Breakthroughs, deadlines, and self-reported progress: Contracting for multistage projects. *American Economic Review*, 106(12):3660–99.

Hall, B., Helmers, C., Rogers, M., and Sena, V. (2014). The choice between formal and informal intellectual property: A review. *Journal of Economic Literature*, 52(2):375–423.

Henao-Restrepo, A. M., Camacho, A., Longini, I. M., et al. (2017). Efficacy and effectiveness of an rVSV-vectored vaccine in preventing Ebola virus disease: Final results from the Guinea ring vaccination, open-label, cluster-randomised trial (Ebola Ça suffit!). *The Lancet*, 389(10068):505–518.

Hopenhayn, H. A. and Squintani, F. (2016). Patent rights and innovation disclosure. *Review of Economic Studies*, 83(1):199–230.

Horstmann, I., MacDonald, G. M., and Slivinski, A. (1985). Patents as information transfer mechanisms: To patent or (maybe) not to patent. *Journal of Political Economy*, 93(5):837–858.

Kim, Y. (2022). Managing a project by splitting it into pieces. Available at SSRN: https://ssrn.com/abstract=3450802.

Kocourek, P. and Kováč, E. (2023). Revealing private information in a patent race.

Krasteva, S. (2014). Imperfect patent protection and innovation. *Journal of Industrial Economics*, 62(4):682–708.

Krasteva, S., Sharma, P., and Wang, C. (2020). Patent policy, imitation incentives, and the rate of cumulative innovation. *Journal of Economic Behavior & Organization*, 178:509–533.

Kultti, K., Takalo, T., and Toikka, J. (2007). Secrecy versus patenting. *RAND Journal of Economics*, 38(1):22–42.

Kwon, I. (2012). Patent races with secrecy. *Journal of Industrial Economics*, 60(3):499–516.

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

Lichtman, D., Baker, S., and Kraus, K. (2000). Strategic disclosure in the patent system. *Vanderbilt Law Review*, 53(6):2175–2217.

Loury, G. C. (1979). Market structure and innovation. *Quarterly Journal of Economics*, 93(3):395–410.

Miyagiwa, K. (2015). The 2011 America Invents Act: Does it undermine innovation? *Journal of Economics & Management Strategy*, 24(2):211–227.

Pardi, N., Hogan, M. J., Porter, F. W., and Weissman, D. (2018). mRNA vaccines — a new era in vaccinology. *Nature Reviews Drug Discovery*, 17(4):261–279.

Ponce, C. J. (2011). Knowledge disclosure as intellectual property rights protection. *Journal of Economic Behavior & Organization*, 80(3):418–434.

Scotchmer, S. and Green, J. (1990). Novelty and disclosure in patent law. *RAND Journal of Economics*, 21(1):131–146.

Seierstad, A. and Sydsaeter, K. (1987). *Optimal control theory with economic applications*. Elsevier North-Holland, Inc.

Song, Y. and Zhao, M. (2021). Dynamic R&D competition under uncertainty and strategic disclosure. *Journal of Economic Behavior & Organization*, 181:169–210.

Spiegel, Y. (2008). Licensing interim R&D knowledge. *CSIO Working Paper*. Available at https://www.econstor.eu/bitstream/10419/38651/1/57493765X.pdf.

Takalo, T. (1998). Innovation and imitation under imperfect patent protection. *Journal of Economics*, 67(3):229–241.

Zhang, T. (2012). Patenting in the shadow of independent discoveries by rivals. *International Journal of Industrial Organization*, 30(1):41–49.