Exploration and Exploitation in R&D Competition*

Alan Jaske†

February 2022

This paper is frequently updated. Click here for the latest version.

Abstract

This paper considers a dynamic model of R&D in which firms navigate a trade-off between exploration (i.e. staying in the patent race) and exploitation (i.e. competing in the market). In the model, greater rivalry in the patent race has an ambiguous effect on equilibrium R&D incentives. On the one hand, there is a higher chance of rival success, which raises R&D incentives through a *racing effect*. On the other hand, there is less rivalry in the product market, which lowers R&D incentives through a novel economic force: the *competition effect*. Once considered, the competition effect has significant implications for the dynamics of firm investment; thus, it provides a rich description of real-world behavior compared to models that consider only racing effects. In terms of expected welfare, the total amount of R&D performed in equilibrium is socially insufficient. A change in market structure, specifically a merger to monopoly, may increase R&D incentives through enhanced appropriability. However, if the trade-off between exploration and exploitation is large, then a merger *always* reduces R&D incentives, regardless of its effect on appropriability.

JEL Classification: C73, D25, O31

Keywords: R&D competition, competitive experimentation, patent race, exploitation versus exploitation, market structure

*I thank Attila Ambrus, Jim Anton, Arjada Bardhi, Gary Biglaiser, Silvana Krasteva, Leslie Marx, Curtis Taylor, Dennis Yao, and Huseyin Yildirim for helpful discussions and valuable feedback. I also thank the participants at the Duke-UNC Micro Theory Seminar and the Duke Micro Theory Workshop many helpful comments.

†Department of Economics, Duke University. Email: alan.jaske@duke.edu.

1 Introduction

In highly innovative industries, one of the most important decisions that firms make is the extent to which resources should be allocated toward (i) the *exploitation* of existing technology, to ensure current competitiveness, or (ii) the *exploration* of new technology, to ensure future competitiveness. This dichotomy has long been recognized in economics and management science under various names – innovation versus invention (Schumpeter, 1939), incremental versus radical innovation (Mansfield et al., 1968), exploration versus exploitation (March, 1991) – and both activities appear to be key determinants of industry performance (Caggese, 2019; He and Wong, 2004) as well as economic growth (Akcigit et al., 2021). Striking a balance between exploitation and exploration is, however, difficult in practice because each demands that scarce resources (e.g. labor, capital, and time) be deployed toward distinct and often incompatible tasks (Lavie et al., 2010). It is therefore an inescapable reality that firms, to effectively compete in markets, must trade off short-term profitability (i.e. exploitation) for long-term innovation (i.e. exploration) and vice versa.

In this paper, I model the trade-off between exploration and exploitation in the context of a dynamic R&D competition, and analyze its implications for competition, welfare, and the fundamental relationship between market structure and innovation. Specifically, I consider a duopolistic industry set in continuous time. The game begins when an exogenous scientific discovery triggers a patent race to develop a cost-reducing innovation, the feasibility of which is uncertain.[1] To explore the innovation's feasibility, a firm must to undertake R&D in the form of costly experimentation (c.f. Choi, 1991; Keller et al., 2005); however, a trade-off exists in that undertaking R&D negatively impacts a firm's ability to produce output at low cost.[2] Profit maximization therefore requires that firms consider not only the *explicit cost* of R&D, but also the *opportunity cost* of R&D (i.e. the value of low-cost production).

Compared to the existing literature, the focus of my paper is to highlight the central role of endogenous product market interaction in shaping firms' incentives to invest (or not invest) in R&D. This provides an economic lens through which both applied and theoretical researchers can interpret behavior in real-world R&D races, and explain why observed empirical patterns have contradicted theoretical predictions (Cohen, 2010; Gilbert, 2006).

To illustrate the empirical relevance of my model, consider, for example, the race to commercialize digital camera technology during the late-1980s and early-1990s. A historical leader in film innovation, the Polaroid Corporation made large initial investments into R&D of digital camera inputs (e.g. fiber optics, solar cells, and disk drives) in a deliberate effort to innovate where contemporary rivals would not. During this time, Polaroid adopted a

[1] I focus on the case of cost-reducing (i.e. process) innovation to draw the most direct comparisons to the existing literature. It would be straightforward, however, to extend my analysis to the case of quality-enhancing (i.e. product) innovation; for example, by adapting the approach of (Vives, 2008).

[2] For the sake of generality, I do not model the source of the trade-off between R&D investment and short-term productivity; however, it can be theoretically motivated by any number of organizational frictions, including financing constraints (e.g. Brown et al., 2012; Giebel and Kraft, 2019), limited managerial attention (e.g. Dosi, 1988; Gifford, 1992), and time constraints (e.g. Becker, 1965; Radner and Rothschild, 1975).

technology-driven approach that explicitly favored exploration at the expense of exploitation.[3]

Polaroid's focus changed in the mid-1990s when the company, as part of a major restructuring effort, slashed R&D spending in favor of expanded investment into print film and marketing of existing products. Then-CEO, Gary DiCamillo, justified Polaroid's decision explicitly in terms of the trade-off between exploration and exploration and the desire to prioritize short-term competitiveness in the product market:

> *"Can we be a down and dirty manufacturer at the same time as we're an innovator over here? Can you have two different philosophies running simultaneously in the company?"*
>
> *"[W]e have to focus on what value added we provide that's unique ... Substitute technology such as inkjet or thermal technologies are interesting, but they're not here yet."*
>
> (Tripsas and Gavetti, 2000, pp. 1155-1157)

At the time, Polaroid executives considered the decision to exit the race for digital photography to be a shrewd gamble. Now, most business schools teach it as a case study in poor strategic management, which directly contributed to the the company's 10-year decline and ultimate bankruptcy. My results provide insight as to why Polaroid, an otherwise innovative and forward-looking company, could have found exit to be optimal despite enjoying a sizable lead in the race for digital photography.[4]

Due to the trade-off between exploration and exploitation, a firm must explicitly compare the long-term benefit of R&D investment (i.e. staying in the patent race) to the short-term benefit of exiting the patent race to more aggressively compete in the market (i.e. produce output at low cost). Equilibrium R&D incentives are therefore determined by the interaction of two distinct effects. First, there is a *racing effect* that measures the extent to which R&D rivalry motivates a firm to invest out of fear that its rival will be first to succeed. In most patent race models, the racing effect is the main force that determines equilibrium R&D incentives. This is because R&D investment is assumed to have no effect on a firm's ability to compete in the product market; hence, the short-term benefit from exiting the patent race is, by definition, equal to zero.[5] In my model, however, the short-term benefit of exiting the

[3] In its 1980 annual report to shareholders, Polaroid summarized its mission as follows: "Do not undertake [a] program unless the goal is manifestly important and its achievement nearly impossible. Do not do anything that anyone else can do readily." According to interviews with senior management, Polaroid's mission was motivated by perceived comparative advantage in R&D: "What we were good at was major inventions. Large-scale, lengthy projects that other firms would hesitate to tackle" (Tripsas and Gavetti, 2000, pp. 1150-1151).

[4] Polaroid's R&D units had been involved in digital photography since the 1960s, and filed numerous patents during this time. Even in 1989, nearly 42 percent of all Polaroid R&D spending was devoted to digital photography (see "The Collapse of Polaroid – 4 Reasons Why Polaroid Failed and What We Can Learn." Predictable Profits. 8 October 2020).

[5] For this reason, most patent race models assume that firms are potential entrants into a new market. The market-entry interpretation of a patent race provides a great deal of tractability; however, it completely abstracts from central focus of this paper – namely, the endogenous relationship between product market competition and R&D investment.

patent race is endogenously determined by the (non-zero) value of low-cost production. Thus, the racing effect is dampened by a countervailing force – the *competition effect* – that measures the extent to which R&D rivalry increases a firm's temptation to prematurely exit patent race in order to benefit from a temporary cost advantage in the product market.

Focusing on the Markov-perfect equilibrium, my model generates a rich set of possible dynamics, the fundamental nature of which depends on the strength of the competition effect relative to the racing effect. If the racing effect dominates the competition effect – for example, because the gains from successful innovation are sufficiently large, or because R&D has little impact on a firm's ability to product output at low cost – then equilibrium R&D incentives resemble that of a typical patent race with hazard-rate uncertainty. In particular, there is a unique Markov-perfect equilibrium in which firms perform R&D until beliefs about the innovation's feasibility become sufficiently pessimistic. Provided that firms are symmetric (which I assume in the base model), the belief threshold at which exit occurs will also be symmetric. Hence, equilibrium R&D investment always features *simultaneous exit* from the patent race in finite time.

However, if the competition effect dominates the racing effect – for example, because R&D has a large impact on a firm's ability to produce output at low cost – then equilibrium R&D incentives depart considerably from a typical patent race. In particular, every Markov-perfect equilibrium always features *sequential exit*. This is because the competition effect causes each firm's incentive to perform R&D to be lowest when its rival also invests (i.e. R&D investments are strategic substitutes). Consequently, if one firm exits the patent race, then the remaining firm experiences an instantaneous increase in its own R&D incentives (because it no longer is tempted by the competition effect). So, it strictly prefers to continue performing R&D.

To characterize the dynamics of equilibrium R&D investment when sequential exit occurs, it suffices to determine each firm's preferred order of exit. Depending on the marginal-cost impact of R&D investment, I show that the continuation subgame after one firm exits may be characterized by an *innovator's curse* in which the remaining firm (i) continues to perform R&D, but (ii) receives a strictly lower expected payoff compared to the firm that exited. If the innovator's curse is insignificant (i.e. exists over a small range of beliefs), then each firm strictly prefers to be the last to exit the patent race; hence, the dynamics of sequential exit closely resemble a war of attrition: each firm aggressively performs R&D in hopes of becoming the sole innovator. However, if the innovator's curse is significant (i.e. exists over a wide range of beliefs), then each firm strictly prefers to be the first to exit. In this case, the dynamics of sequential exit resemble a preemption game: each firm is willing to exit the patent race much earlier than would be otherwise optimal in hopes of becoming the sole non-innovator.

The expected total amount of R&D performed in any Markov-perfect equilibrium is less than what would be chosen by a social planner. It is therefore reasonable to ask how a change in market structure – for example, a merger to monopoly – would positively or negatively affect equilibrium R&D incentives. Arrow (1962) famously argued that, all else equal, a monopolist's incentive to perform R&D should be less than that of a competitive firm because innovation cannibalizes existing monopoly profits. Absent a trade-off between exploration and exploitation, the Arrow replacement effect is necessary and sufficient to conclude that monopoly R&D incentives are lower than duopoly incentives. However, when a trade-off between exploration and exploitation exists, I show that the Arrow replacement effect is no longer necessary to obtain this conclusion. This is because a monopolist, due to its enhanced market power, benefits more from low-cost production than does a duopolist. The trade-off between exploration and exploitation, therefore, generates an *opportunity cost* effect that lowers monopoly R&D incentives by more than standard arguments based on appropriability would suggest. In fact, if the only cost of R&D were the opportunity cost of foregone exploitation, then monopoly R&D incentives are *always* lower than duopoly incentives, regardless of whether or not the gains to successful innovation are higher for the monopolist.

Next, I consider the applications of my model to merger review. Taken literally, my analysis of monopoly and duopoly R&D incentives represents the simplest possible merger – namely, a pure reduction in the number of firms. In reality, however, mergers are approved (or rejected) based on numerous considerations; for example, whether merger-specific cost synergies can offset the harmful effects of enhanced market power. In principle, one could argue that greater cost synergies should enhance post-merger R&D incentives through its effect on appropriability (i.e. increasing the gains to successful innovation). I show that this argument is, in general, incomplete because it ignores the fundamental relationship between cost synergies and the opportunity cost of R&D (i.e. the value of low-cost production). While greater cost synergies may increase the gains to successful innovation, I show that the opportunity cost of R&D always increases by an even larger amount. Thus, my results provide a theoretical foundation for recent arguments made by both US and European competition authorities – namely, that a merger can produce benign (or even positive) market structure effects, but may nevertheless be considered anti-competitive due to its negative effects on equilibrium R&D investment (see Denicolo and Polo, 2018; Jullien and Lefouili, 2018).

The remainder of this paper is organized as follows. In the next section, I give a detailed comparison of my model and its results to the existing literature. I describe my base model in Section 3. In Section 4, I characterize the dynamics of equilibrium R&D investment via backward induction and discuss related comparative statics. Section 5 compares equilibrium R&D investment to that which maximizes present-discounted total surplus. Section 6 then contrasts equilibrium R&D investment to that which would be performed under monopoly. I then extend this monopoly analysis to the case of merger review in the presence of both production and R&D-related cost synergies. The paper concludes in Section 7 with a discussion of model robustness, possible extensions, and directions for future work. All proofs are relegated to the Appendix.

2 Related Literature

First and foremost, my paper contributes to the vast literature on firm investment and dynamic R&D competition and, specifically, patent races. The early foundations of the patent race literature are surveyed in Reinganum (1989); see also Gilbert (2006) for a survey that focuses on the relationship between competition and investment. That being said, the mechanism through which product market and R&D incentives interact – namely, the competition effect – is broadly applicable other modes of dynamic R&D competition, including multi-stage R&D races (e.g. Harris and Vickers, 1987), research tournaments (e.g Taylor, 1995), and step-by-step innovation models of endogenous growth (e.g. Aghion et al., 2001).

Firms in my model allocate resources toward R&D in hopes of producing an innovation of known quality but unknown feasibility. Thus, my paper relates to those in the literature on competitive experimentation. Choi (1991) was among the first to consider experimentation in the context of R&D competition. His model introduces hazard-rate uncertainty into the seminal patent race model of Reinganum (1982), and has since been extended to include a number of additional considerations, such as variable R&D intensity (Malueg and Tsutsui, 1997), unobservable actions (Akcigit and Liu, 2016), and private information (Awaya and Krishna, 2021; Moscarini and Squintani, 2010). My paper contributes to this literature by highlighting the inherent sensitivity of Choi's model to the introduction of a trade-off between exploration and exploitation. In doing so, my results can be used to improve future models and better explain observed patterns of R&D investment in terms of industry characteristics, namely the cost structure of R&D and its impact on current competitiveness.

To the best of my knowledge, the only paper to explicitly consider a trade-off between R&D investment (i.e. exploration) and current competitiveness (i.e. exploitation) in the context of a patent race is Besanko and Wu (2013). In their model, firms decide whether to allocate resources toward (i) R&D to develop a product innovation or (ii) marketing to increase demand for an existing product. However, they make a critical assumption that each firm is a monopolist in its respective market. This proves consequential for Besanko and Wu's equilibrium construction, as it guarantees that firms' strategies depend only on current beliefs, and not the actions of rivals. Equilibrium R&D incentives are, therefore, solely determined by the racing effect – the competition effect is zero, by definition, because firms do not compete in the product market.

Thomas (2021) considers a game of strategic experimentation in which players share a common safe arm can be activated by only one player at a time. In equilibrium, competition for access to the safe arm generates a preemption motive similar manner to my competition effect – while a rival player is experimenting, the opportunity cost of experimentation is positive; however, while a rival player is not experimenting, the opportunity cost of experimentation is zero. However, her model assumes that each player owns a private risky arm. Consequently, there is no rivalry in experimentation, and equilibrium dynamics are completely determined by preemption motives. Furthermore, Thomas's notion of congestion that is exogenous: one player *must* experiment if the other occupies the safe arm. By contrast, the competition effect in my model is fully-endogenous. Furthermore, because I consider an explicit market structure, my model can be used to address a number of questions related to welfare and optimal R&D policy that lie beyond the scope of Thomas's model.

Finally, the opportunity cost effect I the analysis of monopoly R&D incentives can be interpreted as a dynamic version of the "switchover disruption effect" identified Holmes et al. (2012) in the context of new technology adoption. In their model, adoption of a new process innovation results in temporary period of higher costs – the switchover disruption. The key difference is that Holmes et al. consider innovation as a one-time adoption decision, whereas I view it as an inherently dynamic process. Furthermore, while their analysis does not directly consider how cost synergies interact with the switchover disruption effect, it is straightforward to extend their results to show that greater cost synergies (as I to model them) always reduce the size of the switchover disruption effect. This stands in contrast to the opportunity cost effect, which is always increasing in the size of cost synergies. This fundamental difference in our results arises because a monopolist that innovates via a one-time technology adoption decision will be inherently less affected by switchover disruptions than a monopolist that must continuously invest in R&D.

3 The Model

The Market. Two firms, $i \in \{1,2\}$, compete in continuous time over an infinite time horizon. Demand is stationary over time, with inverse demand function $P(Q_t) = a - Q_t$, where $Q_t \equiv q_{1,t} + q_{2,t}$ denotes industry output in period t, and $a > 0$ is a known parameter. The game begins in period 0, when a scientific discovery initiates a patent race to develop an innovation. To participate in the race, a firm must choose to enter in period 0 and begin to allocate a flow stock of productive resources toward R&D investment, which may represent financial outlays (e.g. borrowing, spending) or the allocation of physical resources (e.g. scientists, engineers). Exit from the patent race is costless but irreversible; thus, it is without loss of generality to assume that both firms initially enter the race.

If credit is easily accessible, then financial outlays associated with R&D investment are unlikely to interfere with a firm's ability to compete in the product market. The cost of financial outlays may therefore be captured with a flow cost $f \geq 0$. The allocation of physical resources, however, is much more likely to affect a firm's competitiveness. To capture this possibility, let $\phi \geq 0$ denote the extent to which R&D investment allocates resources away from current operations, thereby raising a firm's marginal cost. Thus, Firm i's total cost in period t may be written as $C(q_{i,t}, e_{i,t}) = (c_A + \phi e_{i,t})q_t + f e_{i,t}$, where $e_{i,t} \in \{0,1\}$ indicates whether or not Firm i remains in the patent race at time t.

The Patent Race. The winner of the patent race obtains exclusive rights to the innovation, which allows it to produce at marginal cost $c_B < c_A$ in all future periods. Both firms know the value of c_B, but they are uncertain as to whether success is possible. I let this uncertainty be represented by an unknown state variable $\theta \in \{0,1\}$ that determines the rate at which R&D produces success over time. If $\theta = 1$, then the innovation is feasible, and R&D generates success at Poisson rate $\lambda > 0$; otherwise, success is impossible.

Initial beliefs about θ, denoted $\beta_0 \equiv \mathbb{P}(\theta = 1)$, are taken to be symmetric. I assume that exit from the patent race is publicly observable; thus, posterior beliefs β_t will remain symmetric for all $t > 0$. Conditional on no success, the formula for β_t is derived using Bayes' rule:

$$\beta_t = \frac{\beta_0 e^{-\lambda \int_0^t (e_{1,s} + e_{2,s}) ds}}{\beta_0 e^{-\lambda \int_0^t (e_{1,s} + e_{2,s}) ds} + (1 - \beta_0)}. \tag{1}$$

Because learning is exponential, observe that β_t is a deterministic function of the total R&D performed up to period t, denoted $K_t = \int_0^t (e_{1,s} + e_{2,s}) ds$. Differentiating β_t with respect to t yields the following law of motion of posterior beliefs (conditional on no success):

$$\frac{d\beta_t}{dt} = -\lambda(e_{1,t} + e_{2,t})\beta_t(1 - \beta_t). \tag{2}$$

Clearly, $d\beta_t/dt < 0$ holds whenever $e_{1,t} + e_{2,t} > 0$. Thus, beliefs about the innovation's feasibility become more pessimistic the longer R&D is performed without success – i.e. "no news is bad news".

Solution Concept. I consider Markov-perfect equilibrium. In continuous-time, it is typical to assume that information and reaction lags. Firms may therefore instantaneously respond to one another's actions. This, however, creates technical issues related to the definition of proper subgames (see Simon and Stinchcombe, 1989). To address these issues, I follow the literature and suppose that a tie-break occurs in the event that simultaneous moves occur.[6] In particular, if both firms simultaneously attempt to exit the patent race, then one firm is randomly selected to be the first-mover in exit. The second-mover may then choose to follow the first-mover or revise its action to stay in the patent race.[7]

With this additional structure, I may now formally define an equilibrium. A pure strategy for Firm i consists of a sequence of output decisions $(q_{i,t})_{t \geq 0}$ and a pair of stopping times (τ_1^i, τ_2^i). The stopping time τ_k^i specifies the date at which Firm i will exit the patent race if the number of remaining firms is k. An equilibrium is a strategy profile (σ_1^*, σ_2^*) such that σ_1^* maximizes Firm 1's expected discounted profit given σ_2^* starting in any period $t \geq 0$, and vice versa for Firm 2. Markov-perfection requires that (σ_1^*, σ_2^*) depend only on payoff-relevant states, namely beliefs and the number of firms in the patent race.

Quantity choices affect only current flow profit. Firms, therefore, engage in period-by-period Cournot competition in any Markov-perfect equilibrium. I assume that each firm produces a positive level of output at all times; that is, I suppose that $2\phi < a - c_A$ and $2c_A < a + c_B$. Hence, the formula for Firm i's flow profit before R&D success occurs is

$$\Pi(e_{i,t}, e_{j,t}) - fe_{i,t}, \quad \text{where} \quad \Pi(e_{i,t}, e_{j,t}) \equiv \frac{(a - c_A + \phi(e_{j,t} - 2e_{i,t}))^2}{9}.$$

After R&D success occurs, the winner and loser of the patent race receive flow profit

$$\Pi^L \equiv \frac{(a - 2c_B + c_A)^2}{9} \quad \text{and} \quad \Pi^F \equiv \frac{(a - 2c_A + c_B)^2}{9},$$

respectively, in all future periods.

[6] See, for example, Hopenhayn and Squintani (2011), Dutta et al. (1995), and Dutta and Rustichini (1993).

[7] In my analysis, I take the tie-breaking rule as given; however, it may be fully micro-founded using a suitable extension of subgame-perfect equilibrium, as in Fudenberg and Tirole (1985).

4 Equilibrium Analysis

4.1 Individual R&D Subgame

We proceed via backward induction, beginning with the individual R&D subgame initiated after one firm, say j, exits the patent race. In this subgame, we refer to Firm i and Firm j as the sole innovator and non-innovator, respectively.

Because exit is irreversible, the analysis of the individual R&D subgame is equivalent to a single-agent decision problem. As long as the sole innovator chooses to remain in the patent race, its expected value, denoted $V_{(1,0)}(\beta)$, satisfies the following Hamilton-Jacobi Bellman (HJB) equation:

$$\rho V_{(1,0)}(\beta) = \Pi(1,0) - f + \beta\lambda\left[V^L - V_{(1,0)}(\beta)\right] - \lambda\beta(1-\beta)V'_{(1,0)}(\beta), \tag{3}$$

where $V^L \equiv \Pi^L/\rho$ denotes the value of winning the patent race.

The terms of Equation (3) highlight three distinct sources of value for the sole innovator; in order, they are flow profit, $\Pi(1,0) - f$, the expected gains from R&D, $\beta\lambda[V^L - V_{(1,0)}(\beta)]$, and the effect of learning, $-\lambda\beta(1-\beta)V'_{(1,0)}(\beta)$. Note that this last term enters negatively into the sole innovator's HJB equation because posterior beliefs about the innovation's feasibility are decreasing over time so long as success does not occur.

To derive the sole innovator's optimal exit strategy, we first note that the general solution to Equation (3) takes the form

$$V_{(1,0)}(\beta) = H_{(1,0)}(\beta) + C_{(1,0)}(1-\beta)\Phi(\beta)^{\frac{\rho}{\lambda}}, \tag{4}$$

where $\Phi(\beta) \equiv (1-\beta)/\beta$, and $C_{(1,0)}$ is a constant of integration.

To interpret this general solution, it is helpful to first define

$$H_{(1,0)}(\beta) \equiv (1-\beta)\left(\frac{\Pi(1,0) - f}{\rho}\right) + \beta\left(\frac{\Pi(1,0) - f + \lambda V^L}{\rho + \lambda}\right).$$

The function $H_{(1,0)}(\beta)$ gives the sole innovator's expected value when it never exits the patent race. In other words, $H_{(1,0)}(\beta)$ reflects the commitment value of R&D. The second term of the general solution for $V_{(1,0)}(\beta)$ therefore gives the option value of stopping. By definition, this option value is positive since $\Pi(1,0) - f < \Pi(0,0)$. Thus, $C_{(1,0)} > 0$, and so $V_{(1,0)}(\beta)$ is strictly convex.

Because $V_{(1,0)}(\beta)$ is strictly convex, it follows from standard arguments (e.g. Strulovici and Szydlowski, 2015) that optimal stopping will occur at some threshold belief β_1^*. If interior, the boundary conditions defining β_1^* are (i) $V_{(1,0)}(\beta_1^*) = V_{(0,0)} \equiv \Pi(0,0)/\rho$ (value matching) and (ii) $V'_{(1,0)}(\beta_1^*) = 0$ (smooth pasting).[8] Inserting these conditions into Equation (3) yields the

[8] Value matching is always necessary; smooth pasting is necessary because $V_{(1,0)}(\beta)$ is convex.

following solution:

$$\beta_1^* = \frac{\rho}{\lambda}\left(\frac{f + C(0)}{\Pi^L - \Pi(0,0)}\right), \tag{5}$$

where $C(0) \equiv \Pi(0,0) - \Pi(1,0)$ denotes the implicit cost of R&D.

Intuitively, the sole innovator will remain in the patent race so long as the expected benefits of R&D relative to benefits of exit, namely $\beta\lambda[\Pi^L - \Pi(0,0)]/\rho$, outweigh the opportunity cost of R&D, $f + C(0)$. The threshold β_1^* represents the belief at which the net expected benefit of additional R&D is exactly zero.

Clearly, there is both an exogenous and endogenous component to β_1^*, which we respectively define as

$$\beta_f^* \equiv \frac{\rho}{\lambda}\left(\frac{f}{\Pi^L - \Pi(0,0)}\right) \quad \text{and} \quad \beta_\phi^* \equiv \frac{\rho}{\lambda}\left(\frac{C(0)}{\Pi^L - \Pi(0,0)}\right).$$

By definition β_f^* is the stopping threshold that the sole innovator would use if all costs of R&D were explicit. Likewise, β_ϕ^* is the stopping threshold that the sole innovator would use if all costs of R&D were implicit. Within our Cournot specification, this latter threshold may be written as

$$\beta_\phi^* = \underbrace{\frac{\rho}{\lambda}\left(\frac{\phi}{c_A - c_B}\right)}_{\equiv\,\beta^c}\underbrace{\left(\frac{a - c_A - \phi}{a - c_B}\right)}_{\equiv\,\Gamma}.$$

The term β^c describes the sole innovator's incentive to perform R&D if its sole objective were cost-minimization.[9] In general, however, the sole innovator wishes to maximize profit. So the additional term

$$\Gamma \equiv \frac{a - c_A - \phi}{a - c_B} = \frac{\partial\Pi(1,0)/\partial\phi}{\partial\Pi^L/\partial c_B}$$

appears to account for the difference between these two objectives. Observe that $\Gamma < 1$ holds in all relevant cases. Thus, profit motives lead the sole innovator to perform more R&D than it would under cost-minimization. This is because flow profit is strictly convex in own marginal cost; hence, there is value in waiting to exit (c.f. McDonald and Siegel, 1986).

4.1.1 Innovator (Dis)advantage in R&D

Having characterized the sole innovator's optimal exit strategy, we may now compare each firm's respective payoffs in the individual R&D subgame. In doing so, we establish our first main result, Proposition 1, which states that the most innovative firm (in terms of R&D investment) need not be the most profitable in equilibrium. Specifically, if the implicit cost of

[9]Indeed, if the sole innovator were interested in minimizing present discounted marginal cost, the indifference condition for optimal stopping would be $c_A = c_A + \phi + \beta\lambda[c_B - c_A]/\rho$, which solves to give $\beta = \beta^c$, as desired.

R&D, as measured by ϕ, is sufficiently large, then there exists a nonempty interval of beliefs over which (i) individual R&D occurs with probability one, and (ii) the sole innovator receives a strictly lower expected payoff than the sole non-innovator.

To establish this result, let us first compute the sole innovator's expected payoff before beliefs reach β_1^*. By definition, $V_{(1,0)}(\beta)$ satisfies Equation (4). So all that must be determined is the constant of integration $C_{(1,0)}$. This can be done by inserting the value-matching condition $V_{(1,0)}(\beta_1^*) = V_{(0,0)}$ into Equation (4) and solving for $C_{(1,0)}$. This yields the following exact solution for the sole innovator's expected payoff over $(\beta_1^*, 1)$:

$$V_{(1,0)}^*(\beta) = H_{(1,0)}(\beta) + \frac{(1-\beta)\Phi(\beta)^{\frac{\rho}{\lambda}}}{(1-\beta_1^*)\Phi(\beta_1^*)^{\frac{\rho}{\lambda}}} \left(V_{(0,0)} - H_{(1,0)}(\beta_1^*)\right). \tag{6}$$

To compute the non-innovator's expected payoff, denoted $V_{(0,1)}(\beta)$, we proceed in a similar fashion. For all $\beta \in (\beta_1^*, 1)$, the HJB equation defining $V_{(0,1)}(\beta)$ is

$$\rho V_{(0,1)}(\beta) = \Pi(0,1) + \beta\lambda \left[V^F - V_{(0,1)}(\beta)\right] - \lambda\beta(1-\beta)V_{(0,1)}'(\beta). \tag{7}$$

The general solution to this equation takes the form

$$V_{(0,1)}(\beta) = H_{(0,1)}(\beta) + C_{(0,1)}(1-\beta)\Phi(\beta)^{\frac{\rho}{\lambda}}, \tag{8}$$

where

$$H_{(0,1)}(\beta) \equiv (1-\beta)\left(\frac{\Pi(0,1)}{\rho}\right) + \beta\left(\frac{\Pi(0,1) + \lambda V^F}{\rho+\lambda}\right),$$

and $C_{(0,1)}$ is a constant of integration.

The components of the general solution for $V_{(0,1)}(\beta)$ are analogous to the components of $V_{(1,0)}(\beta)$. The function $H_{(0,1)}(\beta)$ gives the sole non-innovator's expected value if the sole innovator were committed to remain in the patent race indefinitely. Thus, $C_{(0,1)}(1-\beta)\Phi(\beta)^{\frac{\rho}{\lambda}}$ measures the additional value that the sole non-innovator received when the sole innovator has the option to exit. The constant of integration $C_{(0,1)}$ may then be determined from value-matching at β_1^*. This yields the following exact solution for the sole non-innovator's expected payoff over $(\beta_1^*, 1)$:

$$V_{(0,1)}^*(\beta) = H_{(0,1)}(\beta) + \frac{(1-\beta)\Phi(\beta)^{\frac{\rho}{\lambda}}}{(1-\beta_1^*)\Phi(\beta_1^*)^{\frac{\rho}{\lambda}}} \left(V_{(0,0)} - H_{(0,1)}(\beta_1^*)\right). \tag{9}$$

Depending on whether $V_{(0,0)} - H_{(0,1)}(\beta_1^*)$ is positive or negative, the expected payoff $V_{(0,1)}^*(\beta)$ may be strictly convex or strictly concave. However, we always have $\partial H_{(0,1)}(\beta_1^*)/\partial\phi > 0$. Thus, the sole non-innovator assigns a negative value to the sole innovator's ability to exit if and only if ϕ is sufficiently large, namely above some threshold $\bar{\phi}_{(1,0)}(f)$. By definition, the sole non-innovator would be willing to pay a nonzero amount of money to induce the sole innovator to continue performing R&D past β_1^* whenever $\phi < \bar{\phi}_{(1,0)}(f)$. Therefore, as beliefs

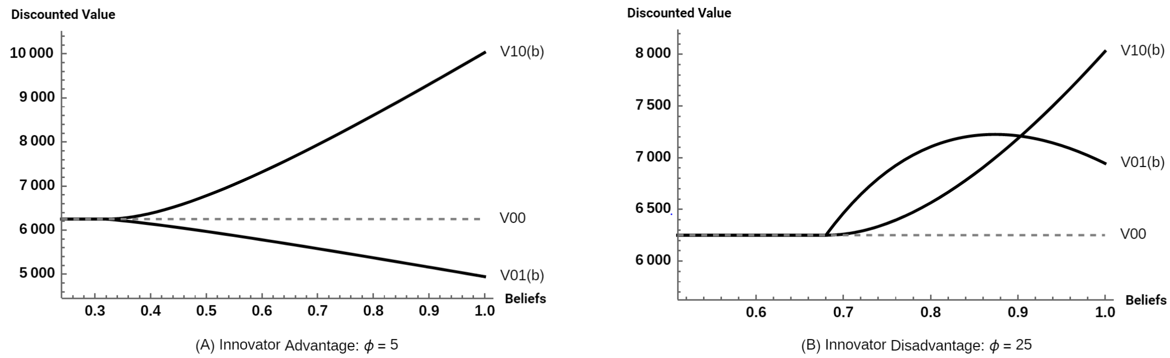

Figure 1: *Innovator Advantage and Disadvantage in the Individual R&D Subgame (Parameter values: $a = 100$, $c_A = 25$, $c_B = 0$, $f = 200$, $\rho = .1$,$\lambda = .1$)*

approach β_1^*, it is possible that the sole non-innovator may benefit more from additional R&D than the sole innovator itself.

Proposition 1. *There exists a threshold $\bar{\phi}_{(1,0)}(f)$ such that $V^*_{(1,0)}(\beta) > V^*_{(0,1)}(\beta)$ (i.e. an innovator advantage exists) for all $\beta \in (\beta_1^*, 1)$ whenever $\phi \leq \bar{\phi}_{(1,0)}(f)$. Otherwise, there exists a belief $\hat{\beta} > \beta_1^*$ such that $V^*_{(0,1)}(\beta) > V^*_{(1,0)}(\beta)$ (i.e. an innovator disadvantage exists) for all $\beta \in (\beta_1^*, \hat{\beta})$.*

Proof. See the Appendix. □

The logic behind the proof of Proposition 1 is that an innovator disadvantage exists in the individual R&D subgame if and only if ϕ is sufficiently large to guarantee that $V(0,1)'(\beta_1^*+) > 0$; that is, the sole non-innovator strictly benefits from additional R&D at β_1^*. The reason why this can occur follows from the fact that R&D investment, by producing marginal cost asymmetry, may increase industry profit through a joint-profit effect (i.e. $\Pi(0,1) + \Pi(1,0) > 2\Pi(0,0)$). Indeed, as the proof shows, an innovator disadvantage exists if and only if

$$\beta_1^* < \frac{\rho}{\lambda}\left(\frac{\Pi(0,1) - \Pi(0,0)}{\Pi(0,0) - \Pi^F}\right).$$

For $f = 0$, this condition is equivalent to

$$\frac{\Pi(0,0) - \Pi(1,0)}{\Pi^L - \Pi(0,0)} < \frac{\Pi(0,1) - \Pi(0,0)}{\Pi(0,0) - \Pi^F},$$

which holds whenever ϕ is large enough to produce a joint-profit effect. Note, however, that a joint-profit effect is not generally sufficient to produce an innovator disadvantage when $f > 0$ because ϕ must now be large enough to guarantee that the non-innovator benefits from additional R&D at $\beta_1^* = \beta_f^* + \beta_\phi^*$ rather than just β_ϕ^*.

Proposition 1 gives us a first glimpse at the novel strategic considerations that emerge when we consider the implications of R&D investment on product market competition. Specifically, depending on the cost structure of R&D (i.e. the share of explicit and implicit costs), firms may posses strategic incentives to underinvest in R&D to receive the short-run benefit of enhanced profitability while its rival continues to allocate resources toward R&D. This serves as a major departure from the standard patent race framework, which assumes $\phi = 0$. In this special case, there is no strategic benefit to early exit. Thus, if given the choice, each firm would strictly prefer to be the last firm to exit the patent race. As we show in the next subsection, this difference will have significant implications for the dynamics of joint R&D investment.

4.2 The Joint R&D Subgame

Now suppose that both firms remain in the patent race. To characterize Firm i's best-response in this subgame, take the action $e_j = 1$ as given and write out the HJB equation defining Firm i's expected payoff from remaining in the patent race:

$$\rho V_i(\beta) = \Pi(1,1) - f + 2\beta\lambda \left[\frac{V^L + V^F}{2} - V_i(\beta) \right] - 2\lambda(1-\beta)V_i'(\beta)V_i'(\beta). \tag{10}$$

As in the individual R&D subgame, Firm i's expected payoff can be written as the sum of (i) flow profit, (ii) the expected gains from additional R&D, and (iii) the effect of learning over time. For all $\phi > 0$, observe that $e_j = 1$ relaxes product market competition (i.e. $\Pi(1,1) > \Pi(1,0)$) but intensifies R&D competition, because Firm i's expected prize in the race is $(V^L + V^F)/2 < V^L$. The former effect is likely to reduce Firm i's incentive to perform R&D, while the latter effect is likely to increase it. Therefore, the overall effect of $e_j = 1$ on Firm i's incentive to perform R&D is unclear; we must explicitly compare stopping thresholds.

To derive Firm i's threshold for optimal stopping given $e_j = 1$, observe that Equation (10) may be rewritten as follows:

$$\rho V_i(\beta) = \Pi(0,1) + \beta\lambda \left[V^F - V_i(\beta) - (1-\beta)V_i'(\beta) \right] + \left[B(\beta, V_i) - f - C(1) \right], \tag{11}$$

where $C(1) \equiv \Pi(0,1) - \Pi(1,1)$ denotes the implicit cost of R&D given $e_j = 1$. Written in this way, we can easily see that Firm i strictly prefers the action $e_i = 1$ if and only if $B(\beta, V_i) \geq f + C(1)$; otherwise, $e_i = 0$ is optimal. Thus, Firm i's indifference condition for exit given $e_j = 1$ is

$$\rho V^*_{(0,1)}(\beta) = D_1(\beta) \equiv \Pi(0,1) + f + C(1) - \beta\lambda \left[V^L - V^F \right]. \tag{12}$$

Provided that firms are not too myopic, Equation (12) possesses a unique solution; denote it β_2^*. By definition, β_2^* is the belief at which Firm i would exit the patent race if Firm j were expected to choose $e_j = 1$ with probability one. For this reason, we call β_2^* the *joint stopping threshold*, as it completely summarizes Firm i's incentive to allow joint R&D to continue given $e_j = 1$.

To compare joint and individual R&D incentives, it suffices to compare β_2^* and β_1^*. Recall

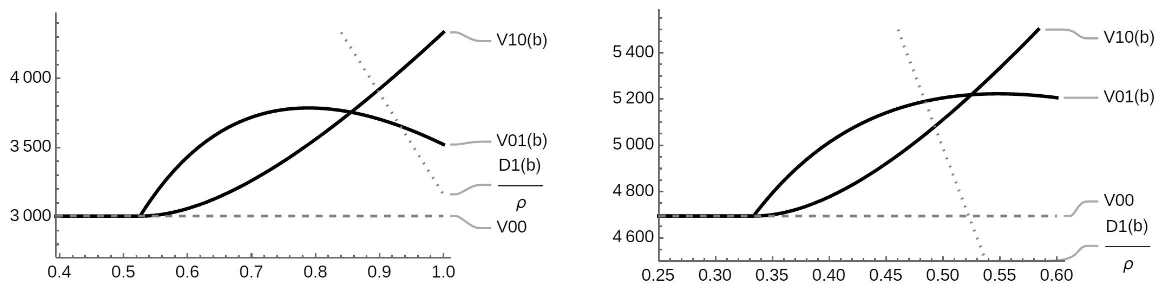

(A) Attrition (a=70,cA=18,ϕ=20,f=10,cB=0,ρ=.1,λ=.1) (B) Preemption (a=90,cA=25,ϕ=15,f=0,cB=0,ρ=.1,λ=.1)

Figure 2: *Attrition and Preemption in the Joint R&D Subgame*

that $\rho V^*_{(0,1)}(\beta) = \Pi(0,0)$ holds for all $\beta \leq \beta^*_1$, and observe that the right-hand side of Equation (12) is decreasing in β. Therefore, a necessary and sufficient condition for $\beta^*_2 \leq \beta^*_1$ is

$$\Pi(0,0) > \Pi(0,1) + f + C(1) - \beta^*_1\lambda \left[V^L - V^F\right] \iff RE(\beta^*_1) > CE, \tag{13}$$

where $RE(\beta) \equiv \beta\lambda[V^L - V^F] - f$, and $CE \equiv [C(1) - C(0)] + [\Pi(0,1) - \Pi(1,0)]$.

We refer to $RE(\beta)$ as the *racing effect*; it measures Firm i's additional incentive to remain in the patent race given $e_j = 1$ out of fear that business-stealing will occur. Racing effects like this one are a common feature winner-take-all patent races (see Reinganum, 1989) and can be responsible for generating strategic complementarity in firms' investments. Indeed, if all costs of R&D were explicit, then joint R&D incentives would exceed individual R&D incentives (i.e. e_i and e_j are strategic complements) whenever business-stealing concerns dominate the fixed cost of R&D for all $\beta \geq \beta^*_1$ – that is, whenever $RE(\beta^*_1) > 0$.

When both explicit and implicit costs of R&D are considered, another effect of $e_j = 1$ appears to alter Firm i's incentive to perform R&D. This effect is captured by the term CE, which we refer to as the *competition effect*. By definition, the competition effect measures the extent to which $e_j = 1$ reduces Firm i's incentive to perform R&D by raising both the implicit cost of R&D and the profitability of exit, as measured by the difference between innovator and non-innovator flow profits in the individual R&D subgame (i.e. $\Pi(0,1) - \Pi(1,0)$).

For all relevant parameter values, it is straightforward to verify that $RE(\beta^*_1) - CE$ is strictly decreasing in ϕ. Thus, we have $RE(\beta^*_1) \geq CE$ if and only if ϕ is sufficiently small, namely below some threshold value $\bar{\phi}_{(1,1)}(f)$. In the following proposition, we describe the structure of Markov-perfect equilibria in this case.

Proposition 2 (Simultaneous Exit MPE)**.** *Suppose that $\phi < \bar{\phi}_{(1,1)}(f)$ (i.e. $RE(\beta^*_1) > CE$). Then there is a unique Markov-perfect equilibrium. In this equilibrium, both firms perform R&D until beliefs reach the threshold β^*_1. Then simultaneous exit occurs.*

Proof. See the Appendix. □

Whenever the implicit cost of R&D is relatively small, Proposition 2 establishes that

Markov-perfect equilibrium R&D investment is practically non-strategic. Both firms perform R&D until beliefs reach the individual stopping threshold, and then both exit. The reason why both firms must perform R&D at least until beliefs reach β_1^* is straightforward. By definition, $RE(\beta_1^*) > CE$ implies that $\beta_2^* \leq \beta_1^*$ holds. Thus, the action $e = 1$ is a dominant strategy for both firms over the entire interval $(\beta_1^*, 1)$. So joint R&D must continue at least until beliefs reach β_1^*.

The reason why joint R&D cannot continue past β_1^*, however, is less obvious. To understand the logic behind the proof, note that if joint R&D were to continue past β_1^*, then each firm could induce simultaneous exit at any point in time by choosing the action $e = 0$. In such an equilibrium, the joint R&D subgame would therefore resemble a single-agent decision problem with a lower expected prize following the first success, namely $(V^L + V^F)/2 < V^L$, but a higher interim flow profit, $\Pi(1,1) - f$. As the proof of Proposition 2 shows, this negative prize-reduction effect of joint R&D always outweighs the benefit of relaxed product market competition. Therefore, if joint R&D were to continue until beliefs reached $\tilde{\beta} < \beta_1^*$, then each firm would have a strict incentive to induce simultaneous exit over a non-empty subset of beliefs between $\tilde{\beta}$ and β_1^*.

For larger values of ϕ, namely above $\bar{\phi}_{(1,1)}(f)$, we have $RE(\beta_1^*) < CE$. Thus, joint R&D incentives are weaker than individual R&D incentives, and sequential exit must occur in any Markov-perfect equilibrium. However, in sharp contrast to Proposition 2, the resulting dynamics of R&D investment now become highly strategic, because now each firm must anticipate that sequential exit will occur, and whether it wishes to be the first or last firm to exit the patent race.

Proposition 3 (Sequential Exit MPE)**.** *Suppose that $\phi > \bar{\phi}_{(1,1)}(f)$ (i.e. $RE(\beta_1^*) < CE$).*

(a) (Attrition Equilibrium.) If $\hat{\beta} < \beta_2^*$, then there is a unique symmetric Markov-perfect equilibrium. In this equilibrium, both firms perform R&D until beliefs reach β_2^*, and then randomization occurs. Specifically, each firm begins to exit with intensity

$$\mu^*(\beta) = \frac{f + C(1) - B(\beta, V^*_{(0,1)})}{V^*_{(1,0)}(\beta) - V^*_{(0,1)}(\beta)}$$

so long as neither firm has exited the patent race. With probability one, exit occurs before beliefs reach $\hat{\beta}$.

(b) (Preemption Equilibrium.) If $\hat{\beta} \geq \beta_2^*$, then there is a unique Markov-perfect equilibrium. In this equilibrium, both firms perform R&D until beliefs reach $\hat{\beta}$. Then with probability one-half Firm 1 exits the patent race; Firm 2 exits with complementary probability.

Proof. See the Appendix. □

The proof of Proposition 3 relies on a single powerful observation: for each firm, the $e = 1$ is a strictly dominant action if and only if $\beta > \max\{\beta_2^*, \hat{\beta}\}$. Therefore, joint R&D must continue at least until beliefs reach $\max\{\beta_2^*, \hat{\beta}\}$. Then firms face a coordination problem, the exact nature of which depends on the sign of $\beta_2^* - \hat{\beta}$.

If $\hat{\beta} < \beta_2^*$, then each firm would stay in the patent race if its rival were expected to exit, but not if its rival were expected to stay. Because $V_{(1,0)}^*(\beta_2^*) > V_{(0,1)}^*(\beta_2^*)$, the resulting continuation game is equivalent to a continuous-time War of Attrition, as defined by Hendricks et al. (1988). As such, there are three possible equilibria: two asymmetric equilibria in which one firm exits at β_2^* and one symmetric equilibrium in which both firms randomize. In the symmetric equilibrium, the rate at which firms exit intuitively depends on the difference between the cost and benefits of R&D, namely $f + C(1) - B(\beta, V_{(0,1)}^*)$, and the difference $V_{(1,0)}^*(\beta) - V_{(0,1)}^*(\beta)$, which measures the size of the innovator advantage at β. If success does not occur, then $f + C(1) - B(\beta, V_{(0,1)}^*)$ will increase over time, while $V_{(1,0)}^*(\beta) - V_{(0,1)}^*(\beta)$ will decrease. Therefore, exit occurs with higher intensity over time as joint R&D continues. Moreover, as β approaches $\hat{\beta}$, we have $\lim_{\beta \to \hat{\beta}} V_{(1,0)}^*(\beta) - V_{(0,1)}^*(\beta) = 0$. Thus, with probability one, exit will occur in $(\hat{\beta}, \beta_1^*)$

If $\hat{\beta} \geq \beta_2^*$, then the dynamics of joint R&D change dramatically. In this case, each firm strictly prefers to be the first firm to exit as beliefs cross $\max\{\beta_2^*, \hat{\beta}\}$. The resulting continuation game is not a War of Attrition, but rather a preemption game (c.f Fudenberg and Tirole, 1985). Indeed, if Firm j were expected to remain in the patent race until the belief $\tilde{\beta} < \hat{\beta}$, then Firm i would have a strict incentive to exit the patent race at beliefs slightly above $\tilde{\beta}$ because $V_{(0,1)}^*(\tilde{\beta} + \epsilon) > V_{(1,0)}^*(\tilde{\beta} + \epsilon)$ holds for all $\epsilon > 0$ sufficiently small. By symmetry, however, Firm j shares identical preemption motivations. Thus, exit must occur so long as $V_{(0,1)}^*(\beta) > V_{(1,0)}^*(\beta)$, which leaves only one possibility, namely that one firm exits at $\hat{\beta}$, where the innovator and non-innovator receive the same expected payoffs, namely $V_{(1,0)}^*(\hat{\beta}) = V_{(0,1)}^*(\hat{\beta})$.

4.2.1 Sufficient Conditions and Comparative Statics

Proposition 3 characterizes the set of sequential exit Markov-perfect equilibria solely in terms of $\hat{\beta}$ and β_2^*, two endogenous variables. We make this choice because the requirements for each type of sequential exit equilibrium to exist in terms of market fundamentals, namely ϕ and f, are surprisingly complex. To illustrate this complexity, consider what is required for the preemption equilibrium to exist. The first obvious requirement is that an innovator disadvantage must exist above β_1^*. Consequently, ϕ cannot be too close to zero. However, ϕ cannot be too large either, because the sole innovator must remain in the patent race for a sufficiently long time; otherwise, the benefits of early exit will not have enough time to accrue. Thus, the preemption equilibrium will exist only if ϕ is moderately large.

Figure 3 depicts the structure of Markov-perfect equilibria in (ϕ, f)-space for a representative selection of model parameters. In line with the preceding discussion, the preemption equilibrium exists only for moderate values of ϕ. If $f = 0$, then preemption occurs only if ϕ is not too large. As f increases, observe that equilibrium incentives always begin to favor attrition, and then simultaneous exit. The set of Markov-perfect equilibria may thus be viewed as monotonic with

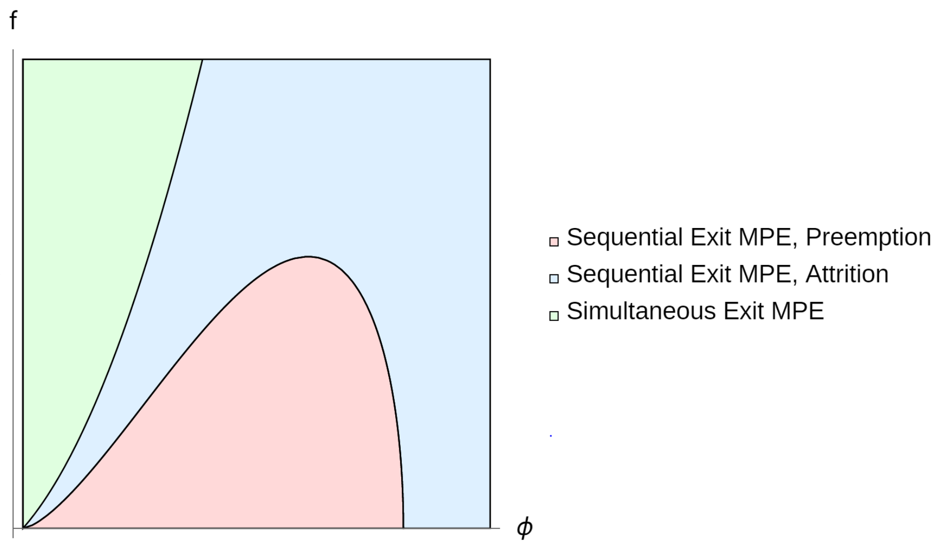

Figure 3: *Existence Regions for Simultaneous and Sequential Exit MPE in the Joint R&D Subgame*

respect to changes in f. The same cannot be said, however, for changes in ϕ. Indeed, with respect to an increase in ϕ, equilibrium incentives may first favor simultaneous exit, then favor attrition, then favor preemption, and then favor attrition once again. This surprising feature of equilibrium incentives is driven by the combination of positive and negative effects of an increase in ϕ, namely $\partial\Pi(0,1)/\partial\phi > 0$ (profitability of exit effect) and $\partial\beta_1^*/\partial\phi > 0$ (length of individual R&D effect).

In general, Markov-perfect equilibrium incentives are non-monotonic with respect to changes in any parameter that affects the profitability of exit and the length of individual R&D in opposite directions, such as the value of successful innovation.

Example 1. Fix $a = 100$, $c_A = 25$, $\phi = 20$, $f = 0$, $\rho = .2$, $\lambda = .25$. Suppose the marginal cost of the new technology is $c_B = 9$. Then we have $\beta_1^* = 0.604$, $\hat{\beta} = 0.855$, and $\beta_2^* = 0.889$. Therefore, joint R&D is characterized by attrition. Now suppose that c_B falls to 0. Then $\beta_1^* = 0.352$, $\hat{\beta} = 0.563$, $\beta_2^* = 0.535$. In this case, both β_2^* and $\hat{\beta}$ have decreased, but strategic incentives in equilibrium are reversed: joint R&D is now characterized by preemption.

5 Welfare

Let us now explore the welfare consequences of Propositions 2 and 3 by deriving the socially optimal R&D policy. We say that an R&D policy is socially optimal whenever it maximizes expected discounted total surplus, $TS = PS + CS$, net the flow cost of R&D. In our Cournot specification, flow total surplus (ignoring the fixed cost of R&D) is given by

$$\begin{cases} TS(e_1, e_j) = \Pi(e_i, e_j) + \Pi(e_j, e_i) + \frac{1}{2}Q(e_i, e_j)^2 & \text{before success occurs,} \\ TS^+ = \Pi^L + \Pi^F + \frac{1}{2}(Q^+)^2 & \text{after success occurs,} \end{cases} \tag{14}$$

where $Q(e_i, e_j) \equiv \frac{1}{3}(a - c_A - \phi(e_1 + e_2))$ and $Q^+ \equiv \frac{1}{3}(a - c_B - c_A)$ denote flow industry output before and after success occurs, respectively.

To characterize the dynamics of the optimal R&D policy, we proceed in a similar manner to the analysis of the strategic problem. Let $V_S(\beta)$ denote the social planner's value in the optimal R&D policy. Then the indifference condition for optimal stopping by Firm i, given e_j, is

$$\rho V_S(\beta) = TS(0, e_j) + e_j C_S(e_j), \tag{15}$$

where $C_S(e_j) \equiv TS(0, e_j) - TS(1, e_j)$ denotes the implicit social cost of R&D.

The right-hand side of Equation (15) measures the planner's incentive to have Firm i perform R&D given e_j. Clearly, the planner's joint R&D incentive exceeds its individual R&D incentive if and only if

$$TS(0, 1) + C_S(1) < TS(0, 0) \iff 2TS(0, 1) > TS(0, 0) + TS(1, 1).$$

By symmetry, note that $TS(0, 1) = TS(1, 0)$. Thus, social incentives for joint R&D are weaker than individual incentives whenever $TS(\cdot, \cdot)$ is strictly submodular. This is true whenever $\phi > 0$. From this, we conclude that the socially optimal R&D policy takes the following form:

Proposition 4 (Socially Optimal R&D)**.** *For all $\phi > 0$, sequential exit occurs in the socially optimal R&D policy. Furthermore, the last firm to exit always does so when beliefs reach the threshold*

$$\beta_1^S = \frac{\rho}{\lambda}\left(\frac{f + C_S(0)}{TS^+ - TS(0, 0)}\right).$$

If $\phi = 0$, then the socially optimal R&D policy recommends simultaneous exit at β_1^S.

Proof. See the Appendix. □

In contrast to the strategic problem, the socially optimal R&D features sequential exit for all $\phi > 0$. This is because the social value of successful innovation does not depend on which firm succeeds. Therefore, the social racing effect vanishes from the planner's indifference condition, and all that matters is the social competition effect, $CE_S \equiv TS(0, 1) + C_S(1) - TS(0, 0) = C_S(1) - C_S(0) \geq 0$. For all $\phi > 0$, the social planner therefore views (e_i, e_j) as strict substitutes. Consequently, the socially optimal R&D policy will recommend that individual R&D be performed over a non-empty interval of intermediate beliefs (β_1^S, β_2^S) because this allows the planner to optimally balance the trade-off between innovation and current productive efficiency.

In our setting, equilibrium R&D investment may be socially inefficient along two dimensions. The first dimension is the total amount of R&D performed before both firms exit the patent race. This is measured by the comparison between β_1^S and β_1^*. The second dimension is the amount of R&D that is performed jointly (i.e. the intensity of R&D). This is measured by the difference between β_2^S and $\max\{\beta_2^*, \hat{\beta}\}$. In the following proposition, we characterize the extent to which equilibrium R&D is inefficient along each dimension.

Proposition 5 (Equilibrium Inefficiency)**.** *Suppose that $\beta_0 > \beta_1^S$. Then:*

(i) For all $\phi \geq 0$, the total amount of R&D performed in any Markov-perfect equilibrium is, on average, less than the socially desirable amount.

(ii) For all $\phi > 0$, the intensity of R&D performed in any Markov-perfect equilibrium is, on average, higher than the socially desirable intensity whenever firms are sufficiently patient.

Proof. See the Appendix. □

The sources of equilibrium inefficiency are two-fold. First, there is an under-appropriability effect that reduces individual R&D incentives to an inefficiently low level because firms do not fully internalize the social benefits of successful innovation. Second, there is distortion in joint R&D incentives caused by inefficient private incentives for either attrition or preemption. Consequently, intensity of equilibrium R&D may be higher or lower than is socially desirable. When firms are sufficiently patient, however, we find that equilibrium R&D is always, on average, more intense than the socially optimum. Interestingly, this occurs regardless of whether equilibrium R&D features simultaneous, attrition, or even preemption.

To understand why this small-discounting result holds, consider how a change in ρ affects the relative magnitudes of CE and $RE(\beta)$. By definition, CE depends only on current flow profit. Thus, it is invariant to changes in ρ. The racing effect, however, depends on the difference $V^L - V^F$, which increases without bound as $\rho \to 0$. Therefore, the dominant consideration that shapes joint R&D incentives as firms become sufficiently patient is the racing effect. Formally, this implies $\lim_{\rho\to 0} \max\{\beta_2^*, \hat{\beta}\} = 0$. By contrast, the social competition always dominates the social racing effect, and each is invariant to the scale of ρ. Therefore, the planner's choice of β_2^S is bounded away from zero for any level of patience, thus establishing our desired result.

6 Market Structure and R&D Incentives

Our results thus far have illustrated the significant impact of product market considerations (as measured by ϕ) on equilibrium R&D investment and social welfare. In this section, we explore the policy implications of our framework by considering the effect of a change in market structure – specifically, a merger to monopoly – on the positive and normative aspects of equilibrium R&D investment.

To isolate the impact of market structure on R&D incentives, let us first consider the simplest possible model of a merger: a pure reduction in the number of firms from two (duopoly) to one (monopoly). We assume that the monopoly firm, labeled M, possesses identical capabilities to that of a single duopoly firm. So all changes in equilibrium R&D investment are, by definition, due to the elimination of competition in both the product market and in the patent race.

With this setup, all that changes between the monopoly problem and the individual R&D subgame is flow profit before and after successful innovation. Before success occurs, Firm M receives flow profit $\Pi_M(e) = (a - c_A - \phi e)^2/(4b)$. After success occurs, Firm M receives $\Pi_M^+ = (a - c_B)^2/(4b)$ in all future periods. Letting $V_M^+ = \Pi_M^+/\rho$ denote the value of successful innovation, the HJB equation defining Firm M's value in the patent race is

$$\rho V_M(\beta) = \Pi_M(1) - f + \beta\lambda \left[V_M^+ - V_M(\beta)\right] - \lambda\beta(1-\beta)V_M'(\beta). \tag{16}$$

The interpretation of Equation (16) is analogous to Equation (3) from the individual R&D subgame. By advancing similar arguments, we find that Firm M will continue to perform R&D until beliefs reach the threshold

$$\beta^M = \frac{\rho}{\lambda}\left(\frac{f + C_M}{\Pi_M^+ - \Pi_M(0)}\right), \tag{17}$$

where $C_M \equiv \Pi_M(0) - \Pi_M(1)$ denotes the implicit cost of monopoly R&D.

As before, the monopoly stopping threshold can be interpreted as having an explicit cost component, β_f^M, and an implicit cost component, β_ϕ^M. Observe that $\beta_f^M < \beta_f^*$ holds if and only if $\Pi_M^+ - \Pi_M(0) > \Pi^L - \Pi(0,0)$. In other words, if all costs of R&D were explicit, then we recover the familiar result that monopoly R&D incentives exceed duopoly R&D incentives whenever the returns to successful innovation are higher under monopoly.

However, when we account for both the explicit and implicit costs of R&D, this condition is no longer sufficient to generate enhanced monopoly innovation incentives. To demonstrate this, observe that β_ϕ^M simplifies to become

$$\beta_\phi^M = \frac{\rho}{\lambda}\left(\frac{\phi}{c_A - c_B}\right)\left(\frac{a - c_A - \phi + (a - c_A)}{a - c_B + (a - c_A)}\right), \tag{18}$$

which exceeds β_ϕ^* for all relevant parameter values. Compared to a duopoly market structure, enhanced market power allows a monopolist to enjoy greater profitability whether or not it successfully innovates. Therefore, the difference between the profit that Firm M receives while performing R&D and the profit it could receive if it were to exit, namely $C_M = \Pi_M(0) - \Pi_M(1)$ increases. We call this phenomenon the opportunity cost effect. After accounting for this effect, we see that monopoly R&D incentives exceed duopoly incentives only if the increase in returns to successful innovation are large enough to offset the negative impact of the opportunity cost effect on R&D incentives. Specifically, we have

$$\beta^M < \beta_1^* \iff \Pi_M^+ - \Pi_M(0) > \left(\Pi^L - \Pi(0,0)\right)\left(1 + \frac{\beta_\phi^M - \beta_\phi^*}{\beta_f^M}\right).$$

The following proposition illustrates the extent to which the opportunity cost effect may change the welfare effects of a change in market structure from monopoly to duopoly.

Proposition 6 (Monopoly vs Duopoly Welfare)**.** *Suppose that $TS_M^+ > TS(0,0)$. Then:*

(i) If all costs of R&D are explicit (i.e. $f > 0 = \phi$), then long-run average total welfare is highest under a monopoly market structure if and only if $k \equiv (\Pi_M^+ - \Pi_M(0))/(\Pi^L - \Pi(0,0))$ is sufficiently larger than one.

(ii) If all costs of R&D are implicit (i.e. $f = 0 < \phi$), then long-run average total welfare is always lower under a monopoly market structure compared to duopoly.

Proof. See the Appendix. □

A common defense of mergers is that cost or R&D synergies will generate benefits that will be passed along to consumers in the form of lower prices or increased innovation incentives. To explore the effects of each type of synergy in the context of our framework, let us now suppose that Firm M achieves greater production and research capability compared to a pre-merger duopolist. Specifically, let Firm M's marginal cost before success be $c^M(e) = (1 - s_C)c(e)$, where $s_C \in [0, 1]$ parameterizes the size of cost synergies following the merger. Likewise, let Firm M's marginal cost after success occurs be $c_B^M = (1 - s_C)c_B$.

In addition to production-related cost synergies, we allow the possibility that Firm M can perform R&D more cheaply than before the merger. To account for such R&D synergies, let us suppose that the explicit (fixed) cost of R&D investment now becomes $f^M = (1 - s_{RD})f$, where $s_{RD} \in [0, 1]$ parameterizes the size of R&D cost synergies.

In this positive synergies case, the monopoly stopping threshold becomes

$$\beta^M = \frac{\rho}{\lambda}\left(\frac{(1 - s_{RD})f + (1 - s_C)C_M}{\Pi_M^+ - \Pi_M(0)}\right). \tag{19}$$

Observe that R&D cost synergies unambiguously increase R&D incentives, i.e. $\partial\beta^M/\partial s_{RD} < 0$. So let us henceforth assume $s_{RD} = 0$ to focus purely on production synergies, which have an ambiguous effect on monopoly R&D incentives.

In terms of the explicit cost component of β^M, we have (i) $\partial\beta_f^M/\partial s_C < 0$ if and only if $s_C < (a - c_A - c_B)/(c_A + c_B)$. In other words, greater production-related cost synergies increase the returns to successful innovation only if s_C is not too large; otherwise, a non-innovating monopolist can be nearly as efficient without successful innovation than it can be if it were to succeed. Indeed, if $s_C = 1$, then $\beta_f^M \to 0$ because success in the patent race is no longer necessary to achieve the lowest possible cost: $c_A^M = c_B^M = 0$.

In terms of the implicit component of β^M, we have $\partial\beta_\phi^M/s_C > 0$ for all relevant parameter values. This comparative static highlights a surprising feature of the interaction between cost synergies and the implicit cost of R&D, namely that greater cost synergies, by enhancing current profitability, increase the size of the opportunity cost effect. Consequently, it is not valid to always consider cost synergies as being conducive to enhanced innovation incentives, even if the synergies in question would increase the returns to successful innovation. In terms of merger control, it follows that a merger with large anticipate product-related cost synergies may, in fact, require greater scrutiny than a merger with smaller synergies, depending on the relative cost structure of R&D.

To describe these merger implications formally, consider once again the effect of a merger on long-run average welfare. The total effect of the merger is the difference between long-run average welfare under monopoly, labeled $\nu_S^M(\beta_0 \mid \beta^M)$, and long-run average welfare under duopoly, labeled $\nu_S^D(\beta_0 \mid \beta_1^*)$. By adding and subtracting the term $\nu_S^M(\beta_0 \mid \beta_1^*)$, i.e. long-run average welfare under monopoly if the monopolist were to exit the patent race at β_1^*, we can write the total welfare effect of a merger as

$$\nu_S^M(\beta_0 \mid \beta^M) - \nu_S^D(\beta_0 \mid \beta_1^*) = \underbrace{\left[\nu_S^M(\beta_0 \mid \beta_1^*) - \nu_S^D(\beta_0 \mid \beta_1^*)\right]}_{\text{Market Structure Effect } (MSE)} + \underbrace{\left[\nu_S^M(\beta_0 \mid \beta^M) - \nu_S^M(\beta_0 \mid \beta_1^*)\right]}_{\text{R\&D Incentive Effect } (RDE)} \tag{20}$$

The market structure effect describes the long-run welfare effect of a merger based solely on how the merger affects market structure, holding innovation incentives constant. Clearly, this effect is always increasing in s_C. The R&D incentive effect, by comparison, measures the welfare impact of changing R&D incentives, holding market structure fixed. This effect may be positive or negative, depending on the comparison between β^M and β_1^*. In particular, we have $RDE > 0$ if and only if $\beta^M < \beta_1^*$. Using this decomposition, we may now summarize the implications of the opportunity cost effect on the welfare effects of a merger with synergies.

Proposition 7. *Suppose that s_C is such that $MSE = 0$; that is, the merger is welfare-neutral from the standpoint of market structure.*

(i) If all costs of R&D are explicit, then a merger with cost synergies s_C improves long-run average welfare if and only if $\Pi_M^+ - \Pi_M(0) > \Pi^L - \Pi(0,0)$.

(ii) If all costs of R&D are implicit, then a merger with cost synergies s_C always reduces long-run average welfare.

Proof. See the Appendix. □

Depending on the cost structure of R&D and the implications of firm investment on product market competitiveness, Propositions 6 and 7 highlight that welfare-maximizing merger enforcement policy will necessarily apply stricter standards toward proposed mergers in industries characterized by relatively large implicit costs of R&D are relatively large compared to explicit costs. This is because the opportunity cost effect of monopolization is most likely to negatively affect innovation incentives and, as a result, welfare. Consequently, arguments based solely on enhanced appropriability are not generally sufficient to justify a merger as welfare improving. As Proposition 7 indicates, a merger may generate benign (or even positive) market structure benefits may nevertheless be blocked for the reason that it will harm innovation incentives. Our results therefore provide theoretical support for an innovation theory of harm, which has become a central consideration of recent merger cases (see Denicolo and Polo, 2018; Jullien and Lefouili, 2018).

7 Concluding Remarks

To successful innovate, it is often required that firms allocate scarce resources to *"support the search for new knowledge and prospective opportunities instead of leveraging currently available knowledge to address immediate needs"* (Lavie et al., 2010). Economic models should therefore not only consider the explicit costs of R&D, but also the opportunity cost of foregone exploitation. In this paper, I consider the strategic implications of this opportunity cost for firm investment the context of dynamic R&D competition – a patent race. My model shows that equilibrium R&D incentives can significantly depart from the predictions of existing models in terms of investment dynamics, social welfare, and the fundamental relationship between market structure and the incentives to innovate.

Extensions. My model can be extended in several directions. First, the assumption of Cournot competition is not essential for any of our equilibrium results. Indeed, our framework may easily accommodate other modes of product market competition, including Bertrand competition, Hotelling competition, or the Singh and Vives (1984) model of differentiated price or quantity competition. All that my results require is that products are strict substitutes and that R&D investment raises firms' marginal costs in such a way that reduced-form profits are decreasing in own R&D investment and increasing in rival R&D investment. Likewise, my base assumption that patent protection is infinite and perfect is not essential. My results extend to the case of finite patent life or imperfect protection without modification whenever the patent's duration is sufficiently long and imitation is sufficiently difficult. If patent duration is short or imitation is easy, however, then the gains to successful innovation may be destroyed. In this case, the competition effect will dominate the racing effect for all relevant parameter values, and so the unique equilibrium outcome in this case will involve preemption (if entry into the patent race at all).

Because I assume the winner of the patent race obtains a non-drastic innovation, it possible that firms could mutually benefit from the sale of a technology license. If gains from trade exist, then it is easy to see that licensing will increase the gains to successful innovation and, therefore, increase the total amount of R&D performed in equilibrium. However, licensing may also reduce the strength of the racing effect, depending on how the gains from trade are distributed between licensor and licensee. If the licensor captures a most of the gains from trade, then licensing will result in "faster" innovation by making the attrition or simultaneous exit more likely to occur. However, if the licensee captures a sufficiently large amount of the gains from trade, it is possible that licensing can result in "slower" innovation by decreasing the strength of the racing effect and, thus, making preemption more likely to occur. The effects of licensing on total R&D incentives are, thus, unambiguously positive, but its effects on the speed of innovation will generally depend on the terms of the anticipated licensing agreement.

Directions for Future Work. Despite the flexibility of my model, there are several obvious directions for future work. First, I consider the case of a single innovation. The implicit assumption here is that the expected rate of future innovation is zero. One might wish to extend to allow subsequent innovation in order to study the dynamics of experimentation incentives over multiple generations of new technology. Such an extension is beyond the scope of the present paper, but I acknowledge that it may produce interesting insights related to debate surrounding the persistence of monopoly (see Gilbert and Newbery, 1982; Reinganum, 1983). In particular, one may wish to study whether or not increasing dominance is an inevitable feature of competitive experimentation.

Second, additional insights may be gleaned from an extension of my duopoly model to N-firm oligopoly because the strategic incentives to invest in R&D will naturally evolve as firms drop out over time. Intuitively, the product market is be most contestable when many firms participate in the patent race. Therefore, one might conjecture that preemption incentives are strongest at the start of the patent race and then become weaker as more firms drop out, ultimately causing strategic incentives to favor attrition. Empirically, we should therefore expect to see a U-shaped distribution in firms' exit times from the patent race (i.e. many firms exiting at the beginning and end of the patent race), which complements existing work on the relationship between the innovation and the intensity of product market competition (Aghion et al., 2005). I leave this and the previous question to future work.

References

Aghion, P., N. Bloom, R. Blundell, R. Griffith, and P. Howitt (2005): “Competition and Innovation: an Inverted-U Relationship*,” *The Quarterly Journal of Economics*, 120, 701–728.

Aghion, P., C. Harris, P. Howitt, and J. Vickers (2001): “Competition, Imitation and Growth with Step-by-Step Innovation,” *The Review of Economic Studies*, 68, 467–492.

Akcigit, U., D. Hanley, and N. Serrano-Velarde (2021): “Back to Basics: Basic Research Spillovers, Innovation Policy, and Growth,” *The Review of Economic Studies*, 88, 1–43.

Akcigit, U. and Q. Liu (2016): “The Role of Information in Innovation and Competition,” *Journal of the European Economic Association*, 14, 828–870.

Arrow, K. (1962): “Economic Welfare and the Allocation of Resources for Invention,” in *The Rate and Direction of Inventive Activity: Economic and Social Factors*, Princeton University Press, 609–626.

Awaya, Y. and V. Krishna (2021): "Startups and Upstarts: Disadvantageous Information in R&D," *Journal of Political Economy*, 129, 534–569.

Becker, G. S. (1965): "A Theory of the Allocation of Time," *The Economic Journal*, 75, 493–517.

Besanko, D. and J. Wu (2013): "The Impact of Market Structure and Learning on the Tradeoff between R&D Competition and Cooperation," *The Journal of Industrial Economics*, 61, 166–201.

Brown, J., G. Martinsson, and B. Petersen (2012): "Do financing constraints matter for R&D?" *European Economic Review*, 56, 1512–1529.

Caggese, A. (2019): "Financing Constraints, Radical versus Incremental Innovation, and Aggregate Productivity," *American Economic Journal: Macroeconomics*, 11, 275–309.

Choi, J. P. (1991): "Dynamic R&D Competition under "Hazard Rate" Uncertainty," *The RAND Journal of Economics*, 22, 596–610.

Cohen, W. M. (2010): "Fifty Years of Empirical Studies of Innovative Activity and Performance," in *Handbook of the Economics of Innovation*, Elsevier, vol. 1, 129–213.

Denicolo, V. and M. Polo (2018): "The Innovation Theory of Harm: An Appraisal," SSRN Scholarly Paper ID 3146731, Social Science Research Network, Rochester, NY.

Dosi, G. (1988): "Sources, Procedures, and Microeconomic Effects of Innovation," *Journal of Economic Literature*, 26, 1120–1171.

Dutta, P. K., S. Lach, and A. Rustichini (1995): "Better Late than Early: Vertical Differentiation in the Adoption of a New Technology," *Journal of Economics & Management Strategy*, 4, 563–589.

Dutta, P. K. and A. Rustichini (1993): "A Theory of Stopping Time Games with Applications to Product Innovation and Asset Sales," *Economic Theory*, 3, 743–763.

Fudenberg, D. and J. Tirole (1985): "Preemption and Rent Equalization in the Adoption of New Technology," *The Review of Economic Studies*, 52, 383–401.

Giebel, M. and K. Kraft (2019): "External Financing Constraints and Firm Innovation," *The Journal of Industrial Economics*, 67, 91–126.

Gifford, S. (1992): "Allocation of entrepreneurial attention," *Journal of Economic Behavior & Organization*, 19, 265–284.

Gilbert, R. (2006): "Looking for Mr. Schumpeter: Where Are We in the Competition–Innovation Debate?" *Innovation Policy and the Economy*, 6, 159–215.

Gilbert, R. J. and D. M. G. Newbery (1982): "Preemptive Patenting and the Persistence of Monopoly," *The American Economic Review*, 72, 514–526.

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

He, Z.-L. and P.-K. Wong (2004): "Exploration vs. Exploitation: An Empirical Test of the Ambidexterity Hypothesis," *Organization Science*, 15, 481–494.

Hendricks, K., A. Weiss, and C. Wilson (1988): "The War of Attrition in Continuous Time with Complete Information," *International Economic Review*, 29, 663–680.

Holmes, T. J., D. K. Levine, and J. A. Schmitz (2012): "Monopoly and the Incentive to Innovate When Adoption Involves Switchover Disruptions," *American Economic Journal: Microeconomics*, 4, 1–33.

Hopenhayn, H. and F. Squintani (2011): "Preemption Games with Private Information," *The Review of Economic Studies*, 78, 667–692.

Jullien, B. and Y. Lefouili (2018): "Horizontal Mergers and Innovation," *Journal of Competition Law & Economics*, 14, 364–392.

Keller, G., S. Rady, and M. Cripps (2005): "Strategic Experimentation with Exponential Bandits," *Econometrica*, 73, 39–68.

Lavie, D., U. Stettner, and M. L. Tushman (2010): "Exploration and Exploitation Within and Across Organizations," *Academy of Management Annals*, 4, 109–155.

Malueg, D. A. and S. O. Tsutsui (1997): "Dynamic R&D Competition with Learning," *The RAND Journal of Economics*, 28, 751–772.

Mansfield, E., Yale University, and Cowles Foundation for Research in Economics (1968): *Industrial research and technological innovation; an econometric analysis.*, New York: Norton.

March, J. G. (1991): "Exploration and Exploitation in Organizational Learning," *Organization Science*, 2, 71–87.

McDonald, R. and D. Siegel (1986): "The Value of Waiting to Invest," *The Quarterly Journal of Economics*, 101, 707–728.

Moscarini, G. and F. Squintani (2010): "Competitive experimentation with private information: The survivor's curse," *Journal of Economic Theory*, 145, 639–660.

Radner, R. and M. Rothschild (1975): "On the allocation of effort," *Journal of Economic Theory*, 10, 358–376.

Reinganum, J. F. (1982): "A Dynamic Game of R and D: Patent Protection and Competitive Behavior," *Econometrica*, 50, 671–688.

——— (1983): "Uncertain Innovation and the Persistence of Monopoly," *The American Economic Review*, 73, 741–748.

——— (1989): "The Timing of Innovation: Research, Development, and Diffusion," in *Handbook of Industrial Organization*, Elsevier, vol. 1, 849–908.

Schumpeter, J. (1939): *Business Cycles*, vol. 1, McGraw Hill New York.

Simon, L. K. and M. B. Stinchcombe (1989): "Extensive Form Games in Continuous Time: Pure Strategies," *Econometrica*, 57, 1171–1214.

Singh, N. and X. Vives (1984): "Price and Quantity Competition in a Differentiated Duopoly," *The RAND Journal of Economics*, 15, 546–554.

Strulovici, B. and M. Szydlowski (2015): "On the Smoothness of Value Functions and the Existence of Optimal Strategies in Diffusion Models," *Journal of Economic Theory*, 159, 1016–1055.

Taylor, C. R. (1995): "Digging for Golden Carrots: An Analysis of Research Tournaments," *The American Economic Review*, 85, 872–890.

Thomas, C. D. (2021): "Strategic Experimentation with Congestion," *American Economic Journal: Microeconomics*, 13, 1–82.

Tripsas, M. and G. Gavetti (2000): "Capabilities, Cognition, and Inertia: Evidence from Digital Imaging," *Strategic Management Journal*, 21, 1147–1161.

Vives, X. (2008): "Innovation and Competitive Pressure*," *The Journal of Industrial Economics*, 56, 419–469.

A Appendix

Proof of Proposition 1. Observe that both $H_{(0,1)}(\beta)$ and $(1-\beta)\Phi(\beta)^{\frac{\rho}{\lambda}}$ are decreasing over $(0,1)$. Therefore, $V^*_{(0,1)}(\beta)$ is decreasing over $(\beta_1^*, 1)$ whenever $C_{(0,1)} \geq 0$. Thus, $C_{(0,1)} < 0$ is necessary for an innovator disadvantage to exist. By our earlier observation, $(V^*_{(0,1)})'(\beta_1^*+) > 0$ is sufficient for $C_{(0,1)} < 0$. Therefore, an innovator disadvantage exists iff $(V^*_{(0,1)})'(\beta_1^*+) > 0$. Furthermore, because $V^*_{(0,1)}(\beta)$ is concave and $V^*_{(1,0)}(\beta)$ is convex over $(\beta_1^*, 1)$, the set of beliefs at which an innovator disadvantage exists must take the form of an interval $(\beta_1^*, \hat{\beta})$.

To identify a condition under $(V^*_{(0,1)})'(\beta_1^*+) > 0$ holds, take the limit of Equation (7) as beliefs approach β_1^*+ to conclude that a necessary and sufficient condition for $(V^*_{(0,1)})'(\beta_1^*+) > 0$ is

$$\Pi(0,0) < \Pi(0,1) + \beta_1^* \lambda \left[V^F - V_{(0,0)}\right] \iff \beta_1^* < \tilde{\beta}_{(1,0)} \equiv \frac{\rho}{\lambda}\left(\frac{\Pi(0,1) - \Pi(0,0)}{\Pi(0,0) - \Pi^F}\right).$$

The threshold $\tilde{\beta}_{(1,0)}$ denotes the highest belief at which the non-innovator would pay in order to prolong the individual R&D subgame. Thus, we conclude an intuitive fact: an innovator disadvantage exists iff the non-innovator prefers that exit not occur at β_1^*.

To relate the above necessary and sufficient condition to ϕ, observe that

$$\begin{aligned}\frac{\partial}{\partial \phi}\beta_1^* &= \frac{\rho}{\lambda}\left(\frac{1}{\Pi^L - \Pi(0,0)}\right)\frac{\partial}{\partial \phi}(\Pi(0,0) - \Pi(1,0)) \\ &< \frac{\rho}{\lambda}\left(\frac{1}{\Pi(0,0) - \Pi^F}\right)\frac{\partial}{\partial \phi}(\Pi(0,1) - \Pi(0,0)) = \frac{\partial}{\partial \phi}\tilde{\beta}_{(1,0)}\end{aligned}$$

always holds. Intuitively, an increase in ϕ raises the opportunity cost of R&D as well as the profitability of non-investment; however, the latter effect is always larger than the former. This is due to a joint profit effect, namely that an increase in ϕ raises $\Pi(0,1)$ by more than it lowers $\Pi(1,0)$. Therefore, $\tilde{\beta}_{(1,1)} - \beta_1^*$ is increasing in ϕ and is positive whenever ϕ is sufficiently large.

To complete the proof, simply note that $\tilde{\beta}_{(1,0)} = 0 \leq \beta_1^*$ holds whenever $\phi = 0$. Thus, there exists a threshold $\bar{\phi}_{(1,0)}(f)$ (> 0 if $f > 0$) such that $\beta_1^* < \tilde{\beta}_{(1,0)}$ holds iff $\phi > \bar{\phi}_{(1,0)}(f)$, as desired. □

Proof of Proposition 2. The proof consists of two steps. First, we establish that if simultaneous exit occurs in Markov-perfect equilibrium, then it must happen when beliefs reach β_1^*. Then we establish an equivalence between simultaneous exit in general and $CE < RE(\beta_1^*)$.

Suppose that simultaneous exit occurs at the belief $\tilde{\beta}$ in a particular Markov-perfect equilibrium. Then clearly we must ave $\tilde{\beta} \leq \beta_1^*$; otherwise, one firm would have a strict incentive to deviate by remaining in the patent race. Let $V_{(1,1)}(\beta)$ denote the (symmetric) equilibrium payoff in such an equilibrium. If it were the case that $\tilde{\beta} < \beta^*_{(1,0)}$, then exit by Firm i will induce immediate exit by Firm j for all $\beta \in (\tilde{\beta}, \beta_1^*)$. Thus, we must have $V_{(1,1)}(\beta) \geq V_{(0,0)}$ over this entire interval. By strict convexity, it follows that $V'_{(1,1)}(\tilde{\beta}+) \geq 0$ must hold.

To determine the possibility of $V'_{(1,1)}(\tilde{\beta}+) \geq 0$, take the limit of Equation (10) as beliefs approach $\tilde{\beta}+$ to conclude that a necessary and sufficient condition for $V'_{(1,1)}(\tilde{\beta}+) \geq 0$ is

$$\Pi(0,0) \leq \Pi(1,1) - f + \tilde{\beta}\lambda \left[V^L + V^F - 2V_{(0,0)}\right] \iff \tilde{\beta} \geq \frac{\rho}{\lambda}\left(\frac{\Pi(0,0) - \Pi(1,1) + f}{\Pi^L + \Pi^F - 2\Pi(0,0)}\right).$$

Therefore, to establish that $V'_{(1,1)}(\tilde{\beta}+) \geq 0$ is impossible, the above inequality implies that it suffices to show that $\tilde{\beta}_{(1,1)} \geq \beta_1^*$ always holds. This condition, in turn, is equivalent to

$$\frac{\Pi(0,0) - \Pi^F}{\Pi^L - \Pi(0,0)} \geq \frac{\Pi(1,1) - \Pi(1,0)}{\Pi(0,0) - \Pi(1,0) + f}.$$

This condition clearly holds whenever $\phi = 0$. For $\phi > 0$, an upper bound on the right-hand side of our desired condition is obtained by setting $f = 0$. In this case, it can be easily checked $(\Pi(1,1) - \Pi(1,0))/(\Pi(0,0) - \Pi(1,0))$ is decreasing in ϕ. Thus, our our desired condition holds for all $\phi > 0$. From this, we conclude that simultaneous exit at $\tilde{\beta} < \beta_1^*$ cannot occur in any Markov-perfect equilibrium, which leaves only the possibility that $\tilde{\beta} = \beta_1^*$.

Having this intermediate result, it now suffices to establish necessary and sufficient conditions for simultaneous exit at β_1^* to be an equilibrium. To do this, suppose that both firms perform R&D until beliefs reach $\beta^*_{(1,0)}$, and then both firms exit. Once again, let $V_{(1,1)}(\beta)$ denote the (convex) value function induced by these strategies. Then, at any belief $\beta > \beta^*_{(1,0)}$, Firm i can exit the patent race to receive $V^*_{(0,1)}(\beta)$. For simultaneous exit to occur at β_1^*, such exit must be unprofitable for all $\beta \in (\beta_1^*, 1)$. Thus, we require $V'_{(1,1)}(\beta_1^*+) \geq V'_{(0,1)}(\beta_1^*+)$. Using Equations (7) and (10), this necessary and sufficient condition is equivalent to

$$\frac{1}{2}\left(\Pi(1,1) - f + \beta_1^*\lambda\left[V^L + V^F - 2V_{(0,0)}\right]\right) \geq \Pi(0,1) - \Pi(0,0) + \beta_1^*\left[V^F - V_{(0,0)}\right]$$

which simplifies to $CE \leq RE(\beta^*_{(1,0)})$, as desired. □

Proof of Proposition 3. Suppose that $\phi > \bar{\phi}_{(1,1)}(f)$. Then we know that sequential exit must occur in any Markov-perfect equilibrium because $\beta_2^* > \beta_1^*$. To determine the timing of sequential exit, consider two cases.

Case 1: $\hat{\beta} < \beta_2^*$.

In this case, if Firm j were to exit at the belief $\beta^j < \beta_2^*$, then Firm i would exit at β_2^*. Thus, exit must occur with positive probability whenever $\beta < \beta_2^*$. In a symmetric equilibrium, it cannot be the case that both firms attempt to exit the patent race with positive probability for $\beta \in (\hat{\beta}, \beta_2^*)$ because then $V^*_{(1,0)}(\beta) > V^*_{(0,1)}(\beta)$ would imply that a profitable deviation toward staying in the patent race would exist in order to avoid a tie-break. Thus, continuous randomization must occur over this interval.

Let the intensity at which each firm exits be denoted $\mu_{(1,1)}(\beta)$. To determine a formula for

$\mu^*(\beta)$, write out the HJB equation defining each firm's value during randomization:

$$\tilde{V}_{(1,1)}(\beta) = \Pi(1,1) - f + \beta\lambda \left[V^L + V^F - 2\tilde{V}_{(1,1)}(\beta) - 2(1-\beta)\tilde{V}'_{(1,1)}(\beta) \right] + \mu^*(\beta) \left[V^*_{(1,0)}(\beta) - \tilde{V}_{(1,1)}(\beta) \right].$$

For all $\beta \in (\hat{\beta}, \beta^*_{(1,1)})$, each firm must be indifferent to the possibility of exit. Therefore, $\tilde{V}_{(1,1)}(\beta) = V^*_{(0,1)}(\beta)$ and $\tilde{V}'_{(1,1)}(\beta) = (V^*_{(0,1)})'(\beta)$ hold for all $\beta \in (\hat{\beta}, \beta^*_2)$. Using these boundary conditions, the above HJB equation simplifies to give the desired expression for $\mu^*(\beta)$.

To finish the proof of this case, that the denominator of $\mu^*(\beta)$ approaches zero as $\beta \to \hat{\beta}$, while the numerator remains positive. Thus, $\mu^*(\beta)$ tends to infinity as β approaches $\hat{\beta}$. This implies that exit from the patent race will, with probability one, occurs sometime before beliefs reach $\hat{\beta}$.

Case 2: $\hat{\beta} > \beta^*_{(1,1)}$.

In this case, if Firm j were to exit at the belief $\beta^j < \hat{\beta}$, then Firm i would have a strict preference for preemptive exit at the belief $\beta^j + \epsilon$ for sufficiently small $\epsilon > 0$ because $V^*_{(0,1)}(\beta^j) > V^*_{(1,0)}(\beta^j)$. However, in equilibrium, Firm j possesses the same incentives, which implies that both firms will attempt to exit the race at some belief that is no less than $\hat{\beta}$. For $\beta > \hat{\beta}$, there is a strict innovator advantage. Thus, if Firm j were to exit at such a belief, then Firm i would have a strict incentive to remain in the patent race. Therefore, the only possible belief at which exit occurs is $\hat{\beta}$ itself, where $V^*_{(1,0)}(\hat{\beta}) = V^*_{(0,1)}(\hat{\beta})$. Because both firms attempt to exit at the same instant of time, a tie-break will occur. Because $\hat{\beta} > \beta^*_1$, it follows that the loser of this tie-break (i.e. the second-mover) will revise its action to stay in the patent race. Thus, Firm 1 exits the patent race $\hat{\beta}$ with one-half probability, and Firm 2 exits with complementary probability, as desired. □

Proof of Proposition 4. Given $\beta \in (0,1)$ and $(e_i, e_j) \in \{0,1\}$, the HJB equation defining the social planner's value, denoted $V_S(\beta)$, is

$$\rho V_S(\beta) = TS(e_i, e_j) - f(e_i + e_j) + \beta\lambda(e_i + e_j) \left[V_S^+ - V_S(\beta) - (1-\beta)V'_S(\beta) \right], \tag{21}$$

where $V_S^+ \equiv TS^+/\rho$ denotes the social value of successful innovation.

To derive the indifference condition for optimal stopping given $e_j = 1$, note that Equation (21) can be re-written as

$$\rho V_S(\beta) = TS(0, e_j) + e_j(B(\beta, V_S) - f) + e_i(B(\beta, V_S) - f - C_S(e_j)), \tag{22}$$

where $B(\beta, V_S) \equiv \beta\lambda[V_S^+ - V_S(\beta) - (1-\beta)V'_S(\beta)]$, and $C_S(e_j) \equiv TS(0, e_j) - TS(1, e_j)$. The planner is indifferent between the actions $e_i \in \{0,1\}$ if and only if $B(\beta, V_S) = f + C_S(e_j)$, which is equivalent to the condition $\rho V_S(\beta) = TS(0, e_j) + e_j C_S(e_j)$.

By direct inspection, the planner's joint R&D incentive is lower than its individual R&D incentive if and only if $TS(0,1) + C_S(1) > TS(0,0)$, which simplifies to give $2TS(0,1) > TS(0,0) + TS(1,1)$. Consequently, the optimal R&D policy features sequential whenever

$TS(e_i, e_j)$ is strictly submodular. We know that both $\Pi(e_i, e_j)$ is strictly submodular whenever $\phi > 0$. Therefore, all we must check is whether $\partial^2 \frac{1}{2}(Q(e_i, e_j))^2 / \partial e_i \partial e_j \leq 0$ holds. Computing this derivative, we obtain

$$\frac{\partial^2}{\partial e_i \partial e_j} \frac{1}{2} Q(e_i, e_j)^2 = -\frac{\phi^2}{9} \leq 0.$$

Thus, sequential exit occurs whenever $\phi > 0$.

To determine the belief threshold at which the last firm exits in the optimal policy, we have two boundary conditions: (i) $V_S(\beta_1^S) = TS(0,0)$ (value matching), and (ii) $V_S'(\beta_1^S) = 0$ (smooth pasting). Inserting these conditions into Equation (21) yields the desired value for β_1^S:

$$TS(0,0) = TS(1,0) - f + \beta\lambda[V_S^+ - TS(0,0)/\rho] \iff \beta_1^S = \frac{\rho}{\lambda}\left(\frac{f + C_S(0)}{TS^+ - TS(0,0)}\right).$$

□

Proof of Proposition 5. To determine whether the total amount of R&D performed in equilibrium is, on average, higher or lower than the social optimum, it suffices to compare β_1^S to β_1^*. Using the decomposition $\beta_1^S = \beta_f^S + \beta_\phi^S$, where β_f^S and β_ϕ^S denote the explicit and implicit components of β_1^S, we can compare the implicit components as follows:

$$\beta_\phi^S = \beta^c \left(\frac{8(a - c_A - \phi) - 3\phi}{8(a - c_B) + 3(c_A - c_B)}\right) < \beta_\phi^*.$$

Furthermore, through direct inspection, we have $TS^+ - TS(0,0) > \Pi^L - \Pi(0,0)$ for all relevant parameter values. Therefore, we have $\beta_f^S < \beta_f^*$ as well, which allows us to unambiguously conclude that $\beta_1^S < \beta_1^*$ holds for all relevant parameter values.

To determine the efficiency of equilibrium R&D intensity for small discount rates, we make use of the following limit results, each of which can be easily verified:

(i) $\lim_{\rho \to 0} \beta_1^* = 0$.

(ii) $\lim_{\rho \to 0} \rho V_{(0,1)}^*(\beta) = (1-\beta)\Pi(0,0) + \beta\Pi^F$ (i.e. discounted value approaches the long-run average).

(iii) $\lim_{\rho \to 0} \rho V_{(1,0)}^*(\beta) = (1-\beta)\Pi(0,0) + \beta\Pi^L$ (i.e. discounted value approaches the long-run average).

For all $\rho > 0$, we have $\rho V_{(0,1)}^*(\beta_2^*) = \Pi(0,1) + C(1) - \beta_2^*\lambda(V^L - V^F)$, by definition. In the limit as $\rho \to 0$, the left-hand side of this equation remains finite (by the above limit result). Therefore, the right-hand side must also remain finite, which implies that $\lim_{\rho \to 0} \beta_2^* = 0$ because $V^L - V^F$ increases without bound as discounting becomes small. Likewise, $\hat{\beta}$ is defined by the equation $\rho V_{(1,0)}^*(\hat{\beta}) = \rho V_{(0,1)}^*(\hat{\beta})$. As $\rho \to 0$, this inequality must continue to hold. However, by our above limit result, the only way this is possible is if $\lim_{\rho \to 0} \hat{\beta} = 0$. From this, we conclude that $\max\{\beta_2^*, \hat{\beta}\}$ approaches zero as discounting vanishes.

Now consider the limit $\lim_{\rho \to 0} \beta_2^S$. By definition, β_2^S is defined by the indifference condition $\rho V_S(\beta_2^S) = TS(0,1) + C_S(0,1) > 0$. As $\rho \to 0$, the planner's discounted value (for analogous reasons to our second limit result) converges to $(1-\beta)TS(0,0) + \beta TS^+$. Consequently, we have

$$\lim_{\rho \to 0} \beta_2^S = \frac{TS(0,1) - TS(0,0) + C_S(0,1)}{TS^+ - TS(0,0)} > 0.$$

Therefore, as $\rho \to 0$, we will eventually have $\beta_2^S > \max\{\beta_2^*, \hat{\beta}\}$. So the intensity of equilibrium R&D, in any Markov-perfect equilibrium, will be, on average, higher than is socially desired. $\square$