WORKING PAPERS

N° 377

February 2013

Researcher's Dilemma

Catherine Bobtcheff, Jérôme Bolte and Thomas Mariotti

INSTITUT D'ECONOMIE INDUSTRIELLE

Toulouse School of Economics

Researcher's Dilemma*

Catherine Bobtcheff[†] Jérôme Bolte[‡] Thomas Mariotti[§]

February 13, 2013

Abstract

We model academic competition as a game in which researchers fight for priority. Researchers privately experience breakthroughs and decide how long to let their ideas mature before making them public, thereby establishing priority. In a two-researcher, symmetric environment, the resulting preemption game has a unique equilibrium. We study how the shape of the breakthrough distribution affects equilibrium maturation delays. Making researchers better at discovering new ideas or at developing them has contrasted effects on the quality of research outputs. Finally, when researchers have different innovative abilities, speed of discovery and maturation of ideas are positively correlated in equilibrium.

Keywords: Academic Competition, Preemption Games, Private Information.
JEL Classification: C73, D82.

*We would like to thank Bruno Biais, Jacques Crémer, Pierre Dubois, Leo Ferraris, Daniel Garrett, Faruk Gul, Johannes Hörner, Bruno Jullien, Yassine Lefouili, Lucas Maestri, Philippe Mahenc, David Martimort, Martin Peitz, Andrzej Skrzypacz, and Jean Tirole for very valuable feedback. We also thank seminar audiences at Toulouse School of Economics, Università degli Studi di Roma "Tor Vergata," Universität Mannheim, Université de Franche-Comté, Université de Lausanne, Université de Strasbourg, and Université Montpellier 1, as well as conference participants at the 12th Conference of the Society for the Advancement of Economic Theory for many useful discussions. Financial support from the Agence Nationale de la Recherche (ANR–10–BLAN–0112 JEUDY) and the European Research Council (Starting Grant 203929–ACAP) is gratefully acknowledged.

[†]Toulouse School of Economics (CNRS, LERNA).

[‡]Toulouse School of Economics (GREMAQ).

[§]Toulouse School of Economics (CNRS, GREMAQ, IDEI).

1 Introduction

As pointed out by Stigler (1963), and more recently evidenced by Card and DellaVigna (2013), competition is no less fierce in academic circles than in the market place. One dimension, in particular, along which researchers are very eager to compete is *priority*: in the field of scientific discovery, there is often relatively little value in being a follower. This leads researchers to choose with great care when to publicize new results or theories. A case in point—and a warning—is the publication by Charles Darwin of his theory of evolution through natural selection. After his attention was drawn in 1856 to a paper by the naturalist Alfred Russel Wallace on the "introduction of new species," Darwin was torn between the desire to produce a complete account of his theory and its applications, and the urgency of publishing a short paper summarizing its main insights. It is only when, upon receiving in 1858 a second parcel from Wallace, Darwin realized that he had been "forestalled" and thus was running the risk of losing priority, that he decided to "publish a sketch of [his] general views in about a dozen pages or so" (Darwin (1887, pages 116–117)).[1] The tradeoff faced by Darwin is familiar to any researcher. On the one hand, postponing the publication of his results allowed him to present a more mature theory, to increase the amount of evidence in favor of it, and to answer most likely objections before publishing. On the other hand, it also increased the risk of his being preempted, depriving him of the fruit of his efforts. The present paper investigates this tradeoff in a strategic preemption model in which researchers compete for priority. In this context, we aim at highlighting the following questions: How does competition affect the way researchers let their ideas mature? Does a more talented scientific community tend to produce more accomplished works? Do technological progress or human-capital accumulation necessarily foster research quality?

An important feature of academic competition is that it is often difficult for a researcher working on a new idea to identify her competitors, and even to ascertain whether there are any. The reason is that experiencing a breakthrough is necessary to be an active player in the field, but breakthroughs are typically observed privately by those who experience them—if only because there is a strong incentive to keep them secret to avoid being imitated and let the corresponding ideas mature optimally. Thus, to some extent, a researcher developing a new idea and waiting to publish her results works in the dark: she does not know whether she has any active competitor until it is too late and she already has been preempted. As a result, the more she waits, the more she runs the risk that an opponent experiences a breakthrough

[1] In the end, Darwin's and Wallace's papers were jointly read at the Linnean Society on July 1, 1858 (Desmond and Moore (1991)).

and therefore finds herself in a position to publish first. Our model explicitly incorporates this *potential competition* feature: starting from some initial time that can be interpreted as the date of a common knowledge event, such as a pioneering discovery, that opens up a new research field, each researcher effectively comes into play at some exogenous random time at which she experiences a breakthrough and that is her private information. From then on, she can make a move whenever she likes, letting her idea mature before making it public. Her payoff from moving first depends on the corresponding maturation delay. A key feature of our model is the existence of an optimal maturation delay. This may for instance result from a tradeoff between discounting and the expected return from publishing in terms of academic citations, career prospects, future funding for research, or self-esteem. In an alternative interpretation of the model in which the maturation delay affects the number and length of refereeing rounds—for instance because a more mature paper is less liable to attract criticism—, the optimal maturation delay minimizes the total delay from initial breakthrough to journal publication. Another important feature of our model is that the competitive pressure each researcher is subjected to need not stay constant, in the sense that the rate at which her opponents experience breakthroughs may vary. This may for instance reflect the evolution of technology or human capital, the growth of the research community, or exogenous fashion trends that dictate whether a given scientific field becomes "hotter" or "colder" over time.

For simplicity, we first consider a two-researcher environment with symmetric payoff functions and breakthrough distributions. In this context, there exists a unique, symmetric equilibrium that is described by a differential equation. This equation reflects the tension between the rate at which a researcher's payoff from moving first grows as her project becomes more mature, and the rate at which her opponent experiences breakthroughs and thereby becomes active. How this second rate evolves over time is key for the qualitative properties of equilibrium and, in particular, for how the equilibrium maturation delay evolves as a function of the breakthrough time. In a benchmark case in which the breakthrough rate is constant, each researcher waits the same amount of time before making her research public after experiencing a breakthrough. By contrast, if the breakthrough rate tends to increase over time, for instance as a result of technological progress or human capital accumulation, each researcher's equilibrium maturation delay decreases as a function of the time at which she experiences a breakthrough: the more time elapses, the more researchers tend to engage in defensive publication, attempting for instance to publish incomplete results to establish priority. In this case, our model predicts a negative correlation between the breakthrough time since the opening of a new field and the quality of the corresponding research output: the most accomplished contributions tend to occur relatively early. Assuming that young researchers are typically the ones who invest in new fields, and also experience breakthroughs less often than once they have become more experienced, our model thus provides a new mechanism explaining why most influential work may be done early in a scientist's career (see, in the case of physics, Levin and Stephan (1991) or, in the case of medical science, Evans (2007)); interestingly, this explanation relies on competition between increasingly innovative researchers and not, for instance, on a negative effect of tenure.[2] Accordingly, at the other end of the time spectrum, the perspective is rather grim: if asymptotically the breakthrough rate becomes large, the risk of being preempted erodes most of the gains from letting one's ideas mature; then, conditionally on not having experienced breakthroughs earlier on, researchers only produce "quick-and-dirty" papers. Our model does not commit us to this conclusion, however, and is flexible enough to allow for alternative scenarios. This is especially important as the widespread sentiment that great scientific advances are typically made by young researchers (see, for instance, Zuckerman and Merton (1973) and Simonton (1988)) has been recently challenged in the empirical literature on age and scientific creativity. Jones (2010) argues that great scientific achievements tend to occur at later and later ages throughout the last century, an effect he attributes to the lengthening of the training period during which researchers mostly undertake educational investments (see also Jones (2009), Jones and Weinberg (2011), and Dubois, Rochet, and Schlenker (2012)). Interestingly, this effect is magnified if their creativity is greatest when young. Under this alternative assumption, our model predicts that within a cohort of researchers, the fear of preemption is highest among young researchers, and fades out when they grow older. In those circumstances, the more accomplished contributions tend to occur relatively late in a researcher's career. Our model is useful in that it predicts how the shape of the breakthrough distribution affects equilibrium maturation delays and thus the quality of research outputs.

Because of the simplicity of its equilibrium characterization, our model lends itself very naturally to comparative statics analyses. We consider two such exercises. First, we study how equilibrium maturation delays evolve when researchers become more innovative—an overall shift in their breakthrough rates. Second, we study how equilibrium maturation delays evolve when researchers work on projects with a higher growth potential—an overall shift in the rate at which a researcher's payoff from moving first grows as her project becomes more mature. These two kinds of changes in the efficiency of scientific activity, which can be

[2]As pointed out below, this does not mean that our model cannot shed light on tenure effects.

respectively described as changes in the discovery process and changes in the development process, have contrasted effects on equilibrium maturation delays and thus on the quality of research outputs. On the one hand, when researchers become more efficient at discovering ideas, this increases their fear of being preempted and thus speeds up the maturation process, leading to less accomplished works. On the other hand, when researchers become more efficient at developing ideas, this increases their marginal gain of letting their ideas mature relative to the potential loss of being preempted and thus slows down the maturation process, leading to more accomplished works.[3] When a longer maturation delay reduces the number and length of refereeing rounds, this result leads to the prediction that, everything else equal, if competition takes place between researchers eager to publish quickly, such as assistant professors on a tenure track, equilibrium maturation delays should be longer and publication delays shorter than if it takes place between less impatient researchers, who for instance already enjoy tenure.

An interesting extension of our analysis consists in studying asymmetric contests in which one researcher (the "hare") tends to experience breakthroughs at a higher rate than the other (the "tortoise"). Such differences in aptitudes particularly matter when the competing entities are research labs instead of individuals, in which case an advantage in innovative ability may typically result from factors such as better funding, higher interdisciplinarity, stronger leadership, and other organizational features (see, in the case of medical science, Hollingsworth and Hollingsworth (2000)). In the simplest scenario in which the hare and the tortoise have constant but different breakthrough rates, we show that there exists a unique continuous equilibrium. In this equilibrium, the hare always lets her ideas mature more than the tortoise, no matter when they experience their breakthroughs: the hare endogenously behaves more ambitiously than the tortoise and thus succeeds or fails more spectacularly than her, echoing a theme in March (1991). This leads to the prediction that within a group of competing researchers of unequal aptitudes, speed of discovery and maturation of ideas should be positively correlated. A closely related implication of our analysis is that there is more heterogeneity in the quality of differently apt researchers' outputs when breakthroughs occur relatively late; in our model, this is because the hare features a flight to quality when she experiences breakthroughs later on, whereas the reverse holds true for the tortoise. Finally, we show that an increase in the gap between the hare's and the tortoise's innovative abilities unambiguously deteriorates the quality of the tortoise's research outputs, and under

[3]Note that a monopolistic researcher would not be interested in the *rate* at which ideas mature per se: she would only care about the maximum achievable payoff *level*. By contrast, a competing researcher fears he might be preempted and thus compares marginal gains to absolute losses.

certain circumstances, that of the hare as well. Hence becoming more innovative is a mixed blessing, because it may create a race to the bottom by increasing one's opponent fear of being preempted. Our results on asymmetric contests may help shed light on the findings of Borjas and Doran (2012) on the post-1992 influx of highly-skilled Soviet mathematicians on the scientific production of their US counterparts. Interestingly, they not only show that there is a large drop in the publication rate of mathematicians whose research agenda overlapped most with the Soviets, but also that the quality of their papers, as measured by the number of citations they generated or by their likelihood of becoming "home runs" also fell significantly. This is consistent with our finding that, when she faces a hare, a tortoise tends to behave more cautiously than if she were facing an opponent of equal strength, which reduces the quality of her research output.

Related Literature Few theoretical attempts have been made at modeling the behavior of academic researchers. Sharon and Levin (1991) study a model of research productivity over the life cycle, in which a scientist with both extrinsic and intrinsic motivation for research dynamically allocates her time between research and nonresearch activities such as teaching and consulting. They find that research activity should decline over the life cycle, a prediction which is broadly supported by the data they consider. Our analysis neglects life-cycle effects to focus on the strategic interaction between researchers competing for priority.

Ellison (2002b) studies a model in which researchers decide how much time to invest to improve the quality of their paper's central contribution or to enhance other aspects of quality, such as generality or robustness, that are typically dealt with in revisions. Given scarce journal space, a social norm for publication determines the relative weights attributed to these two quality dimensions, as well as the minimum overall level of quality required for publication. In a static version of his model, Ellison (2002b) shows that there exists a continuum of equilibrium social norms. He then studies a dynamic model of the evolution of such norms that is consistent with the slowdown of the publishing process and the increased length of papers in the economics profession (Ellison (2002a)). Our paper emphasizes priority as a complementary dimension along which researchers compete. Quality of research outputs is then determined by the interplay between the dynamics of innovative ability and the growth potential of research projects.

From a theoretical viewpoint, this paper belongs to the literature on preemption games, that is, timing games with a first-mover advantage. In a seminal paper on the strategic adoption of new technology, Fudenberg and Tirole (1985) show in a complete-information setting that there always exists a subgame-perfect equilibrium in which firms' payoffs are equalized and rents are fully dissipated. This does not arise in our setting (except in the limit case where breakthrough rates become arbitrarily large) because, as in much of the literature, wasteful competition is alleviated by the asymmetry of information between players. As a result, there is a genuine tradeoff between the gains from letting one's project mature and the risk of preemption.

In a real-option context, Lambrecht and Perraudin (2003) study how firms trade off the benefit from learning about the future return of an uncertain project and the risk of being preempted when they do not know each others' investment costs. Under some assumptions on the distribution of costs, they show that there exists a unique equilibrium in the case of a symmetric duopoly. Anderson, Friedman, and Oprea (2012) extend this model to the case of an arbitrary finite number of identical firms, and characterize a symmetric equilibrium. Moreover, they run an experiment which confirms that competition hastens investment and that the lowest-cost investor in a triopoly usually preempts the others. A key feature of our model that distinguishes it from such technology-adoption or investment models is that competition is only potential, in the sense that private information bears on the time at which each player becomes an active competitor. Moreover, unlike Lambrecht and Perraudin (2003) and Anderson, Friedman, and Oprea (2012), we also address the case of asymmetric players with different innovative abilities.

To capture the fact that R&D competitors usually prefer to keep their innovations secret before filing for a patent, Hopenhayn and Squintani (2011) study a preemption game in which players' private information stochastically increases over time and the value of the innovation is uncertain. They construct an equilibrium in which each player terminates the game when her state is above a time-decreasing threshold that is the solution of a differential equation. In our model, breakthroughs are secret, but the payoff from making one's research public is a deterministic function of the maturation delay. In this simpler setting, we are able to show the existence of a unique equilibrium, and to study how changes in the breakthrough distribution and in the payoff function affect researchers' equilibrium incentives to let their projects mature.

The idea that players in preemption games may face uncertainty about whether they have active competitors has first been introduced by Hendricks (1992), who extends Fudenberg and Tirole's (1985) analysis to the case where it is determined at the outset of the game whether firms are innovators or imitators, in which case they cannot move first. Innovators have an incentive to build a reputation for being imitators, which alleviates rent dissipation; they reveal their information gradually by playing according to a mixed strategy. Bobtcheff and Mariotti (2012) consider a setting closer to the one developed in the present paper, in which players randomly and secretly come into play; they show that all equilibria give rise to the same distribution for each player's moving time. However, in their model, as in Fudenberg and Tirole (1985) and Hendricks (1992), the payoff a player derives from making a move first only depends on calendar time and not, as in our model, on the time elapsed since she experienced a breakthrough; as a result, a player who comes into play late is not at a disadvantage relative to one who came into play earlier on but did not make a move in the meanwhile. By contrast, we consider situations in which ideas take time to mature, a more appropriate assumption in the case of academic competition.

Closely related to this paper, Hopenhayn and Squintani (2010) consider a sequential model of R&D races in which research builds on previously patented products. In a given race, firms experience breakthroughs at a constant common rate, and decide when to make them public and file for a patent. In each race, a constant proportion of firms is randomly selected to participate. Therefore, a firm that joins in a race does not know whether it will take part in the following ones. Hopenhayn and Squintani (2010) derive the unique symmetric equilibrium, that consists for firms in waiting a constant amount of time following a breakthrough. They provide several comparative statics results as well as a comparison with the social optimum. Our model is simpler in that we focus on a single race, but we allow for arbitrary breakthrough distributions. We can thus study how researchers' strategies evolve in nonstationary environments in which, for instance, technological progress or human-capital accumulation change their innovative abilities over time. Another distinctive feature of our analysis is that we also consider asymmetric contests.

Abreu and Brunnermeier (2003) and Brunnermeier and Morgan (2010) analyze models in which each player receives, at a random and secret time, a signal about some payoff-relevant state variable. Players must decide when to exit the game, given that waiting increases their gains but also the probability of their being preempted. Because of the lack of synchronization between players' actions, usual unravelling arguments do not apply. Abreu and Brunnermeier (2003) show that, in a financial market with some behavioral agents, this mechanism can explain the persistence of bubbles even in the presence of rational and financially unconstrained arbitrageurs: because their clocks are desynchronized, each of them is ready to ride the bubble as it continues to grow and generates high returns. Accordingly, Abreu and Brunnermeier (2003) show the existence of a unique trading equilibrium in which arbitrageurs liquidate their positions a constant amount of time after becoming aware that the price has departed from fundamentals. A similar result is derived by Brunnermeier and Morgan (2010) in their analysis of clock games. In our model, the researchers' clocks are synchronized, but their optimal publication times are not. Moreover, unless in the benchmark case in which breakthrough rates are constant and identical across researchers, the equilibrium is nonstationary: because the competitive environment changes over time, researchers typically wait different amounts of time before making their research public depending on when they experience their breakthroughs.

The paper is organized as follows. Section 2 describes the model. Section 3 analyzes the pure-strategy equilibrium of the symmetric game, and provides comparative statics results. Section 4 extends the analysis to a class of asymmetric contests. Section 5 concludes. All the proofs are gathered in a technical appendix.

2 The Basic Model

Time is continuous, and indexed by $t \geq 0$. There are two symmetric players, a and b.[4] In what follows, i refers to an arbitrary player and j to her opponent. Player i comes into play at some privately observed random time $\tilde{\tau}^i \geq 0$. For instance, $\tilde{\tau}^i$ can be interpreted as the time at which player i experiences a breakthrough and discovers a new and promising idea. An important assumption is that calendar time is common knowledge; hence players' clocks are synchronized, unlike in Abreu and Brunnermeier (2003) or Brunnermeier and Morgan (2010). In this context, time zero can be interpreted as the date of a common knowledge event, such as a pioneering discovery, that enables the players to make progress and, in turn, experience breakthroughs.

Actions and Payoffs As in standard timing games, each player i has a single opportunity to make a move. We take the simplifying view that in the case of academic competition, making a move, and thereby establishing priority, consists for player i in making her research public. A key feature of the model, that we share with Hopenhayn and Squintani (2010) and Bobtcheff and Mariotti (2012), is that this must occur at some time $t^i \geq \tilde{\tau}^i$, reflecting that player i cannot make a move before experiencing a breakthrough.

Both players are risk-neutral and discount future payoffs at the same rate. Payoffs are defined as follows. If player i, having experienced a breakthrough at time τ^i, makes a move first at time t^i, she obtains a payoff $L(t^i - \tau^i)$ in time-τ^i terms, whereas player j's payoff is zero. Therefore, player i's payoff from preempting her opponent is a function of the

[4] Extending the analysis to an arbitrary number of symmetric players is straightforward, see Section 3.6. The case of asymmetric players is more challenging; a special case is explored in Section 4.

difference $m^i \equiv t^i - \tau^i$, which we call player i's *maturation delay*. Finally, if players i and j, having experienced breakthroughs at times τ^i and τ^j, simultaneously make a move at time t, their payoffs are $\alpha L(t - \tau^i)$ and $\alpha L(t - \tau^j)$ in time-τ^i and time-τ^j terms, where $1 - \alpha \in [0, 1]$ measures each player's proportional loss from making a simultaneous move rather than preempting her opponent.[5]

It is useful to contrast this payoff structure with that considered in Bobtcheff and Mariotti (2012). In their model, each player's payoff from preempting one's opponent only depends on the time at which she makes a move, and not on the time at which she experienced a breakthrough. This intuitively corresponds to a situation in which payoffs are driven by purely exogenous factors, such as the evolution of market demand, as in Gilbert and Harris (1984), or the advancement of a new technology, as in Reinganum (1981) or Fudenberg and Tirole (1985), and thus, ultimately, only depend on calendar time. By contrast, the present model depicts a world in which ideas take time to mature. This typically arises in environments where innovation is at stake. For instance, experiencing a breakthrough is hardly enough to deliver an accomplished academic work, because, to be valuable, any innovation necessitates a development period. We will maintain the following assumption throughout the paper.[6]

Assumption 1 *The function* $L : [0, \infty) \to [0, \infty)$ *is continuous and twice continuously differentiable over* $(0, \infty)$, *and there is a maturation delay* $M > 0$ *such that*

$$\begin{aligned} L(0) &= 0, \\ L(m) &> 0 \ \text{ if } \ m > 0, \\ \dot{L}(m) &> 0 \ \text{ if } \ m < M, \\ \dot{L}(m) &< 0 \ \text{ if } \ m > M, \\ \ddot{L}(m) &< 0 \ \text{ if } \ m < M. \end{aligned}$$

The function L only vanishes at the origin: thus a positive maturation delay is required for a breakthrough to have any value. The function L then increases, reaches a maximum at M, after which it decreases. These monotonicity properties imply that experiencing a breakthrough relatively early on gives a player a double advantage over her opponent: she can not only make a move earlier than her opponent does, but she also need not wait as much in order to enjoy the same payoff. Finally, the function L exhibits decreasing returns over

[5] Whereas specifying payoffs in case of a tie is necessary to fully specify the model, it turns out that players never simultaneously make a move in equilibrium.

[6] We use dots to denote time derivatives.

the part of its domain where these returns are positive: in applications, this typically reflects discounting of future payoffs, and possibly decreasing returns over time of the development technology. As we discuss in Section 3.1, a key implication of this concavity property, coupled with the assumption that the payoffs from making a move depend on maturation delays and not on calendar time, is that player i's preferences satisfy a strict single-crossing property in (t^i, τ^i). This stands in sharp contrast with Bobtcheff and Mariotti (2012) where, because payoffs only depend on calendar time, this property fails to hold.

By construction, M is the optimal maturation delay for a player that would not be threatened by preemption, and could thus act as a monopolist. In equilibrium, no player will wait more than M units of time to make a move after experiencing a breakthrough, see Lemma 1(i) below. At the other extreme, a zero maturation delay also has a natural game-theoretic interpretation. Suppose indeed that it were common knowledge that both players simultaneously experienced a breakthrough at time τ. Then, from time τ on, we would have a standard complete-information preemption game. In this game, each player is indifferent at time τ between making a move or abstaining, because both options yield a zero payoff. In this context, one can extend Fudenberg and Tirole's (1985) arguments to show that in any subgame-perfect equilibrium, equilibrium maturation delays are zero and rents are fully dissipated.[7]

Example 1 A parametric case in which Assumption 1 holds is when

$$L(m) = \exp(-rm)[\exp(\mu m) - 1], \quad m \geq 0, \tag{1}$$

where $r > \mu > 0$, which arises when players discount future payoffs at rate r and can invest in a research project that involves a setup cost 1 (reflecting for instance the cost of wrapping up a paper before sending it for publication), and generates a payoff which grows at rate μ (reflecting for instance that more time allows researchers to conduct additional experiences or robustness checks that increase the impact of the paper). Then

$$M = \frac{1}{\mu} \ln\left(\frac{r}{r-\mu}\right), \tag{2}$$

and it is easy to verify that L is strictly concave over $[0, M]$.[8]

[7] It should be observed that the proper formulation of this result requires an appropriate extension of the players' strategy spaces, see Fudenberg and Tirole (1985) or Simon and Stinchcombe (1989).

[8] Note that one need not assume in this example that the payoff from making a move grows at a constant rate μ: to capture learning-curve effects, one may for instance assume that this payoff is a logistic function $(1+a)/[1+a\exp(-\mu m)]$ of the maturation delay m, where a is a positive constant.

Example 2 Example 1 relies on the interplay between the discounting of future rewards and costs and the growth potential of a research idea. An alternative story is that the benefit of a longer maturation delay may come in the form of faster publication. Specifically, let us now interpret m as the delay between a breakthrough and the first submission of the corresponding paper to a journal. Assume for simplicity that sending the paper to a journal ensures eventual publication and priority. However, this does not ensure that the paper will be published right away: a publication delay now adds to the initial maturation delay.[9] This publication delay may for instance reflect the time it takes referees to understand the paper's contribution, and thus the number and length of refereeing rounds. It is natural to postulate that the publication delay is a function $p(m)$ of the initial maturation delay m. In particular, the higher m, the more elaborate the first submission, and thus the shorter $p(m)$ is likely to be. We assume that $\lim_{m \downarrow 0} p(m) = \infty$, so that very immature papers ultimately get stuck in the publication process, and that the total delay $m + p(m)$ from breakthrough to publication is minimized for $m = M$. We also normalize publication delays so that $p(m) = 0$ for all $m \geq M$, which implies that a maturation delay as long as M guarantees immediate publication. Players discount at rate r a future reward from publication normalized to 1 which accrues at time $m + p(m)$. A parametric case in which Assumption 1 holds is when $p(m) = M \ln(M/m) \vee 0$ and thus[10]

$$L(m) = \exp\left(-r\left\{m + \left[M \ln\left(\frac{M}{m}\right) \vee 0\right]\right\}\right), \quad m \geq 0, \tag{3}$$

where the condition $M < 1/r$ ensures that L is strictly concave over $[0, M]$.[11]

Information We hereafter assume that the players' breakthrough times $\tilde{\tau}^a$ and $\tilde{\tau}^b$ are independently and identically distributed, with a distribution function F that is continuously differentiable and has a positive density $\dot{F}$ over $[0, \infty)$. In light of our interpretation of time zero, the independence of breakthrough times should be understood conditionally on the event that takes place at this time and makes breakthroughs possible. Whereas the breakthrough distribution is common knowledge, the breakthrough time of each player is her private information, or *type*. In particular, when a player experiences a breakthrough, she does not know whether or not her opponent already experienced one; nor does she observe when her opponent experiences a breakthrough. As a result, the only public information that

[9] Alternatively, one may assume that the publication prospects are in fact uncertain, and that priority is not ensured with probability 1.

[10] For any real numbers x and y, $x \vee y$ denotes the maximum of x and y, and $x \wedge y$ denotes the minimum of x and y.

[11] Strictly speaking, the function (3) does not satisfy Assumption 1 as it is not differentiable at M. Yet, this is inconsequential, as $\dot{L}(M-) = 0$ and L is strictly decreasing over $[M, \infty)$, see the proof of Lemma 1(i).

accrues to each player during the course of the game is whether and when her opponent makes a move, which effectively terminates the game. Thus, as in Hendricks (1992), Hopenhayn and Squintani (2010), or Bobtcheff and Mariotti (2012), a distinctive feature of the present model compared to most analyses of preemption games is that competition is only potential: a player never knows for sure whether she actually has an active competitor until it is too late and she already has been preempted.

Strategies and Equilibria We focus on pure-strategy equilibria in which the equilibrium moving time of player i is described by a function $\sigma^i : [0, \infty) \to [0, \infty)$ that specifies, for each possible value τ^i of her breakthrough time, the time $\sigma^i(\tau^i)$ at which she plans to make a move. Of course, an admissible strategy must satisfy $\sigma^i(\tau^i) \geq \tau^i$ for all $\tau^i \geq 0$, reflecting the feasibility constraint that the maturation delay $\sigma^i(\tau^i) - \tau^i$ must be nonnegative. Player i's payoff if her type is τ^i, player j's strategy is σ^j, and player i plans to make a move at time $t^i \geq \tau^i$ if player j has not made a move by then is

$$V^i(t^i, \tau^i, \sigma^j) \equiv \{\mathbf{P}[\sigma^j(\tilde{\tau}^j) > t^i] + \alpha \mathbf{P}[\sigma^j(\tilde{\tau}^j) = t^i]\} L(t^i - \tau^i). \tag{4}$$

A pair (σ^a, σ^b) is a Bayesian equilibrium if for each i, $\tau^i \geq 0$, and $t^i \geq \tau^i$,

$$V^i(\sigma^i(\tau^i), \tau^i, \sigma^j) \geq V^i(t^i, \tau^i, \sigma^j). \tag{5}$$

Our formulation of the model allows players' breakthrough times and hence players' moving times to be arbitrarily large. Therefore, in equilibrium, the only zero-probability events could be some player unexpectedly making a move. As this effectively terminates the game, we need not be concerned with how players behave after such events, so that our Bayesian equilibria are actually perfect Bayesian equilibria.

Remark As the definition (4) of payoffs makes clear, our preemption game bears a strong formal analogy with a first-price procurement auction with risk-averse bidders (Arozamena and Cantillon (2004)): in this context, τ^i would be bidder i's monetary cost of providing a good, t^i the price offered by bidder i, and L the bidders' common von Neumann–Morgenstern utility function over final wealth. However, our model has two novel features that make the analysis quite different. First, whereas costs in procurement-auction models are typically distributed over a bounded interval, or there is a reserve price set by the buyer, breakthrough times in our model can take arbitrarily large values. Second, unlike a standard von Neumann–Morgenstern utility function over final wealth, the payoff function L reaches a maximum at the optimal maturation delay M. As we shall now see, these two features imply that, whereas the differential equation that describes equilibrium play is the same in our model as in a first-price procurement auction, the boundary conditions it is subject to are of a different nature and call for a specific analysis.

3 Equilibrium Analysis

3.1 Preliminaries

To begin with, we establish three intuitive yet useful properties that hold true for any equilibrium of the game.

Lemma 1 *In any equilibrium,*

(i) $0 < \sigma^i(\tau^i) - \tau^i \leq M$ *for all* i *and* τ^i.

(ii) σ^i *is strictly increasing for all* i.

(iii) $\sigma^a(0) = \sigma^b(0)$.

The first result is very intuitive. Because $L(0) = 0$, a zero maturation delay for player i cannot arise in equilibrium, because it can only be rationalized by the belief that player j will never make a move after player i's breakthrough time. This belief, however, is inconsistent with the fact that with positive probability, player j experiences a breakthrough later than player i does. Similarly, a maturation delay for player i above the optimal maturation delay M cannot arise in equilibrium, because reducing her maturation delay would strictly increase player i's payoff from moving first, while not increasing her risk of being preempted.

The second result consists of two parts. (1) That player i's equilibrium strategy is nondecreasing follows from a standard revealed-preference argument, using the property that given any strategy σ^j of player j, player i's payoff function as defined in (4) satisfies strict single crossing of incremental returns in (t^i, τ^i): for $\hat{\tau}^i > \tau^i$ and $\hat{t}^i > t^i$ in the relevant range, $V^i(\hat{t}^i, \tau^i, \sigma^j) \geq V^i(t^i, \tau^i, \sigma^j)$ implies that $V^i(\hat{t}^i, \hat{\tau}^i, \sigma^j) > V^i(t^i, \hat{\tau}^i, \sigma^j)$.[12] Intuitively, this reflects that because there are decreasing returns to letting one's idea mature, a player with a relatively early breakthrough time has more to lose from being preempted, and less to gain at the margin from slightly delaying her move. (2) That player i's equilibrium strategy is actually strictly increasing is more subtle to establish. If this were not the case, there would be an atom in the distribution of player i's moving time. One can then show that player j would never find it profitable to make a move during some interval of time following this

[12] Observe for instance that $(\partial V^i / \partial \tau^i)(t^i, \tau^i, \sigma^i) = -\{\mathbf{P}[\sigma^j(\tau^j) > t^i] + \alpha \mathbf{P}[\sigma^j(\tau^j) = t^i]\}\dot{L}(t^i - \tau^i)$ is strictly decreasing in $t^i \in [\tau^i, \tau^i + M]$ as L is strictly concave over $[0, M]$.

atom, which in turn implies that some types of player i who make a move at this atom would find it profitable to slightly delay their move, as they locally would not fear to be preempted. Such an atom cannot therefore arise in equilibrium, so that player i's equilibrium moving time $\sigma^i(\tau^i)$ must be a strictly increasing function of her breakthrough time τ^i. Observe that this result says nothing about how player i's equilibrium maturation delay $\sigma^i(\tau^i) - \tau^i$ varies as a function of her breakthrough time τ^i.

Given these two results, the third one is immediate. Suppose indeed that $\sigma^i(0) < \sigma^j(0)$, so that player i, when she experiences a breakthrough at time zero, waits less to make a move than player j would under the same circumstances. Then, as the equilibrium strategy of player j is strictly increasing, player i when she experiences a breakthrough at time zero faces no risk of preemption from player j. Moreover, as the equilibrium maturation delay $\sigma^j(0)$ of player j when she experiences a breakthrough at time zero is no greater than the optimal maturation delay M, the equilibrium maturation delay $\sigma^i(0)$ of player i when she experiences a breakthrough at time zero is strictly less than M. These two observations imply that if one had $\sigma^i(0) < \sigma^j(0)$, player i, were she to experience a breakthrough at time zero, would have a strict incentive to wait slightly longer to make a move: she could thereby let her idea mature more, while still avoiding any preemption risk. This contradiction shows that in equilibrium the two players' maturation delays must coincide when they experience a breakthrough at time zero. Define accordingly $\sigma(0) \equiv \sigma^a(0) = \sigma^b(0)$.

Based on Lemma 1, one can establish two useful regularity properties of equilibria. For each i, let $\phi^i \equiv (\sigma^i)^{-1}$ be the inverse of σ^i, which is well defined and continuous over $\sigma^i([0,\infty))$ by Lemma 1(ii). Then the following holds.

Lemma 2 *In any equilibrium,*

(i) σ^i *is continuous, so that* $\sigma^i([0,\infty)) = [\sigma(0),\infty)$.

(ii) ϕ^i *is differentiable over* $[\sigma(0),\infty)$.

The first result is very intuitive in the case of a symmetric equilibrium in which the two players play the same strategy $\sigma : [0,\infty) \to [0,\infty)$. Indeed, if the common equilibrium strategy σ in a symmetric equilibrium were discontinuous at some breakthrough time τ, there would be a gap $(\sigma(\tau-), \sigma(\tau+))$ in the distribution of moving times. But then one should have $\sigma(\tau-) < \tau + M$, for, otherwise, a player who experiences a breakthrough just after time τ would have a maturation delay strictly longer than M, which is ruled out by Lemma 1(i). This in turn implies that a player who experiences a breakthrough just before time τ would be strictly better off delaying her move until some time in $(\sigma(\tau-), \sigma(\tau+))$. Indeed, as $\sigma(\tau-) - \tau < M$, this would significantly increase her payoff from moving first, while only marginally increasing the probability of being preempted because, by Lemma 1(ii), there is no atom in the distribution of her opponent's moving time and the probability that the latter makes a move during $(\sigma(\tau-), \sigma(\tau+))$ is zero. The corresponding argument in the case of asymmetric equilibria is more complex, and shows the continuity of strategies by establishing that their inverses satisfy appropriate regularity properties.

From the second result, we obtain a system of differential equations for the inverses ϕ^a and ϕ^b of σ^a and σ^b. This system can be heuristically derived as follows. Because σ^a and σ^b are strictly increasing according to Lemma 1(ii), and the breakthrough distribution has no atom by assumption, the probability of a tie is zero. The problem faced by type τ^i of player i can then be written as

$$\max_{t^i \in [\tau^i, \infty)} \{\mathbf{P}[\sigma^j(\tilde{\tau}^j) > t^i] L(t^i - \tau^i)\},$$

or equivalently

$$\max_{t^i \in [\tau^i, \infty)} \{[1 - F(\phi^j(t^i))] L(t^i - \tau^i)\}. \tag{6}$$

We know from Lemma 1(i) that a zero maturation delay is inconsistent with equilibrium, so that the solution to problem (6) must be interior. The first-order condition is

$$[1 - F(\phi^j(t^i))]\dot{L}(t^i - \tau^i) = \dot{F}(\phi^j(t^i))\dot{\phi}^j(t^i)L(t^i - \tau^i). \tag{7}$$

Intuitively, (7) expresses that type τ^i of player i cannot increase her payoff by making a move at time $t^i + dt$ instead of t^i, for dt positive and small enough. Indeed, the probability that player j makes a move during $[t^i, t^i + dt]$ is $\dot{F}(\phi^j(t^i))\dot{\phi}^j(t^i)dt$. Hence, for type τ^i of player i, delaying her move by an infinitesimal amount of time dt comes at an expected cost $\dot{F}(\phi^j(t^i))\dot{\phi}^j(t^i)L(t^i - \tau^i)dt + o(dt)$. With probability $1 - F(\phi^j(t^i + dt))$, however, player j plans to make a move after time $t^i + dt$. Thus, by delaying her move by dt, type τ^i of player i increases her payoff by $[1 - F(\phi^j(t^i))]\dot{L}(t^i - \tau^i)dt + o(dt)$ on average. Equating these expected marginal costs and benefits and neglecting terms of order $o(dt)$, we obtain (7). In equilibrium, this first-order condition must hold for $\tau^i = \phi^i(t^i)$, leading to the following system of nonautonomous ordinary differential equations (ODEs):

$$\dot{\phi}^j(t) = \frac{1 - F(\phi^j(t))}{\dot{F}(\phi^j(t))} \frac{\dot{L}(t - \phi^i(t))}{L(t - \phi^i(t))}, \quad t \geq \sigma(0), \quad i = a, b. \tag{8}$$

Now, observe that $\phi^a(\sigma(0)) = \phi^b(\sigma(0)) = 0$ as $\sigma^a(0) = \sigma^b(0) = \sigma(0)$ by Lemma 1(iii). Together with the differentiability of the functions ϕ^a and ϕ^b, this common initial condition rules out asymmetric equilibria.

Lemma 3 *Every equilibrium is symmetric.*

Consider accordingly a symmetric equilibrium where both players play the strategy σ, and let $\phi \equiv \sigma^{-1}$. The system (8) then reduces to the nonautonomous ODE

$$\dot{\phi}(t) = f(t, \phi(t)), \quad t \geq \sigma(0), \tag{9}$$

where by definition

$$f(t, \tau) \equiv \frac{1 - F(\tau)}{\dot{F}(\tau)} \frac{\dot{L}(t - \tau)}{L(t - \tau)}, \quad t > \tau. \tag{10}$$

Much of what follows is an extensive study of the ODE (9). As announced in Section 2, this ODE is formally similar to that arising in a procurement-auction model. However, three distinctive features of our analysis are worth noting at this stage. (1) First, the ODE (9) holds for all t larger than $\sigma(0)$, which is endogenous as $\sigma(0) = \phi^{-1}(0)$. In turn, $\sigma(0)$ must be chosen carefully, so as to ensure that in line with Lemma 1(i), the maturation delay $\sigma(\tau) - \tau$ remains in $(0, M]$ for all $\tau \geq 0$. Equivalently, we must choose $\sigma(0)$ so that $\phi(t) \in [t - M, t)$ for all $t \geq \sigma(0)$. Hence ϕ must never leave the domain

$$\mathcal{D} \equiv \{(t, \tau) : 0 \leq \tau < t \leq \tau + M\}.$$

This boundary condition is specific to our preemption model and reflects that types may take arbitrarily large values and that the function L reaches a maximum at M. (2) Next, Assumption 1 implies that $\lim_{\tau \uparrow t} f(t, \tau) = \infty$ for all $t \geq 0$ and $\lim_{\tau \downarrow t - M} f(t, \tau) = 0$ for all $t \geq M$. Hence the vector field induced by (10) (using both the breakthrough time τ and the moving time t as state variables) is outward-pointing both on the upper boundary $\tau = t$ and on the lower boundary $\tau = t - M$ of $\mathcal{D}$. Moreover, because L is strictly concave over $[0, M]$, the mapping $t \mapsto f(t, \tau)$ is strictly decreasing. Hence the vector field induced by (10) is monotone in the moving time dimension. These features reinforce the presumption that the choice of $\sigma(0)$ is indeed critical to ensure that ϕ never leaves $\mathcal{D}$. (3) Finally, observe that a maturation delay $t - \phi(t)$ equal to the optimal maturation delay M cannot arise in equilibrium, because this would imply that $\phi(t) = t - M$, which according to our second observation would cause ϕ to leave $\mathcal{D}$ through its lower boundary at time t. As a result, in equilibrium, ϕ must actually remain in the interior $\operatorname{Int} \mathcal{D}$ of $\mathcal{D}$. Intuitively, this reflects that the risk of being preempted is never fully eliminated in equilibrium.

Before we proceed with the analysis, it is worth noting that any solution ϕ to the ODE (9) that remains in $\operatorname{Int} \mathcal{D}$ also satisfies the second-order condition for problem (6). Indeed, for each $t \geq \sigma(0)$, one has

$$\frac{\dot{F}(\phi(t))}{1 - F(\phi(t))}\,\dot{\phi}(t) = \frac{\dot{L}(t - \phi(t))}{L(t - \phi(t))}.$$

Now, ϕ is strictly increasing over $[\sigma(0), \infty)$ as, by construction, $t - \phi(t) \in (0, M)$ for all $t \geq \sigma(0)$. Moreover, for each $\hat{t} \in [\sigma(0), \infty)$, the mapping $\tau \mapsto (\dot{L}/L)(\hat{t} - \tau)$ is strictly increasing over $[\hat{t} - M, \hat{t})$ as L is strictly concave and positive over $(0, M]$. It follows that

$$\frac{\dot{L}(\hat{t} - \phi(t))}{L(\hat{t} - \phi(t))} \gtreqless \frac{\dot{F}(\phi(\hat{t}))}{1 - F(\phi(\hat{t}))}\,\dot{\phi}(\hat{t})$$

if $t \gtreqless \hat{t}$. That is, for a player with type $\phi(t)$, the expected incremental payoff of marginally delaying her move is everywhere positive for $\hat{t} < t$, and everywhere negative for $\hat{t} > t$. As a result, the second-order condition for problem (6) is satisfied, and a player with type $\phi(t)$ optimally makes a move at time t. Observe that these global incentive compatibility conditions hold irrespective of the shape of the breakthrough rate $\dot{F}/(1 - F)$.

3.2 Constant Breakthrough Rate

A benchmark case for our analysis arises when the players' breakthrough times have an exponential distribution, that is, when the breakthrough rate $\dot{F}/(1 - F)$ is equal to some constant $\lambda > 0$, as in Hopenhayn and Squintani (2010). The ODE (9) then becomes

$$\dot{\phi}(t) = f_\lambda(t, \phi(t)) \equiv \frac{1}{\lambda}\,\frac{\dot{L}(t - \phi(t))}{L(t - \phi(t))}, \quad t \geq \sigma(0). \tag{11}$$

An obvious solution to this ODE is

$$\phi_\lambda(t) = t - \left(\frac{\dot{L}}{L}\right)^{-1}(\lambda), \tag{12}$$

corresponding to a constant maturation delay $M_\lambda \equiv (\dot{L}/L)^{-1}(\lambda)$. It is easy to see that ϕ_λ is the only solution to the ODE (11) that remains in $\operatorname{Int}\mathcal{D}$. Indeed, from the definition (11) of f_λ along with the fact that the mapping $\tau \mapsto (\dot{L}/L)(t - \tau)$ is strictly increasing over $[t - M, t)$, we obtain that $f_\lambda(t, \tau) \gtreqless 1$ if $\tau \gtreqless \phi_\lambda(t)$. Because $\dot{\phi}_\lambda = 1$, it follows that any solution to the ODE (11) that lies above ϕ_λ will eventually leave $\mathcal{D}$ through its upper boundary $\tau = t$, whereas any solution to the ODE (11) that lies below ϕ_λ will eventually leave $\mathcal{D}$ through its lower boundary $\tau = t - M$. This illustrates the instability of the ODE (11), and shows that in the constant-breakthrough-rate case there exists a unique equilibrium given by $\sigma_\lambda(\tau) = \tau + M_\lambda$ for all $\tau \geq 0$.

The equilibrium maturation delay M_λ has a natural interpretation. Suppose player i expects player j to wait a constant amount of time before making a move. Then, because breakthroughs occur at a constant rate λ, she expects player i to make a move at a constant rate λ. By waiting herself an amount of time M_λ, she equalizes the growth rate of her payoff from preempting player j to the rate at which player j makes a move, conditional on not having made a move yet. Hence, at this point, the marginal benefit for player i of further delaying her move by an infinitesimal amount of time dt, $\dot{L}(M_\lambda)\,dt$, exactly compensates the corresponding expected marginal loss, which is equal to the probability that player makes a move during $[\tau^i + M_\lambda, \tau^i + M_\lambda + dt)$ conditional on her not having done so yet, λdt, multiplied by the foregone benefit, $L(M_\lambda)$.

The unique equilibrium in the constant-breakthrough-rate case is illustrated on Figure 1, which also represents the phase portrait associated to (11). On the t-axis is the time at which a player makes a move, and on the τ-axis is the time at which a player experiences a breakthrough.

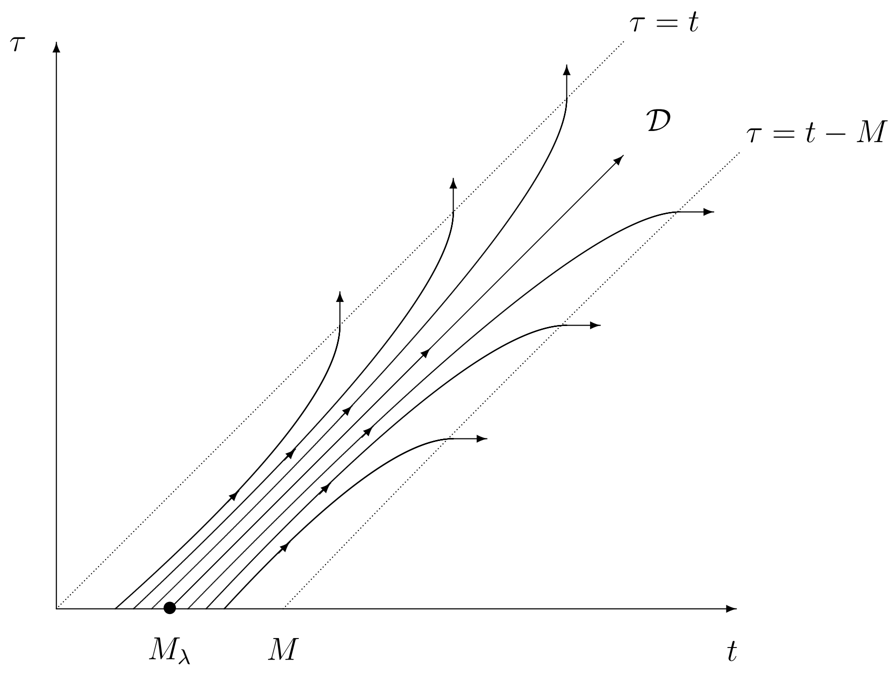

Figure 1 The unique equilibrium when the breakthrough rate is constant.

Comparative Statics As the payoff function L exhibits decreasing returns over $(0, M)$, the comparative statics analysis is straightforward: the higher the breakthrough rate λ, the shorter the equilibrium maturation delay $M_\lambda = (\dot{L}/L)^{-1}(\lambda)$. This reflects that if player i anticipates that player j keeps waiting the same constant amount of time before making a move, while experiencing breakthroughs at a higher rate, her risk of being preempted by player j increases. This in turn leads player i to behave more cautiously by reducing her own maturation delay. When λ is close to zero, M_λ is close to the optimal maturation delay $M = M_0$, reflecting low preemption risk. When λ tends to infinity, M_λ goes to zero, reflecting extreme preemption risk.

The impact of competition on payoffs is also easy to assess in the constant-breakthrough-rate case. Suppose that each player discounts future payoffs at rate r. When player i acts as a monopolist, her ex-ante payoff is

$$\mathbf{E}[\exp(-r\tilde{\tau}^i)L(M)] = \frac{\lambda}{\lambda + r}\, L(M). \tag{13}$$

By contrast, in the duopoly case, player i's ex-ante payoff is

$$\mathbf{E}[\exp(-r\tilde{\tau}^i)L(M_\lambda)\,|\,\tilde{\tau}^i < \tilde{\tau}^j] = \frac{\lambda}{2\lambda + r}\, L(M_\lambda). \tag{14}$$

As is clear from (13)–(14), there are two reasons why competition reduces players' rents relative to the monopoly situation: first, because of the risk of being preempted, and second, because the payoff conditional on not being preempted is lower than in the monopoly case. It should be noted that whereas an increase in the breakthrough rate λ always has a positive impact on a monopolist's ex-ante payoff, it has an ambiguous impact on a duopolist's ex-ante payoff: an increase in λ improves the efficiency of the research process (though, taking into account preemption risk, by a factor that is roughly one-half of that in the monopoly case), but it also reduces the payoff from preempting one's opponent. When λ is large enough, the second effect dominates, and all rents are fully dissipated in the limit.

Parametric Examples To get an order of magnitude for the decline in maturation delays and payoffs resulting from competition, let us consider the specification (1) of Example 1, for which

$$M_\lambda = \frac{1}{\mu}\,\ln\!\left(\frac{\lambda + r}{\lambda + r - \mu}\right). \tag{15}$$

The comparison with the optimal maturation delay M given by (2) is immediate: competition raises the effective discount factor from r to $\lambda + r$. It is worth noting here is that it may be more realistic to assume in practice that λ is large relative to r and μ. If this is the case, it follows from (15) that $M_\lambda \simeq 1/\lambda$. Suppose for instance that the annual discount rate r is 10% and that the annual payoff growth rate μ is 5%. Then the optimal maturation delay M is 13.86 years. Now, let us assume that on average, it takes players one year to experience a breakthrough, so that $\lambda = 1$. Then, in line with the above approximation, the equilibrium maturation delay drops to 0.93 years. Given these numbers, a duopolist's ex-ante payoff (14) amounts to only 9.1% of the monopolist's ex-ante payoff (13).

Similar results hold for the specification (3) of Example 2, for which

$$M_\lambda = \frac{r}{\lambda + r} M. \tag{16}$$

With the above parametrization, the equilibrium maturation delay under a duopoly decreases by 90.9% relative to the monopoly benchmark, while the total delay from breakthrough to publication increases by 148.8%.

3.3 Existence and Uniqueness of Equilibrium

We now turn to the general model of Section 2. We start with an existence result.

Proposition 1 *An equilibrium exists.*

According to the discussion in Section 3.1, we only need to show that for a carefully chosen initial condition, there exists a solution to the ODE (9) that remains in $\operatorname{Int}\mathcal{D}$. The proof of this fact relies on a connectedness argument, that is, a form of intermediate value theorem. Specifically, imagine that any solution to the ODE (9), indexed by its initial condition $(t, \tau) = (\sigma_0, 0)$ for $\sigma_0 \in [0, M]$, eventually leaves $\mathcal{D}$. Then, because the vector field induced by (10) is continuous, we would have a continuous and onto mapping from the space of initial conditions $[0, M]$ to the space consisting of the lower and upper boundaries of $\mathcal{D}$. This, however, is impossible, as the former space is connected, whereas the latter space has two connected components. As a result, there must exist a solution to the ODE (9) with an appropriate initial condition that never leaves $\mathcal{D}$. By construction, any such solution corresponds to a symmetric equilibrium.

As observed in Section 3.1, our game satisfies the single-crossing condition for games of incomplete information. One may, therefore, have good reasons to think that a more direct approach to equilibrium existence would exploit this monotonicity structure, along the lines of Athey (2001). The short answer is that Athey's (2001) existence results do not apply in our setup, because neither action spaces nor type spaces are compact. As for the former, it could seem that a simple fix would be to redefine player i's actions to be maturation delays $m^i \in [0, M]$. This, however, would overlook two difficulties. (1) The first difficulty is technical, namely that type spaces would still remain noncompact. One could of course avoid this problem by assuming that breakthroughs must take place before some fixed and common knowledge time T. This assumption, however, would be somehow artificial. Moreover, an interesting feature of the model is that in equilibrium, the behavior of a player's late types has an indirect effect on the behavior of her opponent's early types, as in overlapping-generation models: although type τ^i of player i is not directly concerned by the behavior of types $\tau^j > \tau^i + M$ of player j, for type τ^i makes a move in equilibrium before these types even have an opportunity to do so, yet the behavior of any such type τ^j is relevant for types $\tau^{i\prime} < \tau^j - M$ of player i and thus, by mutual contagion, ultimately for type τ^i as well. (2) The second difficulty is more substantial, namely that player i's payoff function need not satisfy single crossing of incremental returns in (m^i, τ^i). The monotonicity structure required to apply Athey's (2001) results is then lost. Intuitively, this reflects the fact that variations in the breakthrough rate may make it profitable for a player who experiences a breakthrough relatively late to let her idea mature less than if she had experienced this breakthrough earlier on. In Proposition 3, we provide conditions on the breakthrough rate that guarantee that the equilibrium maturation delay varies monotonically with the breakthrough time. However, one can think of situations that do not exhibit this pattern, and in which maturation delays fluctuate over time in equilibrium. Proposition 1 covers these cases as well.[13]

The proof of Proposition 1 shows that there exists a compact interval $\Sigma_0 \subset (0, M)$ such that for each $\sigma_0 \in \Sigma_0$, the solution to the ODE (9) starting at $(t, \tau) = (\sigma_0, 0)$ remains in $\operatorname{Int} \mathcal{D}$, and thus corresponds to a symmetric equilibrium. These solutions are ordered in a natural way, with a higher value of σ_0 corresponding to uniformly higher maturation delays. In the constant-breakthrough-rate case, Σ_0 is reduced to a point, so that there exists a unique symmetric equilibrium and hence, according to Lemma 3, a unique equilibrium. The following proposition shows that this is always the case.

Proposition 2 *The equilibrium is unique.*

The logic of Proposition 2 differs from that of the uniqueness results derived by Fudenberg and Tirole (1986) or Décamps and Mariotti (2004) in the context of war-of-attrition games with incomplete information. In such games, the potential nonuniqueness problem stems from the fact that whereas the initial condition is the same in any equilibrium, the differential equation describing the equilibrium typically fails to satisfy a Lipschitz condition at time zero given this initial condition (Riley (1980)). Uniqueness of equilibrium is established by imposing a condition on players' preferences that acts as a boundary condition at infinity.

[13]Proposition 1 is very similar in content to the existence theorem for *antifunnels* defined by outward-pointing vector fields, see Hartman (1964, Chapter X, Theorem 2.1) or Hubbard and West (1991, Theorem 4.7.3). An important difference, however, is that the function f in (10) does not satisfy a global Lipschitz condition over $\mathcal{D}$, so that a specific argument is needed.

In Fudenberg and Tirole's (1986) model of exit, this boundary condition states that each firm's cost may be so low that staying in forever is a dominant strategy. In Décamps and Mariotti's (2004) model of investment, this boundary condition states that each firm's cost may be so high that refraining from investing is a dominant strategy even if the investment project is known to generate high payoffs.

By contrast, in the present model, the potential nonuniqueness problem stems from the fact that a priori, multiple initial conditions may be consistent with equilibrium, because there is no terminal condition to pin down the behavior of late types. In that respect, things would be different if breakthroughs were to take place with probability 1 before some fixed and common knowledge time T. Then types τ close to T would face extreme preemption risk: conditional on no player having moved by time τ, each would expect her opponent to make a move in the next instant. As a result, the value of waiting for these types would vanish, and their equilibrium maturation delays would converge to zero.[14] When the breakthrough distribution has an unbounded support, however, such an ad-hoc terminal condition is not available, which makes the uniqueness of equilibrium more surprising. Indeed, player i's decision to uniformly increase her moving time intuitively increases the value of waiting for player j. Given these strategic complementarities, it would be natural to think that multiple equilibria may arise in our model.

The intuition for our uniqueness result is that two solutions to the ODE (9), which correspond to different initial conditions, are not only naturally ordered, but also tend to drift apart because the vector field induced by (10) is monotone in the moving time dimension. Therefore, if there were two symmetric equilibria with $\sigma_1 > \sigma_2$, one also would have $\dot{\sigma}_1 > \dot{\sigma}_2$: the equilibrium in which players fear less preemption also would be the one in which maturation delays respond more to breakthrough times at the margin. To get a contradiction, we need to establish that the rate at which σ_1 and σ_2 drift apart is sufficiently large: in that case, the distance between σ_1 and σ_2 cannot asymptotically remain bounded, as is required by the equilibrium condition that $\phi_1 \equiv \sigma_1^{-1}$ and $\phi_2 \equiv \sigma_2^{-1}$ must never leave $\mathcal{D}$.[15] Specifically, we prove that the rate at which the candidate equilibria σ_1 and σ_2 drift apart is at least as large as the breakthrough rate $\dot{F}/(1-F)$, that is, for each $\overline{\tau} \geq 0$,

$$\sigma_1(\overline{\tau}) - \sigma_2(\overline{\tau}) \geq [\sigma_1(0) - \sigma_2(0)] \exp\left(\int_0^{\overline{\tau}} \frac{\dot{F}(\tau)}{1-F(\tau)}\, d\tau\right).$$

[14]See Lambrecht and Perraudin (2003) and Anderson, Friedman, and Oprea (2010) for a related boundary condition in a real option context.

[15]This method of proof builds on uniqueness results for antifunnels (Hubbard and West (1991)).

Now, for any (absolutely continuous) distribution, the breakthrough rate $\dot{F}/(1-F)$ is not integrable over $[0,\infty)$, even though it may tend to zero at infinity. As a result, even arbitrarily small differences between initial maturation delays $\sigma_1(0)$ and $\sigma_2(0)$ would asymptotically translate into arbitrarily large differences in maturation delays, which is inconsistent with equilibrium behavior in the limit.

A key feature of our model is thus that the unique possible value for the initial condition $\sigma(0)$ in equilibrium is pinned down by the players' behavior as their breakthrough times become arbitrarily large—more precisely, by the weak rationality requirement that their equilibrium maturation delays can never exceed the optimal maturation delay M. It should be noted that this uniqueness result holds independently of the breakthrough distribution.

3.4 Monotone Breakthrough Rate

In Section 3.2 we investigated the case when the breakthrough rate $\dot{F}/(1-F)$ remains constant. In this stationary environment, the competitive pressure to which players are exposed does not vary over time. As a result, the equilibrium maturation delay remains constant as well. Our model, however, is flexible enough to depict situations in which the breakthrough rate varies over time. For instance, the breakthrough rate may increase over time as a result of exogenous technological progress or general human capital accumulation. A case in point is scientific progress, where advances in a given field gradually spread to sometimes distant and a priori unrelated fields on which they shed a new light, making it more likely that competing researchers make progress there as well. (A similar mechanism is likely to be at play in the case of technological innovation.) The reverse case when breakthroughs occur at a decreasing rate also is interesting. In the case of academic competition, this scenario may arise in fields that gradually go out of fashion, for instance because some hotter topic summons up most of the researchers' time and energy.

To study the impact of such variations on the evolution of equilibrium maturation delays, we introduce the following terminology. We say that F has a *differentiably strictly increasing (decreasing) breakthrough rate* if the time-derivative of $\dot{F}/(1-F)$ is everywhere well defined and positive (negative); more generally, we say that F has a *differentiably strictly monotone breakthrough rate* if it has a differentiably strictly increasing or decreasing breakthrough rate. Then the following result holds.

Proposition 3 *If F has a differentiably strictly increasing (decreasing) breakthrough rate, then the equilibrium maturation delay $\sigma(\tau)-\tau$ is strictly decreasing (increasing) with respect to the breakthrough time τ.*

The intuition of this result is as follows. If breakthroughs become more frequent over time, a player who experiences a breakthrough relatively late is more threatened by preemption, and thus has less incentives to wait and let her idea mature than if she had experienced this breakthrough earlier on: ideas developed later on are less mature than ideas developed earlier on. Our model thus predicts that if researchers become increasingly innovative over time, their more path-breaking and accomplished contributions tend to take place relatively early. The case when breakthroughs become less frequent over time leads to the opposite conclusion: as time goes by, the competitive pressure is gradually alleviated, which leads to longer maturation delays and rarer, but on average more accomplished works or higher quality innovations.

It follows from Proposition 3 that if F has a differentiably strictly monotone breakthrough rate, then the equilibrium maturation delay $\sigma(\tau) - \tau$ for a time-τ idea tends to a well-defined limit as τ goes to infinity. For each $\lambda \geq 0$, define M_λ as in Section 3.2. Then two scenarios can arise.

Corollary 1 *If F has a differentiably strictly monotone breakthrough rate, then*

(i) If the breakthrough rate of F eventually converges to a finite value λ_∞, the equilibrium maturation delay eventually converges to M_{λ_∞}.

(ii) If the breakthrough rate of F eventually diverges to infinity, the equilibrium maturation delay eventually converges to zero.

In case (i), two situations are possible. When $\lambda_\infty > 0$, the asymptotic maturation delay coincides with the equilibrium maturation delay M_{λ_∞} for a constant breakthrough rate λ_∞. When $\lambda_\infty = 0$, the equilibrium maturation delay is asymptotically equal to the optimal maturation delay $M_0 = M$, and asymptotically no rents are dissipated, as in Reinganum (1981). By contrast, in case (ii), the breakthrough rate of F grows without bounds. Then the equilibrium maturation delay eventually converges to zero, and asymptotically all rents are dissipated, as in Fudenberg and Tirole (1985). This intuitively corresponds to a situation in which an innovation eventually becomes so obvious that the first researcher who discovers it almost immediately makes it public, for fear of being preempted in the next instant. Note that this case is qualitatively similar to the case when breakthroughs take place before some fixed and common knowledge time T.

A direct inspection of the proofs of Proposition 3 and Corollary 1 reveals that these two results more generally extend to distributions with eventually differentiably strictly monotone breakthrough rates. Thus, for instance, if F eventually has a differentiably strictly increasing breakthrough rate, then, beyond a finite time, ideas developed later on are less mature than ideas developed early on. Corollary 1(i) can be generalized to the case when the breakthrough rate of F eventually converges to a finite and positive value, even if this convergence is not monotonic.

Corollary 2 *If the breakthrough rate $\dot{F}/(1-F)$ eventually converges to a finite and positive value λ_∞, then the equilibrium maturation delay eventually converges to M_{λ_∞}.*

Corollary 2 generalizes the constant-breakthrough-rate case in a natural way, as it only requires the breakthrough rate to be asymptotically constant.

3.5 Comparative Statics

In this section, we explore how the equilibrium reacts in response to changes in the underlying parameters of the model.

Changes in the Breakthrough Distribution In the constant-breakthrough-rate case, an increase in the breakthrough rate λ leads to a decrease in the equilibrium maturation delay M_λ, reflecting increased competition. This insight can be generalized as follows. Specifically, let F_1 and F_2 be two breakthrough distributions with densities $\dot{F}_1$ and $\dot{F}_2$. In analogy with the first-price-auction literature (Lebrun (1998), Maskin and Riley (2000), Arozamena and Cantillon (2004)), we compare the corresponding equilibrium maturation delays under the assumption that F_1 is a *distributional upgrade* of F_2, that is,

$$\frac{\dot{F}_1}{1-F_1} > \frac{\dot{F}_2}{1-F_2} \tag{17}$$

over $(0,\infty)$. Inequality (17) states that conditional on any minimum level of breakthrough time, F_2 is more likely to yield a larger breakthrough time than F_1. In particular, F_2 dominates F_1 in the sense of first-order stochastic dominance: breakthroughs tend to occur later under F_2 than under F_1. Then the following result holds.

Corollary 3 *If F_1 is a distributional upgrade of F_2, then the equilibrium maturation delay $\sigma_1(\tau)-\tau$ under F_1 is strictly shorter than the equilibrium maturation delay $\sigma_2(\tau)-\tau$ under F_2 for any breakthrough time τ.*

Corollary 3 is very much in line with the uniqueness result stated in Proposition 2. Indeed, a key step in the proof of Corollary 3 consists in establishing a single-crossing property, namely that if one had $\sigma_1(\tau_0) - \tau_0 \geq \sigma_2(\tau_0) - \tau_0$ for some $\tau_0 > 0$, one also would have $\sigma_1(\tau) - \tau > \sigma_2(\tau) - \tau$ for all $\tau > \tau_0$. It would then be possible to construct a solution to the ODE (9) for $F = F_2$ that would strictly lie between σ_2 and σ_1 over (τ_0, ∞). Because neither σ_1 nor σ_2 ever leave $\mathcal{D}$, the same would be true for this solution. Therefore, the latter would correspond to a symmetric equilibrium under F_2 distinct from σ_2, contradicting our uniqueness result.

Corollary 3 predicts that if researchers become more innovative, in that they tend to experience breakthroughs earlier, then the resulting increased fear of preemption leads to an equilibrium with shorter maturation delays, and thus less accomplished works or lower quality innovations, conditional on any value of the breakthrough time.

Because breakthrough times are hard to observe in practice, if only because researchers have strong incentives to keep them secret, it also is worthwhile to investigate how the *average* maturation delay for a successful idea is affected by a change in the breakthrough distribution. Formally, let $F_{(1/2)} \equiv 1 - (1 - F)^2$ be the distribution of the first-order statistic for a sample of two independent breakthrough times. Then the average maturation delay for the first successful idea is

$$\int_0^\infty [\sigma(\tau) - \tau] \, dF_{(1/2)}(\tau) = 2 \int_0^\infty [\sigma(\tau) - \tau][1 - F(\tau)]\dot{F}(\tau) \, d\tau,$$

where σ is the equilibrium strategy given a breakthrough distribution F. The following result holds.

Corollary 4 *If F_1 is a distributional upgrade of F_2 and F_1 or F_2 has a differentiably strictly decreasing or constant breakthrough rate, then the average equilibrium maturation delay under F_1 is strictly shorter than the average equilibrium maturation delay under F_2.*

The logic of this result is as follows. According to Corollary 3, conditional on any given value of the breakthrough time, the equilibrium maturation delay under F_1 is shorter than under F_2. However, breakthroughs occur at a faster rate under F_1 than under F_2. The total impact of a change in the breakthrough distribution on the average maturation delay is thus not a priori obvious and depends on how the equilibrium maturation delay varies under F_1 and F_2. Building on a simple stochastic-dominance argument, Corollary 4 gives a sufficient condition under which researchers let their ideas mature less on average when they become more innovative. This condition ensures that at least one of the equilibrium maturation delays under F_1 and F_2 is nondecreasing in the breakthrough time. It then follows from the comparison result in Corollary 3 that for any fixed breakthrough time τ_2, the equilibrium maturation delay under F_1 at any breakthrough time $\tau_1 \leq \tau_2$ is strictly shorter than the equilibrium maturation delay under F_2 at τ_2: intuitively, the gap between the maturation delays under the two distributions is high enough for a relatively early breakthrough. Thus, as breakthroughs tend to occur earlier under F_1 than under F_2, the average maturation delay is shorter under F_1 than under F_2. By contrast, if the equilibrium maturation delays under F_1 and F_2 were both decreasing in the breakthrough time, two competing effects would be at play: whereas the logic of Corollary 3 would remain intact for any given value of the breakthrough time, breakthroughs under F_1 would tend to occur when maturation delays are relatively long, whereas breakthroughs under F_2 would tend to occur when maturation delays are relatively short. Thus in such a case the total impact of a change in the breakthrough distribution on the average maturation delay is ambiguous.

Changes in the Payoff Function We now consider the impact of a change in the payoff function $L : [0, \infty) \to [0, \infty)$. To do so, we first define an order on such functions. According to (9)–(10), what we need is a criterion to compare the rate at which the payoffs from different projects grow as these projects mature. We adapt to our context the fear-of-ruin index introduced by Aumann and Kurz (1977) and more recently studied by Foncel and Treich (2005). Formally, we say that L_1 is *growth-dominated* by L_2 if they both have positive returns on $[0, M)$ and

$$\frac{\dot{L}_1}{L_1} < \frac{\dot{L}_2}{L_2} \tag{18}$$

over $(0, M)$. That is, L_1 grows at a lower rate than L_2 over the range where both have positive returns. It follows from Foncel and Treich (2005, Proposition 1) that this is equivalent to assuming that

$$L_1 = h \circ L_2, \tag{19}$$

for some differentiable and strictly increasing function $h : [0, \infty) \to [0, \infty)$ such that $h(0) = 0$ and $h(l)/l$ is strictly decreasing in l over $(0, L_2(M))$, with a negative derivative over this interval. Observe that (19) implies that L_2 increases exactly where L_1 increases. Thus if L_2 reaches its global maximum at M, so does L_1. For instance, L_1 is growth-dominated by L_2 if $\dot{L}_1/L_1 = h\dot{L}_2/L_2$ so that $L_1 \propto L_2^h$ for some $h \in (0, 1)$.

It is worth noting that if L_1 is growth-dominated by L_2, this does not mean that L_2 uniformly generates higher payoffs than L_1. Actually, quite the contrary may be true, because both functions vanish at the origin, which implies that for L_2 to growth-dominate L_1, it may be necessary for L_2 to be lower than L_1 near the origin.[16] Indeed, as is clear from (9)–(10), what matters for incentives is not the levels of payoffs, but the rate at which they grow. The following result holds.

Corollary 5 *If L_1 is growth-dominated by L_2, then the equilibrium maturation delay $\sigma_1(\tau) - \tau$ under L_1 is strictly shorter than the equilibrium maturation delay $\sigma_2(\tau) - \tau$ under L_2 for any breakthrough time τ.*

Again, the proof of Corollary 5 relies on similar ideas as the proof of the uniqueness result stated in Proposition 2. Corollary 5 predicts that if the payoff from each player's project grows faster, this counteracts the fear of preemption and leads to more accomplished works or higher quality innovations, conditional on any value of the breakthrough time: intuitively, waiting an additional unit of time is more valuable, while the preemption pressure remains the same. A similar result holds when one averages over breakthrough times, as the distribution of breakthrough times is not affected by a change in the payoff function. It should be noted that this is in spite of the fact that growth-domination as we define it does not affect the optimal maturation delay M for a monopolist.

Example 2 nicely illustrates the logic of Corollary 5. Indeed, it directly follows from (3) that if $r_1 < r_2 < 1/M$, then the corresponding payoff functions L_1 and L_2 are such that L_1 is growth-dominated by L_2 as $L_1 = L_2^{r_1/r_2}$. Thus, in this example, if competition takes place between relatively impatient researchers, such as assistant professors on a tenure track, it leads to longer maturation delays and higher quality initial submissions than if competition takes place between relatively patient researchers, who for instance already have tenure. The intuition is that more impatient researchers are more concerned by having their paper stuck in the publication process than less impatient researchers. As a result, they are more willing to take risks by letting their ideas mature more than the latter.[17]

Innovation versus Growth Potential According to Corollaries 3 and 5, changes that affect the rate at which researchers experience breakthroughs and changes that affect the rate at which their payoffs grow once they have experienced a breakthrough have contrasted effects on their equilibrium maturation strategies: if researchers become more innovative, this tends to speed up the maturation process, leading to less accomplished works, whereas

[16] For instance, if $L_1 = L_2^h$ for some $h \in (0,1)$, then $\dot{L}_1 = hL_2^{h-1}\dot{L}_2 > \dot{L}_2$ in the neighborhood of zero, see the discussion of Example 2 below. A less extreme example arises if $\dot{L}_1/L_1 = \varphi \circ (\dot{L}_2/L_2)$ where $\varphi(\gamma) = \gamma + \exp(-\gamma) - 1$ for $\gamma > 0$. Then $\dot{L}_1/L_1 = \dot{L}_2/L_2 - 1 + o(1)$ so that $L_1 = \exp(C + o(1))L_2 > L_2$ in the neighborhood of zero whenever the constant of integration C is positive.

[17] Of course, another reason why young researchers may take more time to complete their research projects is their relative lack of experience.

if researchers work on projects with a higher payoff growth potential, this tends to slow down the maturation process, leading to more accomplished works. Thus there is a contrast between the pure research phase during which researchers experience breakthroughs, and the development phase during which they let their ideas mature. If the social value of research is in part measured by the yardstick of ideas' maturation, this suggests that researchers should be induced to engage in projects with a high growth potential rather than in projects in which breakthroughs are more frequent. As noted above, however, the former need not be the ones that generate the highest private rents for the researchers.

3.6 Multiple Symmetric Players

It is straightforward to extend our analysis to the case of $N > 2$ symmetric players whose breakthrough times are independently distributed according to the distribution function F. The only adjustment to the model consists in specifying players' payoffs when some of them simultaneously make a move; without loss of generality, assume that player i's payoff from making a move at time t^i along with J other players is $\alpha^J L(t^i - \tau^i)$ for some $\alpha \in [0, 1]$. We shall focus on symmetric pure-strategy equilibria, in which all players play the same strategy σ. In line with Lemma 1(i), it is easy to see that $0 < \sigma(\tau) - \tau \leq M$ for all $\tau \geq 0$. Assume that σ is strictly increasing, so that no tie problem can arise, and that σ has a differentiable inverse. The problem faced by type τ^i of player i can now be written as

$$\max_{t^i \in [\tau^i, \infty)} \left\{ \prod_{j \neq i} \mathbf{P}[\sigma(\tilde{\tau}^j) > t^i] L(t^i - \tau^i) \right\},$$

or, letting $\phi \equiv \sigma^{-1}$, and taking advantage of the independence between players' types,

$$\max_{t^i \in [\tau^i, \infty)} \{[1 - F(\phi(t^i))]^{N-1} L(t^i - \tau^i)\}. \tag{20}$$

Because a zero maturation delay is inconsistent with equilibrium, the solution to problem (20) must be interior. The first-order condition is

$$[1 - F(\phi(t^i))]\dot{L}(t^i - \tau^i) = (N-1)\dot{F}(\phi(t^i))\dot{\phi}(t^i)L(t^i - \tau^i),$$

the interpretation of which is similar to that of (7). In a symmetric equilibrium, this first-order condition holds for $\tau^i = \phi(t^i)$, leading to the nonautonomous ODE

$$\dot{\phi}(t) = \frac{1}{N-1} f(t, \phi(t)), \quad t \geq \sigma(0), \tag{21}$$

where f is given by (10). The only difference between (9) and (21) is thus that in the latter case, the breakthrough rate $\dot{F}/(1 - F)$ is multiplied by $N - 1 \geq 2$, as if each player believed she were facing a single opponent whose breakthrough time would be distributed according to the distribution of the first-order statistic for a sample of $N-1$ independent breakthrough times, $F_{(1/N-1)} \equiv 1-(1-F)^{N-1}$. The analysis of equilibrium remains otherwise the same. In particular, as increasing the number of players essentially amounts to increasing the breakthrough rate in a multiplicative way, it follows from Corollary 3 that more competition leads to an increased fear of preemption and thus shorter maturation delays conditional on any value of the breakthrough time.

4 Asymmetric Players: The Hare and the Tortoise

Let us return to the two-player case, and suppose now that players have constant, but different breakthrough rates $\lambda^a > \lambda^b$. Thus a is a relatively more innovative researcher (the hare), whereas b is a relatively less innovative researcher (the tortoise). The description of the game otherwise remains the same as in the symmetric model of Section 2. In particular, payoffs and equilibria are still defined by (4)–(5), the only difference being that $\tilde{\tau}^a$ and $\tilde{\tau}^b$ now have different distributions. As explained in the appendix, Lemma 1 carries over to this asymmetric context. Our proof of Lemma 2, however, does not extend to the case of asymmetric players; in particular, we have not been able to rule out discontinuous equilibria. We shall nevertheless disregard this possibility, and hereafter focus on continuous equilibria. In any such equilibrium (σ^a, σ^b), $\phi^a \equiv (\sigma^a)^{-1}$ and $\phi^b \equiv (\sigma^b)^{-1}$ solve the following system of nonautonomous ODEs:

$$\dot{\phi}^j(t) = \frac{1}{\lambda^j} \frac{\dot{L}(t-\phi^i(t))}{L(t-\phi^i(t))}, \quad t \geq \sigma(0), \quad i = a, b. \tag{22}$$

As in the symmetric case, the initial condition $\sigma(0) = \sigma^a(0) = \sigma^b(0)$ of that system must be chosen in such a way that neither ϕ^a nor ϕ^b leave $\mathcal{D}$. One first has the following result.

Lemma 4 *In any continuous equilibrium, the hare's equilibrium maturation delay $\sigma^a(\tau) - \tau$ is strictly longer than the tortoise's equilibrium maturation delay $\sigma^b(\tau) - \tau$ for any breakthrough time $\tau > 0$.*

Lemma 4 reflects that the hare is more immune to preemption risk than the tortoise is, because the latter experiences fewer breakthroughs per unit of time. Conditional on any value of the breakthrough time, she can thus afford letting her ideas mature longer than the tortoise can. Observe also that if the tortoise and the hare experience a breakthrough simultaneously or sufficiently close in time, it is the tortoise who will make a move first, thereby preempting the hare, because she will let her idea mature less than the latter.

Recall from Section 3.2 that if player i anticipated player j to make a move at a constant rate λ^j, it would be optimal for her to wait an amount of time M_{λ^j} before making a move. Of course, this cannot be quite true in equilibrium because one must have $\sigma^a(0) = \sigma^b(0) = \sigma(0)$. Nevertheless, the delays M_{λ^a} and M_{λ^b} play an important role in the analysis, as we shall now see. First, we can use them to narrow down the set of possible values for $\sigma(0)$.

Lemma 5 *In any continuous equilibrium,*

$$\sigma^a(\tau) - \tau > M_{\lambda^a} \text{ and } \sigma^b(\tau) - \tau < M_{\lambda^b}, \quad \tau \geq 0, \tag{23}$$

so that in particular

$$M_{\lambda^a} < \sigma(0) < M_{\lambda^b}. \tag{24}$$

According to (23), the hare's equilibrium maturation delay is strictly longer than M_{λ^a}, and the tortoise's equilibrium maturation delay is strictly shorter than M_{λ^b}. In light of our analysis of the constant-breakthrough-rate case in Section 3.2, this means that the hare tends to behave less cautiously, and the tortoise more cautiously, than if they were each facing an opponent of equal strength.

Lemma 5 and its proof lead to the natural conjecture that ϕ^a and ϕ^b never leave the restricted domain

$$\mathcal{D}^{a,b} \equiv \{(t, \tau) : M_{\lambda^a} \leq \tau + M_{\lambda^a} < t < \tau + M_{\lambda^b}\} \subset \mathcal{D}.$$

This intuition is indeed correct and underlies Proposition 4 below. To see that, it is helpful to rewrite the system (22) using as new variables the maturation delays $\mu^a(t) = t - \phi^a(t)$ and $\mu^b(t) = t - \phi^b(t)$. This yields

$$\dot{\mu}^j(t) = 1 - \frac{1}{\lambda^j} \frac{\dot{L}(\mu^i(t))}{L(\mu^i(t))}, \quad t \geq \sigma(0), \quad i = a, b. \tag{25}$$

Compared to (22), this change of variables brings two simplifications. The first is that we are now dealing with an *autonomous* system: thanks to the constant-breakthrough-rate assumption, time does not show up as an independent variable in (25). This means in particular that we can, modulo a time translation of length $\sigma(0)$, write down the system (25) for any time $t \geq 0$.[18] Second, proving that an equilibrium exists now amounts to proving that there exists a solution (μ^a, μ^b) to (25) with initial condition $\mu^a(0) = \mu^b(0)$ in $(M_{\lambda^a}, M_{\lambda^b})$ that is entirely contained in a *bounded* set, namely, $(0, M] \times (0, M]$. Taking advantage of these two simplifications, one can derive the following result.

[18] Formally, this consists in redefining maturation delays as $\mu_0^i(t) = \mu^i(t+\sigma(0))$. Without risk of confusion, we identify 'μ^i' and 'μ_0^i' so as to simplify notation.

Proposition 4 *A continuous equilibrium exists. Moreover, in any such equilibrium,*

(i) $M_{\lambda^a} < \sigma^b(\tau) - \tau < \sigma^a(\tau) - \tau < M_{\lambda^b}$ *for all* $\tau > 0$.

(ii) The hare's equilibrium maturation delay $\sigma^a(\tau) - \tau$ *is strictly increasing with respect to the breakthrough time* τ, *and* $\lim_{\tau\to\infty} \sigma^a(\tau) - \tau = M_{\lambda^b}$.

(iii) The tortoise's equilibrium maturation delay $\sigma^b(\tau) - \tau$ *is strictly decreasing with respect to the breakthrough time* τ, *and* $\lim_{\tau\to\infty} \sigma^b(\tau) - \tau = M_{\lambda^a}$.

In line with the proof of Proposition 1, we use a simple connectedness argument to establish Proposition 4. The inequality in (i) generalizes (24) by showing that equilibrium maturation delays always remain in the interval $(M_{\lambda^a}, M_{\lambda^b})$. Given (25), the monotonicity properties of the hare's and the tortoise's equilibrium maturation delays stated in (ii)–(iii) are then a direct consequence of this inequality. Proposition 4 has several implications.

First, the gap between the hare's and the tortoise's equilibrium maturation delays, given a common breakthrough time, increases over time and eventually converge to $M_{\lambda^b} - M_{\lambda^a}$. Thus, if a breakthrough occurs relatively early, then the corresponding maturation delay is less sensitive to the identity of the player who experiences it than if it were to occur later on. As a result, there is less heterogeneity in the quality of research outputs when breakthroughs occur relatively early.

Second, the hare always chooses a longer maturation delay than the tortoise, no matter when they experience breakthroughs. This is a stronger claim than Lemma 4, which stated a similar result conditional on a given value of the breakthrough time. Intuitively, this is because the hare features a flight to quality when she experiences breakthroughs later on, whereas the reverse holds true for the tortoise. In the case of academic competition, this leads to the prediction that more innovative researchers, who tend to experience breakthroughs relatively early, always produce more accomplished works than less innovative researchers when they are not preempted by the latter. Thus, within a group of competing researchers, speed of discovery and maturation of ideas should be positively correlated. In the present model, this effect only results from the competition among researchers; indeed, if they were not threatened by preemption, differently apt researchers would wait the same amount of time M before making a move after experiencing a breakthrough, leading to identical high-quality research outputs. Note that this correlation would be reversed if it were computed across noncompeting groups of researchers with different aptitudes: indeed, groups with more innovative researchers and thus higher speed of discovery tend to be more competitive, leading to shorter maturation delays and lower-quality research outputs.

Third, because the hare's and the tortoise's equilibrium maturation delays respectively converge to M_{λ^b} and M_{λ^a} as their breakthrough times get large, the nonstationarity of the equilibrium induced by the initial condition $\sigma^a(0) = \sigma^b(0)$ asymptotically vanishes. Indeed, each player i asymptotically behaves as if her opponent let her ideas approximatively mature an amount of time M_{λ^i}, thus making a move at an approximatively constant rate λ^j. This in turn induces player i to let her ideas approximatively mature an amount of time M_{λ^j}, thus making a move at an approximatively constant rate λ^i. The fact that any continuous equilibrium must exhibit this asymptotic behavior is a strong hint that there exists at most one such equilibrium. This intuition is correct, as the following result shows.

Proposition 5 *There exists a unique continuous equilibrium.*

This equilibrium is illustrated on Figure 2, which also represents the phase portrait associated to (25).

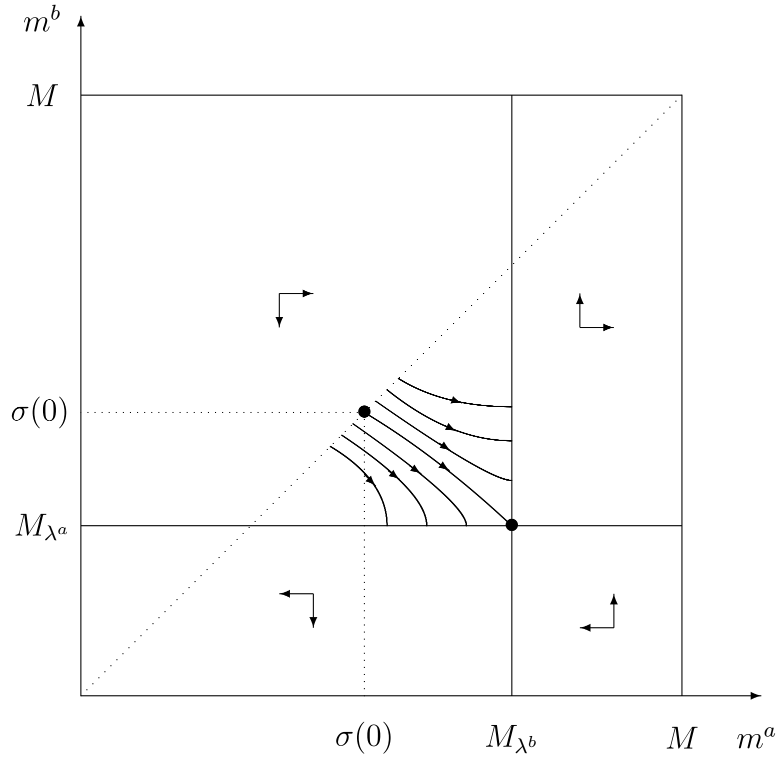

Figure 2 The unique continuous equilibrium when the breakthrough rates are constant.

The equilibrium corresponds to the unique trajectory of this system that converges to $(M_{\lambda^b}, M_{\lambda^a})$, which is a critical point of (25) as

$$1 - \frac{1}{\lambda^j} \frac{\dot{L}(M_{\lambda^j})}{L(M_{\lambda^j})} = 0$$

by definition of M_{λ^j}. This critical point is a saddle: the Jacobian associated to (25) at this point has two real eigenvalues of opposite signs, $\delta > 0$ and $-\delta < 0$. The equilibrium trajectory corresponds to the latter, and coincides with the upper branch of the stable manifold of (25) at $(M_{\lambda^b}, M_{\lambda^a})$. It is a consequence of a theorem of Hartman (1960) that if L is sufficiently regular, which we shall henceforth assume, then

$$\lim_{t\to\infty} \exp(\delta t)[(\mu^a(t), \mu^b(t)) - (M_{\lambda^b}, M_{\lambda^a})] = C\boldsymbol{\xi}_{-\delta}, \tag{26}$$

where C is a nonzero scalar constant and $\boldsymbol{\xi}_{-\delta}$ is a nonzero eigenvector associated to $-\delta$. Hence the equilibrium maturation delays $\mu^a(t)$ and $\mu^b(t)$ converge exponentially fast to their limit values M_{λ^b} and M_{λ^a}.

It is relatively straightforward to study how the equilibrium characterized in Propositions 4 and 5 reacts to changes in the breakthrough rates λ^a and λ^b. Suppose for instance that the hare's breakthrough rate increases, whereas the tortoise's breakthrough rate remains the same. Then, according to Proposition 4(ii), the tortoise's equilibrium maturation delay asymptotically becomes shorter as M_{λ^a} is a decreasing function of λ^a. This comparative statics result actually holds for any value of the tortoise's breakthrough time, and not only for large values thereof, as the following result shows.

Corollary 6 *If the hare's breakthrough rate increases from λ^a_1 to $\lambda^a_2 > \lambda^a_1$, whereas the tortoise's breakthrough rate λ^b_1 remains the same, then the tortoise's equilibrium maturation delay $\sigma^b_2(\tau) - \tau$ under $(\lambda^a_2, \lambda^b_1)$ is strictly shorter than her equilibrium maturation delay $\sigma^b_1(\tau) - \tau$ under $(\lambda^a_1, \lambda^b_1)$ for any breakthrough time τ.*

A symmetric result holds when the tortoise's breakthrough rate is modified, keeping the hare's breakthrough rate constant. Corollary 6 has a natural interpretation: if the hare, who already enjoys an advantage over the tortoise, were to become even more innovative, then the tortoise's increased fear of being preempted would lead her to let her ideas mature less, whatever the time at which she would experience a breakthrough. Indeed, the direct effect of an increase in the hare's breakthrough rate is immediate upon writing (22) for $i = b$: formally, one has $(\dot{L}/L)(\mu^b) = \lambda^a(1 - \dot{\mu}^a)$ so that, if the hare's behavior as summarized by $\dot{\mu}^a$ is held fixed, an increase in λ^a triggers a downward shift of the function μ^b. Of course, the proof of Corollary 6 is slightly more convoluted, because the hare's behavior also varies when her own breakthrough rate increases. Yet the same conclusion is upheld: the mere fact of facing an increasingly challenging opponent unambiguously deteriorates the quality of the tortoise's research outputs.

As for the hare, becoming even more innovative is a mixed blessing. On the one hand, experiencing breakthroughs more often increases her competitive edge over the tortoise and, other things equal, would allow her to take more time to let her ideas mature. On the other hand, this direct effect is counterbalanced by the fact that, as we have seen, the tortoise reacts to an increase in the hare's breakthrough rate by letting her own ideas mature less, which makes her a tougher opponent from the hare's perspective. Which of these effects dominates is a priori unclear, and depends on when the hare experiences a breakthrough. If this happens relatively early—a more likely scenario in case she becomes more innovative—the second or strategic effect dominates. This is because, according to Lemma 1(iii), $\sigma^a(0)$ must be equal to $\sigma^b(0)$ in equilibrium, whereas, according to Corollary 6, $\sigma^b(0)$ decreases when the hare experiences breakthrough more often. Thus the fact that the tortoise behaves more cautiously in equilibrium compels the hare to do the same, at least when she experiences a breakthrough relatively early. However, unlike in Corollary 6, it is not possible to translate this local insight into a global comparative statics result. Indeed, for larger values of the hare's breakthrough time, an increase in her breakthrough rate has an ambiguous effect on her equilibrium maturation delay: a priori, it may increase or decrease. Both scenarios can arise in equilibrium, as the following result for the specifications (1) and (3) shows.

Corollary 7 *Let the function L be given by (1) (respectively (3)). Then, if $\lambda^a < \sqrt{r(r-\mu)}$ (respectively $\lambda^a < r$), a small increase in the hare's breakthrough rate λ^a shortens the hare's equilibrium maturation delay for large values of her breakthrough time, whereas the opposite is true if $\lambda^a > \sqrt{r(r-\mu)}$ (respectively $\lambda^a > r$).*

The intuition for this result is as follows. According to Proposition 4(iii), the tortoise's equilibrium maturation delay is close to M_{λ^a} when she experiences a late breakthrough. Now, when the payoff function L is given by (1) or (3), it is straightforward to see that the maturation delay M_{λ^a} is convex in λ^a. Therefore, the limit of the tortoise's maturation delay is less sensitive to an increase in the hare's breakthrough rate when the hare's initial breakthrough rate is high than when it is low. When the hare experiences a late breakthrough herself, she is thus less threatened by preemption at the margin in the former case than in the latter, and she is ready to let her idea mature more: in terms of our previous discussion, the direct effect of an increase in the hare's breakthrough rate asymptotically dominates the indirect effect that works through the modification of the tortoise's equilibrium behavior. This prediction is reversed if the hare's initial breakthrough is initially lower, for then an increase in it has a large impact on the tortoise's limit equilibrium behavior. It should however be observed that this second scenario is perhaps less likely to occur, as it may be more realistic to assume in practice that λ^a is large relative to $\sqrt{r(r-\mu)}$ or r.

5 Concluding Remarks

In this paper, we took a first pass at studying academic competition from a strategic point of view. The key feature of the environment we studied is that breakthroughs are secret, which leads to a nontrivial tradeoff between the gains from letting one's project mature and the risk of being preempted. In the symmetric two-researcher case, the unique equilibrium is described by a differential equation. The two forces that drive the equilibrium are the instantaneous rate at which researchers experience breakthroughs and the instantaneous rate at which their payoffs increase as a function of how long they let their projects mature. Comparative statics analyses show that variations in these two measures of a researcher's ability have contrasted effects on the quality of research outputs. Besides, how breakthrough rates evolve over time is crucial to understand how equilibrium maturation delays vary as a function of the breakthrough time. Finally, in asymmetric contests in which researchers have different innovative abilities, we found that speed of discovery and maturation of ideas are positively correlated in equilibrium.

Our model could be extended and enriched along several lines. A limitation of our approach is that breakthroughs are exogenous. This is because we chose to focus our analysis on the maturation of ideas rather than on the process through which ideas are discovered in the first place. Yet researchers devote a lot of time and effort to find new ideas. It would be interesting, in future work, to endogenize the rate at which researchers experience breakthroughs. Another limitation of our analysis is that all breakthroughs are of identical and a priori known quality. Alas, researchers too often spend time on a new idea only to discover that it is devoid of interest, or too hard to develop—indeed, the very fact that no one has ever written on a given topic is sometimes bad news about it. A richer model of idea maturation should take this feature into account.

We emphasized the deleterious impact that competition may have on the quality of research outputs when researchers fear they might lose priority. In doing so, we implicitly adopted the point of view of researchers engaged in such a preemption game, or that of an outside observer solely concerned with intellectual achievements per se. From a social point of view, however, it is unclear whether, even in the context of our model, competition among researchers is necessarily wasteful: whereas competition in our model unambiguously leads to less accomplished works than a monopolistic researcher would produce, it also speeds up the process through which scientific discoveries are made public. How institutions such as publication standards in academic journals or promotion requirements in universities should be designed and evolve to meet this tradeoff between the maturation and dissemination of ideas is a fascinating topic for future research.

Technical Appendix (For Online Publication)

Proof of Lemma 1. (i) Suppose first by way of contradiction that $\sigma^i(\tau^i) = \tau^i$ for some i and τ^i. Then, according to (4), type τ^i's equilibrium payoff is zero. Yet, according to (4) again, type τ^i could secure a payoff $\{\mathbf{P}[\sigma^j(\tilde{\tau}^j) > \tau^i + m] + \alpha\mathbf{P}[\sigma^j(\tilde{\tau}^j) = \tau^i + m]\}L(m)$ by waiting an amount of time $m > 0$. As $L(m) > 0$ for all $m > 0$, this implies that $\mathbf{P}[\sigma^j(\tilde{\tau}^j) > \tau^i + m] + \alpha\mathbf{P}[\sigma^j(\tilde{\tau}^j) = \tau^i + m] = 0$ and hence $\mathbf{P}[\sigma^j(\tilde{\tau}^j) > \tau^i + m] = 0$ for any such m. As a result, one must have $\mathbf{P}[\sigma^j(\tilde{\tau}^j) \leq \tau^i] = 1$, which is impossible as

$$\mathbf{P}[\sigma^j(\tilde{\tau}^j) \leq \tau^i] = \int_0^{\tau^i} 1_{\{\sigma^j(\tau^j) \leq \tau^i\}}\, dF(\tau^j) \leq F(\tau^i) < 1. \tag{27}$$

This contradiction establishes that $\sigma^i(\tau^i) - \tau^i > 0$ for all i and τ^i.

Suppose next by way of contradiction that $\sigma^i(\tau^i) > \tau^i + M$ for some i and τ^i. Then, as L is strictly decreasing over $[M, \infty)$, one has, for each $\varepsilon \in (0, \sigma^i(\tau^i) - \tau^i - M]$ such that $\mathbf{P}[\sigma^j(\tilde{\tau}^j) = \sigma^i(\tau^i) - \varepsilon] = 0$,

$$\begin{aligned}
V^i(\sigma^i(\tau^i), \tau^i, \sigma^j) &= \{\mathbf{P}[\sigma^j(\tilde{\tau}^j) > \sigma^i(\tau^i)] + \alpha\mathbf{P}[\sigma^j(\tilde{\tau}^j) = \sigma^i(\tau^i)]\}L(\sigma^i(\tau^i) - \tau^i) \\
&\leq \mathbf{P}[\sigma^j(\tilde{\tau}^j) \geq \sigma^i(\tau^i)]L(\sigma^i(\tau^i) - \tau^i) \\
&< \mathbf{P}[\sigma^j(\tilde{\tau}^j) \geq \sigma^i(\tau^i) - \varepsilon]L(\sigma^i(\tau^i) - \varepsilon - \tau^i) \\
&= \{\mathbf{P}[\sigma^j(\tilde{\tau}^j) > \sigma^i(\tau^i) - \varepsilon] + \alpha\mathbf{P}[\sigma^j(\tilde{\tau}^j) = \sigma^i(\tau^i) - \varepsilon]\}L(\sigma^i(\tau^i) - \varepsilon - \tau^i) \\
&= V^i(\sigma^i(\tau^i) - \varepsilon - \tau^i, \tau^i, \sigma^j),
\end{aligned}$$

which is ruled out by (5). This contradiction establishes that $\sigma^i(\tau^i) - \tau^i \leq M$ for all i and τ^i. The result follows.

(ii) We first prove that σ^i is nondecreasing for all i, that is, that $\sigma^i(\hat{\tau}^i) \geq \sigma^i(\tau^i)$ for all $\tau^i \geq 0$ and $\hat{\tau}^i > \tau^i$. By Lemma 1(i), the result is obvious if $\hat{\tau}^i \geq \sigma^i(\tau^i)$ or $\sigma^i(\hat{\tau}^i) \geq \tau^i + M$. Hence suppose that $\hat{\tau}^i < \sigma^i(\tau^i)$ and $\sigma^i(\hat{\tau}^i) < \tau^i + M$. It follows from the first of these inequalities that $\sigma^i(\tau^i)$ is a feasible action for type $\hat{\tau}^i$, just like $\sigma^i(\hat{\tau}^i)$ is a feasible action for type τ^i as $\sigma^i(\hat{\tau}^i) \geq \hat{\tau}^i > \tau^i$. Hence, by (5),

$$\begin{aligned}
V^i(\sigma^i(\tau^i), \tau^i, \sigma^j) &\geq V^i(\sigma^i(\hat{\tau}^i), \tau^i, \sigma^j), \\
V^i(\sigma^i(\hat{\tau}^i), \hat{\tau}^i, \sigma^j) &\geq V^i(\sigma^i(\tau^i), \hat{\tau}^i, \sigma^j).
\end{aligned}$$

Summing these two inequalities and rearranging using (4) yields

$$\psi(\sigma^i(\tau^i))[L(\sigma^i(\tau^i) - \tau^i) - L(\sigma^i(\tau^i) - \hat{\tau}^i)] \geq \psi(\sigma^i(\hat{\tau}^i))[L(\sigma^i(\hat{\tau}^i) - \tau^i) - L(\sigma^i(\hat{\tau}^i) - \hat{\tau}^i)], \tag{28}$$

where $\psi(t) \equiv \mathbf{P}[\sigma^j(\tilde{\tau}^j) > t] + \alpha \mathbf{P}[\sigma^j(\tilde{\tau}^j) = t] = \alpha \mathbf{P}[\sigma^j(\tilde{\tau}^j) \geq t] + (1-\alpha)\mathbf{P}[\sigma^j(\tilde{\tau}^j) > t]$ is a nonincreasing function of t. Because $\hat{\tau}^i > \tau^i$, $\hat{\tau}^i < \sigma^i(\tau^i)$, and $\sigma^i(\hat{\tau}^i) < \tau^i + M$, it follows from Lemma 1(i) that the four numbers $\sigma^i(\tau^i) - \tau^i$, $\sigma^i(\tau^i) - \hat{\tau}^i$, $\sigma^i(\hat{\tau}^i) - \tau^i$, and $\sigma^i(\hat{\tau}^i) - \hat{\tau}^i$ belong to $[0, M]$. Now, suppose by way of contradiction that $\sigma^i(\tau^i) > \sigma^i(\hat{\tau}^i)$. Then, as L is strictly increasing and strictly concave over $[0, M]$,

$$0 < L(\sigma^i(\tau^i) - \tau^i) - L(\sigma^i(\tau^i) - \hat{\tau}^i) < L(\sigma^i(\hat{\tau}^i) - \tau^i) - L(\sigma^i(\hat{\tau}^i) - \hat{\tau}^i). \tag{29}$$

It follows from (28)–(29) that $\psi(\sigma^i(\tau^i)) > \psi(\sigma^i(\hat{\tau}^i))$. This, however, is impossible, because $\sigma^i(\tau^i) > \sigma^i(\hat{\tau}^i)$ by assumption and ψ is nonincreasing. This contradiction establishes that σ^i is nondecreasing for all i.

Suppose next by way of contradiction that $\hat{\tau}^i > \tau^i$ and yet $\sigma^i(\hat{\tau}^i) = \sigma^i(\tau^i)$. Then σ^i is constant over $[\tau^i, \hat{\tau}^i]$ and the distribution of player i's equilibrium moving time has an atom at $\sigma^i(\tau^i)$. The proof then relies on the following claim, the proof of which can be found below.

Claim 1 *There exists $\varepsilon_0 > 0$ such that*

$$\sigma^j(\tau^j) \notin (\sigma^i(\tau^i), \sigma^i(\tau^i) + \varepsilon_0), \quad \tau^j \in [0, \sigma^i(\tau^i)). \tag{30}$$

(Observe that the interval $[0, \sigma^i(\tau^i))$ is nonempty as $\sigma^i(\tau^i) \geq \sigma^i(0) > 0$ by Lemma 1(i) along with the fact that σ^i is nondecreasing.) According to Claim 1, the only types of player j who can make a move during $(\sigma^i(\tau^i), \sigma^i(\tau^i) + \varepsilon_0)$ are those such that $\tau^j \geq \sigma^i(\tau^i)$. But it follows from Lemma 1(i) that $\sigma^j(\sigma^i(\tau^i)) = \sigma^i(\tau^i) + \varepsilon_1$ for some $\varepsilon_1 > 0$. Hence, as σ^j is nondecreasing, player j never makes a move during $(\sigma^i(\tau^i), \sigma^i(\tau^i) + \varepsilon_0 \wedge \varepsilon_1)$. As $\sigma^i(\hat{\tau}^i) = \sigma^i(\tau^i) \leq \tau^i + M < \hat{\tau}^i + M$ by Lemma 1(i), one has, letting $\hat{t}^i \equiv \left[\sigma^i(\tau^i) + (\varepsilon_0 \wedge \varepsilon_1)/2\right] \wedge (\hat{\tau}^i + M)$ and using the fact that L is strictly increasing over $[0, M]$,

$$\begin{aligned}
V^i(\sigma^i(\hat{\tau}^i), \hat{\tau}^i, \sigma^j) &= V^i(\sigma^i(\tau^i), \hat{\tau}^i, \sigma^j) \\
&= \{\mathbf{P}[\sigma^j(\tilde{\tau}^j) > \sigma^i(\tau^i)] + \alpha \mathbf{P}[\sigma^j(\tilde{\tau}^j) = \sigma^i(\tau^i)]\} L(\sigma^i(\tau^i) - \hat{\tau}^i) \\
&\leq \mathbf{P}[\sigma^j(\tilde{\tau}^j) \geq \sigma^i(\tau^i)] L(\sigma^i(\tau^i) - \hat{\tau}^i) \\
&< \mathbf{P}[\sigma^j(\tilde{\tau}^j) \geq \sigma^i(\tau^i)] L(\hat{t}^i - \hat{\tau}^i) \\
&= \{\mathbf{P}[\sigma^j(\tilde{\tau}^j) > \hat{t}^i] + \alpha \mathbf{P}[\sigma^j(\tilde{\tau}^j) = \hat{t}^i]\} L(\hat{t}^i - \hat{\tau}^i) \\
&= V^i(\hat{t}^i, \hat{\tau}^i, \sigma^j),
\end{aligned}$$

which is ruled out by (5). This contradiction establishes that σ^i is strictly increasing for all i. The result follows.

To complete the proof of Lemma 1(ii), it remains to prove Claim 1. If $\sigma^j(\sigma^i(\tau^i)-) \leq \sigma^i(\tau^i)$, (30) directly follows from the fact that σ^j is nondecreasing. Hence suppose that $\sigma^j(\sigma^i(\tau^i)-) = \sigma^i(\tau^i) + \varepsilon_1^-$ for some $\varepsilon_1^- > 0$. Then, as σ^j is nondecreasing, there exists $\delta_1 > 0$ such that $\sigma^j(\tau^j) > \sigma^i(\tau^i) + \varepsilon_1^-/2$ for all $\tau^j > \sigma^i(\tau^i) - \delta_1$. Consider now types $\tau^j \leq \sigma^i(\tau^i) - \delta_1$. By Lemma 1(i), among these types, we only need to be concerned by those such that $\tau^j \geq [\sigma^i(\tau^i) - M] \vee 0$. We now show that there exists some $\varepsilon_0^- > 0$ such that each of these types is strictly better off making a move before time $\sigma^i(\tau^i)$ than making a move during $(\sigma^i(\tau^i), \sigma^i(\tau^i) + \varepsilon_0^-)$, from which Claim 1 follows for $\varepsilon_0 \equiv \varepsilon_0^- \wedge (\varepsilon_1^-/2)$. For any type $\tau^j \in [[\sigma^i(\tau^i) - M] \vee 0, \sigma^i(\tau^i) - \delta_1]$, making a move at time $\sigma^i(\tau^i) - \varepsilon$ yields a payoff $\mathbf{P}[\sigma^i(\tilde{\tau}^i) \geq \sigma^i(\tau^i) - \varepsilon]L(\sigma^i(\tau^i) - \varepsilon - \tau^j)$ for all $\varepsilon \in (0, \sigma^i(\tau^i) - \tau^j)$ such that $\mathbf{P}[\sigma^i(\tilde{\tau}^i) = \sigma^i(\tau^i) - \varepsilon] = 0$, whereas making a move at time $\sigma^i(\tau^i) + \varepsilon'$ yields at most a payoff $\mathbf{P}[\sigma^i(\tilde{\tau}^i) \geq \sigma^i(\tau^i) + \varepsilon']L(\sigma^i(\tau^i) + \varepsilon' - \tau^j)$ for all $\varepsilon' > 0$. Suppose by way of contradiction that

$$\begin{aligned}
&\forall \varepsilon_0^- > 0 \;\; \exists \varepsilon' \in [0, \varepsilon_0^-] \;\; \exists \tau^j \in [[\sigma^i(\tau^i) - M] \vee 0, \sigma^i(\tau^i) - \delta_1] \;\; \forall \varepsilon \in (0, \sigma^i(\tau^i) - \tau^j) \\
&\quad \mathbf{P}[\sigma^i(\tilde{\tau}^i) \geq \sigma^i(\tau^i) + \varepsilon']L(\sigma^i(\tau^i) + \varepsilon' - \tau^j) \geq \mathbf{P}[\sigma^i(\tilde{\tau}^i) \geq \sigma^i(\tau^i) - \varepsilon]L(\sigma^i(\tau^i) - \varepsilon - \tau^j).
\end{aligned}$$

Then a fortiori

$$\begin{aligned}
&\forall \varepsilon_0^- > 0 \;\; \exists \varepsilon' \in [0, \varepsilon_0^-] \;\; \exists \tau^j \in [[\sigma^i(\tau^i) - M] \vee 0, \sigma^i(\tau^i) - \delta_1] \;\; \forall \varepsilon \in (0, \sigma^i(\tau^i) - \tau^j) \\
&\quad \{\mathbf{P}[\sigma^i(\tilde{\tau}^i) \geq \sigma^i(\tau^i) - \varepsilon] - \mathbf{P}[\sigma^i(\tilde{\tau}^i) = \sigma^i(\tau^i)]L(\sigma^i(\tau^i) + \varepsilon' - \tau^j)\} \\
&\quad \geq \mathbf{P}[\sigma^i(\tilde{\tau}^i) \geq \sigma^i(\tau^i) - \varepsilon]L(\sigma^i(\tau^i) - \varepsilon - \tau^j),
\end{aligned}$$

so that, letting ε go to zero,

$$\begin{aligned}
&\forall \varepsilon_0^- > 0 \;\; \exists \varepsilon' \in [0, \varepsilon_0^-] \;\; \exists \tau^j \in [[\sigma^i(\tau^i) - M] \vee 0, \sigma^i(\tau^i) - \delta_1] \\
&\quad \mathbf{P}[\sigma^i(\tilde{\tau}^i) \geq \sigma^i(\tau^i)]\left[1 - \frac{L(\sigma^i(\tau^i) - \tau^j)}{L(\sigma^i(\tau^i) + \varepsilon' - \tau^j)}\right] \geq \mathbf{P}[\sigma^i(\tilde{\tau}^i) = \sigma^i(\tau^i)].
\end{aligned}$$

Because the distribution of player i's equilibrium moving time has an atom at $\sigma^i(\tau^i)$, this implies that

$$\inf_{\varepsilon_0^- > 0} \sup_{\varepsilon' \in [0, \varepsilon_0^-]} \sup_{\tau^j \in [[\sigma^i(\tau^i) - M] \vee 0, \sigma^i(\tau^i) - \delta_1]} \left\{1 - \frac{L(\sigma^i(\tau^i) - \tau^j)}{L(\sigma^i(\tau^i) + \varepsilon' - \tau^j)}\right\} > 0. \tag{31}$$

Because L is strictly concave over $[0, M]$ and strictly decreasing over $[M, \infty]$, and because $\sigma^i(\tau^i) - \tau^j \in [\delta_1, M]$ for all $\tau^j \in [[\sigma^i(\tau^i) - M] \vee 0, \sigma^i(\tau^i) - \delta_1]$,

$$\sup_{\tau^j \in [[\sigma^i(\tau^i) - M] \vee 0, \sigma^i(\tau^i) - \delta_1]} \left\{1 - \frac{L(\sigma^i(\tau^i) - \tau^j)}{L(\sigma^i(\tau^i) + \varepsilon' - \tau^j)}\right\} = 1 - \frac{L(\delta_1)}{L(\delta_1 + \varepsilon')}.$$

In turn,

$$\sup_{\varepsilon' \in [0, \varepsilon_0^-]} \left\{ 1 - \frac{L(\delta_1)}{L(\delta_1 + \varepsilon')} \right\} = 1 - \frac{L(\delta_1)}{L(\delta_1 + \varepsilon_0^-)}$$

for all $\varepsilon_0^- \in [0, M - \delta_1]$. As $\delta_1 > 0$, one can let ε_0^- go to zero, and we get that the left-hand side of (31) is zero. This contradiction establishes Claim 1. The result follows.

(iii) Suppose by way of contradiction that $\sigma^i(0) < \sigma^j(0)$ for some i. By Lemma 1(i), $\sigma^j(0) \le M$, and hence $\sigma^i(0) < M$. By Lemma 1(ii), $\sigma^j(\tau) > \sigma^j(0)$ for all $\tau > 0$, and hence $\sigma^j(\tau) > \hat{t}^i$ for all $\tau \ge 0$ and $\hat{t}^i \in (\sigma^i(0), \sigma^j(0) \wedge M)$. It follows that for any such $\hat{t}^i$

$$V^i(\sigma^i(0), 0, \sigma^j) = L(\sigma^i(0)) < L(\hat{t}^i) = V^i(\hat{t}^i, 0, \sigma^j),$$

which is ruled out by (5). This contradiction establishes that $\sigma^a(0) = \sigma^b(0)$. The result follows.

It should be observed for future reference that nothing in the above proofs hinges on the assumption that players have identical breakthrough distributions. Thus the results more generally hold when $\tilde{\tau}^a$ and $\tilde{\tau}^b$ are independently distributed, with distribution functions F^a and F^b that are continuously differentiable and have positive densities $\dot{F}^a$ and $\dot{F}^b$ over $[0, \infty)$. The only change is in (27), where 'F^j' should be substituted to 'F' throughout. ■

Proof of Lemma 2. (i) The proof goes through a series of steps.

Step 1 We first prove that in any equilibrium, if the players' equilibrium strategies have discontinuity points, then the corresponding gaps in the distributions of their moving times $\sigma^a(\tilde{\tau}^a)$ and $\sigma^b(\tilde{\tau}^b)$ cannot overlap. If this did not hold then, because these distributions have no atoms by Lemma 1(ii) along with the assumption that the breakthrough distribution has no atoms, there would exist some player i and some discontinuity point τ^i of σ^i, such that for some $\varepsilon > 0$, with probability 1 player j does not make a move during $[\sigma^i(\tau^i-), \sigma^i(\tau^i-) + \varepsilon]$. One must have $\sigma^i(\tau^i-) < \tau^i + M$, for, otherwise, $\sigma^i(\tau^i+) > \tau^i + M$ as σ^i is discontinuous at τ^i, which would imply that some type of player i close to but above τ^i would have a maturation delay strictly longer than M, which is impossible by Lemma 1(i). As a result, one cannot have $\tau^i = 0$, for, otherwise, player i with type 0 would be strictly better off making a move at time $[\sigma^i(0) + \varepsilon/2] \wedge M$, as she would thereby increase her payoff from moving first, while still avoiding any preemption risk. Thus one can choose $\hat{\tau}^i < \tau^i$ close enough to τ^i such that $\sigma^i(\tau^i-) < \hat{\tau}^i + M$ and

$$L([\sigma^i(\tau^i-) + \varepsilon - \hat{\tau}^i] \wedge M) > \frac{\mathbf{P}[\sigma^j(\tilde{\tau}^j) > \sigma^i(\hat{\tau}^i)]}{\mathbf{P}[\sigma^j(\tilde{\tau}^j) > \sigma^i(\tau^i-)]} L(\sigma^i(\hat{\tau}^i) - \hat{\tau}^i).$$

One then has, letting $\hat{t}^i \equiv [\sigma^i(\tau^i-) + \varepsilon] \wedge (\hat{\tau}^i + M)$ and using the facts that L is strictly increasing over $[0, M]$, and that the distribution of $\sigma^j(\tilde{\tau}^j)$ has no atom and does not charge the interval $[\sigma^i(\tau^i-), \sigma^i(\tau^i-) + \varepsilon]$,

$$\begin{aligned} V^i(\sigma^i(\hat{\tau}^i), \hat{\tau}^i, \sigma^j) &= \mathbf{P}[\sigma^j(\tilde{\tau}^j) > \sigma^i(\hat{\tau}^i)] L(\sigma^i(\hat{\tau}^i) - \hat{\tau}^i) \\ &< \mathbf{P}[\sigma^j(\tilde{\tau}^j) > \sigma^i(\tau^i-)] L(\hat{t}^i - \hat{\tau}^i) \\ &= \mathbf{P}[\sigma^j(\tilde{\tau}^j) > \hat{t}^i] L(\hat{t}^i - \hat{\tau}^i) \\ &= V^i(\hat{t}^i, \hat{\tau}^i, \sigma^j), \end{aligned}$$

which is ruled out by (5). This contradiction establishes that the gaps, if any exists, in the distributions of $\sigma^i(\tilde{\tau}^j)$ and $\sigma^j(\tilde{\tau}^j)$ cannot overlap. As explained in the main text, this notably rules out discontinuous symmetric equilibria.

Step 2 From Step 1 along with the fact that σ^j can only have jump discontinuities, if player i's equilibrium strategy has a discontinuity point at τ^i, then $\tau^i > 0$ and the set $\phi^j((\sigma^i(\tau^i-), \sigma^i(\tau^i+)))$ is well defined. We first prove that in such a case $\sigma^j(\tau^j) = \tau^j + M$ for all $\tau^j \in \phi^j((\sigma^i(\tau^i-), \sigma^i(\tau^i+)))$. Indeed, suppose by way of contradiction that $\sigma^j(\tau^j) < \tau^j + M$ for such a type τ^j. One then has, letting $\hat{t}^j = \sigma^i(\tau^i+) \wedge (\tau^j + M)$ and using the facts that L is strictly increasing over $[0, M]$, and that the distribution of $\sigma^i(\tilde{\tau}^i)$ has no atom and does not charge the interval $[\sigma^i(\tau^i-), \sigma^i(\tau^i+)]$,

$$\begin{aligned} V^j(\sigma^j(\tau^j), \tau^j, \sigma^i) &= \mathbf{P}[\sigma^i(\tilde{\tau}^i) > \sigma^j(\tau^j)] L(\sigma^j(\tau^j) - \tau^j) \\ &< \mathbf{P}[\sigma^i(\tilde{\tau}^i) > \hat{t}^j] L(\hat{t}^j - \tau^j) \\ &= V^j(\hat{t}^j, \tau^j, \sigma^i), \end{aligned}$$

which is ruled out by (5). This contradiction establishes the claim. Now consider type $\overline{\tau}^j \equiv \sup \phi^j((\sigma^i(\tau^i-), \sigma^i(\tau^i+)))$. Because $\sigma^j(\tau^j) = \tau^j + M$ for all $\tau^j < \overline{\tau}^j$ close enough to $\overline{\tau}^j$, it follows from Lemma 1(i)–(ii) that $\sigma^j(\overline{\tau}^j) = \overline{\tau}^j + M$, so that $\overline{\tau}^j = \phi^j(\overline{\tau}^j + M) = \phi^j(\sigma^i(\tau^i+))$. Finally consider type $\underline{\tau}^j \equiv \inf \phi^j((\sigma^i(\tau^i-), \sigma^i(\tau^i+)))$. Because $\sigma^j(\tau^j) = \tau^j + M$ for all $\tau^j > \underline{\tau}^j$ close enough to $\underline{\tau}^j$, σ^j is discontinuous at $\underline{\tau}^i$ if $\sigma^j(\underline{\tau}^j) < \underline{\tau}^j + M = \sigma^i(\tau^i-)$. But then, it follows from the above reasoning that $\sigma^i(\tau^i-) = \tau^i + M$, which is impossible as noted in Step 1. Thus $\sigma^j(\underline{\tau}^j) = \underline{\tau}^j + M$, so that $\underline{\tau}^j = \phi^j(\underline{\tau}^j + M) = \phi^j(\sigma^i(\tau^i-))$. Overall, we have shown that $\sigma^j(\tau^j) = \tau^j + M$ for all $\tau^j \in \phi^j([\sigma^i(\tau^i-), \sigma^i(\tau^i+)])$.

Step 3 Suppose as in Step 2 that player i's equilibrium strategy has a discontinuity point at $\tau^i > 0$. Then consider type $\tilde{\tau} = \sup\{\tau \in [0, \tau^i) : \sigma^i(\tau) \geq \sigma^j(\tau)\}$, which is well defined as $\sigma^i(0) = \sigma^j(0)$, and strictly less than τ^i by Step 2. Observe that σ^j must be continuous over $(\tilde{\tau}, \tau^i]$, for, otherwise, it would follow from Step 2 that $\sigma^i(\tau) = \tau + M \geq \sigma^j(\tau)$ for some $\tau \in (\tilde{\tau}, \tau^i]$, in contradiction with the definition of $\tilde{\tau}$. We now show that $\sigma^i(\tilde{\tau}+) = \sigma^j(\tilde{\tau}+)$. Clearly one must have $\sigma^i(\tilde{\tau}+) \leq \sigma^j(\tilde{\tau}+)$, for, otherwise, one would have $\sigma^i(\tau) > \sigma^j(\tau)$ for some $\tau \in (\tilde{\tau}, \tau^i)$, in contradiction with the definition of $\tilde{\tau}$. Then suppose by way of contradiction that $\sigma^i(\tilde{\tau}+) < \sigma^j(\tilde{\tau}+)$. According to the definition of $\tilde{\tau}$, there exists a nondecreasing sequence $\{\tau_n\}$ converging to $\tilde{\tau}$ such that $\sigma^i(\tau_n) \geq \sigma^j(\tau_n)$ for all n. Thus $\sigma^j(\tilde{\tau}+) > \sigma^i(\tilde{\tau}+) > \sigma^i(\tau_n) \geq \sigma^j(\tau_n)$ which shows that σ^j is discontinuous at $\tilde{\tau}$. But it then follows from Step 2 that $\sigma^i(\tilde{\tau}+) \geq \sigma^j(\tilde{\tau}+)$. Together with the fact that $\sigma^i(\tilde{\tau}+) \leq \sigma^j(\tilde{\tau}+)$, this contradiction shows that $\sigma^i(\tilde{\tau}+) = \sigma^j(\tilde{\tau}+) \equiv \tilde{\sigma}$, as claimed.

Step 4 Define τ^i, $\tilde{\tau}$, and $\tilde{\sigma}$ as in Step 3. Consider the functions ϕ^i and ϕ^j. As for ϕ^j, it is continuous and strictly increasing over $(\tilde{\sigma}, \sigma^i(\tau^i-))$ because σ^j is strictly increasing and continuous over $(\tilde{\tau}, \tau^i)$ by Lemma 1(ii) and Step 3. As for ϕ^i, it may not be defined over the entire interval $(\tilde{\sigma}, \sigma^i(\tau^i-))$ because σ^i may have discontinuity points in $(\tilde{\tau}, \tau^i)$. Yet one can straightforwardly extend ϕ^i to all of $(\tilde{\sigma}, \sigma^i(\tau^i-))$ by requiring it to be constant over any interval $[\sigma^i_-, \sigma^i_+]$ corresponding to a discontinuity point of σ^i. Call $\overline{\phi}^i$ the function generated in this way, which is continuous and nondecreasing. We first establish that $\overline{\phi}^i$ and ϕ^j are Lipschitz over $(\tilde{\sigma}, \sigma^i(\tau^i-))$.

We start with ϕ^j, and study to that effect the incentives of player i. Two cases must be distinguished. First, if $t \in (\tilde{\sigma}, \sigma^i(\tau^i-)) \cap \sigma^i([0,\infty))$, then $\phi^i(t)$ is well defined. Because ϕ^j is continuous over $(\tilde{\sigma}, \sigma^i(\tau^i-))$, so is the maximization problem faced by any type of player i that belongs to $(\tilde{\tau}, \tau^i)$. Hence, by Berge's maximum theorem, the maximizer correspondence is upper hemicontinuous in player i's type. In the present case, this notably implies that we can without loss of generality assume that there exists a strictly increasing sequence $\{t_n\}$ converging to t in $(\tilde{\sigma}, \sigma^i(\tau^i-)) \cap \sigma^i([0,\infty))$. By Lemma 1(i), $t > \phi^i(t)$. Thus for n large enough, type $\phi^i(t)$ could deviate and make a move at time t_n as type $\phi^i(t_n)$ does. It follows from (5) along with the fact that the distribution of $\sigma^j(\tilde{\tau}^j)$ has no atom by Lemma 1(ii) that

$$\begin{aligned}
[1 - F(\phi^j(t))]L(t - \phi^i(t)) &= \mathbf{P}[\sigma^j(\tilde{\tau}^j) > t]L(t - \phi^i(t)) \\
&= V^i(t, \phi^i(t), \sigma^j) \\
&\geq V^i(t_n, \phi^i(t), \sigma^j) \\
&= \mathbf{P}[\sigma^j(\tilde{\tau}^j) > t_n]L(\hat{t} - \phi^i(t)) \\
&= [1 - F(\phi^j(t_n))]L(t_n - \phi^i(t))
\end{aligned}$$

for n large enough. Rearranging and using the fact that ϕ^j is strictly increasing, we get

$$0 < F(\phi^j(t)) - F(\phi^j(t_n)) \leq [1 - F(\phi^j(t))] \, \frac{L(t - \phi^i(t)) - L(t_n - \phi^i(t))}{L(t_n - \phi^i(t))}$$

for n large enough. Dividing through by $t - t_n$ and letting t_n increase to $t > \phi^i(t)$, we conclude that

$$0 \leq D_-[F \circ \phi^j](t) \leq [1 - F(\phi^j(t))] \, \frac{\dot{L}(t - \phi^i(t))}{L(t - \phi^i(t))}, \quad t \in (\tilde{\sigma}, \sigma^i(\tau^i-)) \cap \sigma^i([0, \infty)), \tag{32}$$

where $D_-[F \circ \phi^j](t)$ is the lower left Dini derivative of $F \circ \phi^j$ at t. As $t - \phi^i(t)$ is bounded away from zero over $(\tilde{\sigma}, \sigma^i(\tau^i-)) \cap \sigma^i([0, \infty))$, it follows that $D_-[F \circ \phi^j]$ is bounded over this set. Now, if $t \in (\tilde{\sigma}, \sigma^i(\tau^i-)) \setminus \sigma^i([0, \infty))$, then $\phi^j(t) = t - M$ by Step 2. Because F is continuously differentiable, there exists a constant K such that

$$0 \leq D_-[F \circ \phi^j](t) \leq K, \quad t \in (\tilde{\sigma}, \sigma^i(\tau^i-)) \setminus \sigma^i([0, \infty)). \tag{33}$$

Combining the bounds (32) and (33), we get that $F \circ \phi^j$ is Lipschitz over $(\tilde{\sigma}, \sigma^i(\tau^i-))$ (see for instance Giorgi and Komlósi (1992, Lemma 1.15)). Moreover, because $\dot{F}$ is continuous and positive over $[0, \infty)$, F^{-1} is locally Lipschitz over $[0, 1)$. Hence ϕ^j is Lipschitz over $(\tilde{\sigma}, \sigma^i(\tau^i-))$.

We consider next ϕ^i, and study to that effect the incentives of player j. Again, two cases must be distinguished. First suppose that $t \in (\tilde{\sigma}, \sigma^i(\tau^i-)) \cap \sigma^i([0, \infty))$, and consider an approximating sequence $\{t_n\}$ as above. By Lemma 1(i), $t > \phi^j(t)$. Thus for n large enough, type $\phi^j(t)$ could deviate and make a move at time t_n as type $\phi^j(t_n)$ does. Proceeding in a similar way, and using the fact that $\phi^i = \overline{\phi}^i$ over $(\tilde{\sigma}, \sigma^i(\tau^i-)) \cap \sigma^i([0, \infty))$, we get

$$0 < F(\overline{\phi}^i(t)) - F(\overline{\phi}^i(t_n)) \leq [1 - F(\overline{\phi}^i(t))] \, \frac{L(t - \phi^j(t)) - L(t_n - \phi^j(t))}{L(t_n - \phi^j(t))}$$

for n large enough. Dividing through by $t - t_n$ and letting t_n increase to $t > \phi^i(t) = \overline{\phi}^i(t)$, we conclude that

$$0 \leq D_-[F \circ \overline{\phi}^i](t) \leq [1 - F(\overline{\phi}^i(t))] \, \frac{\dot{L}(t - \phi^j(t))}{L(t - \phi^j(t))}, \quad t \in (\tilde{\sigma}, \sigma^i(\tau^i-)) \cap \sigma^i([0, \infty)), \tag{34}$$

As $t - \phi^j(t)$ is bounded away from zero over $(\tilde{\sigma}, \sigma^i(\tau^i-)) \cap \sigma^i([0, \infty))$, it follows that $D_-[F \circ \overline{\phi}^i]$ is bounded over this set. Now, if $t \in (\tilde{\sigma}, \sigma^i(\tau^i-)) \setminus \sigma^i([0, \infty))$, then $\overline{\phi}^i$ is constant over $(t - \varepsilon, t]$ for some $\varepsilon > 0$. Thus

$$D_-[F \circ \overline{\phi}^i](t) = 0, \quad t \in (\tilde{\sigma}, \sigma^i(\tau^i-)) \setminus \sigma^i([0, \infty)). \tag{35}$$

Combining the bounds (34) and (35), and reasoning as in the case of ϕ^j, we obtain that $\overline{\phi}^i$ is Lipschitz over $(\tilde{\sigma}, \sigma^i(\tau^i-))$.

Step 5 Define the functions $\overline{\phi}^i$ and ϕ^j as in Step 3. Because they both are Lipschitz over $(\tilde{\sigma}, \sigma^i(\tau^i-))$, they both are absolutely continuous and thus almost everywhere differentiable over this interval. Their derivatives whenever they exist can be evaluated as follows. Consider first some $t \in (\tilde{\sigma}, \sigma^i(\tau^i-)) \setminus \sigma^i([0,\infty))$. If $\overline{\phi}^i$ and ϕ^j are differentiable at t, then, by Step 2, $\dot{\overline{\phi}}^i(t) = 0$ and $\dot{\phi}^j(t) = 1$. Consider next some $t \in (\tilde{\sigma}, \sigma^i(\tau^i-)) \cap \sigma^i([0,\infty))$. If $\overline{\phi}^i$ and ϕ^j are differentiable at t, then $\dot{\overline{\phi}}^i(t)$ (respectively $\dot{\phi}^j(t)$) can be evaluated by differentiating the mapping $\hat{t} \mapsto [1 - F(\overline{\phi}^i(\hat{t}))]L(\hat{t} - \phi^j(t))$ (respectively $\hat{t} \mapsto [1 - F(\phi^j(\hat{t}))]L(\hat{t} - \overline{\phi}^i(t))$) and requiring that the resulting derivative, whenever it exists, be equal to zero at $\hat{t} = t$, as implied by optimality. (Observe that in this case $\overline{\phi}^i(t) = \phi^i(t)$.) For any such t, this yields

$$\dot{\overline{\phi}}^i(t) = \frac{1 - F(\overline{\phi}^i(t))}{\dot{F}(\overline{\phi}^i(t))} \frac{\dot{L}(t - \phi^j(t))}{L(t - \phi^j(t))}, \tag{36}$$

$$\dot{\phi}^j(t) = \frac{1 - F(\phi^j(t))}{\dot{F}(\phi^j(t))} \frac{\dot{L}(t - \overline{\phi}^i(t))}{L(t - \overline{\phi}^i(t))}. \tag{37}$$

Define now the quantity

$$R(t) \equiv \ln\left(\frac{1 - F(\phi^j(t))}{1 - F(\overline{\phi}^i(t))}\right), \quad t \in (\tilde{\sigma}, \sigma^i(\tau^i-)). \tag{38}$$

Using that $\overline{\phi}^i$ and ϕ^j are absolutely continuous over $(\tilde{\sigma}, \sigma^i(\tau^i-))$, that $\dot{F}$ is continuous, that $F \circ \overline{\phi}^i$ and $F \circ \phi^j$ are bounded away from 1 over $(\tilde{\sigma}, \sigma^i(\tau^i-))$, and thus that the logarithm function is Lipschitz over the corresponding range of $(1 - F \circ \phi^j)/(1 - F \circ \overline{\phi}^i)$, we get that R is absolutely continuous over $(\tilde{\sigma}, \sigma^i(\tau^i-))$, and hence is equal to the integral of its derivative, which is well defined almost everywhere. Now, for each $t \in (\tilde{\sigma}, \sigma^i(\tau^i-)) \setminus \sigma^i([0,\infty))$ such that $\overline{\phi}^i$ and ϕ^j are differentiable at t, we have $\dot{\overline{\phi}}^i(t) = 0$ and $\dot{\phi}^j(t) = 1$ and thus

$$\dot{R}(t) = -\frac{\dot{F}(\phi^j(t))}{1 - F(\phi^j(t))} < 0.$$

Similarly, for each $t \in (\tilde{\sigma}, \sigma^i(\tau^i-)) \cap \sigma^i([0,\infty))$ such that $\overline{\phi}^i$ and ϕ^j are differentiable at t,

$$\begin{aligned}\dot{R}(t) &= \frac{\dot{F}(\overline{\phi}^i(t))}{1 - F(\overline{\phi}^i(t))} \dot{\overline{\phi}}^i(t) - \frac{\dot{F}(\phi^j(t))}{1 - F(\phi^j(t))} \dot{\phi}^j(t) \\ &= \frac{\dot{L}(t - \phi^j(t))}{L(t - \phi^j(t))} - \frac{\dot{L}(t - \overline{\phi}^i(t))}{L(t - \overline{\phi}^i(t))} \\ &< 0,\end{aligned}$$

where the second equality follows from (36)–(37), and the inequality from the fact that $\phi^j < \overline{\phi}^i$ over $(\tilde{\sigma}, \sigma^i(\tau^i-))$ by the definition of $\tilde{\sigma}$ in Step 2, along with the assumption that L is strictly concave over $[0, M]$. We thus obtain that R is strictly decreasing over $(\tilde{\sigma}, \sigma^i(\tau^i-))$.

Now, we have

$$\phi^j(\sigma^i(\tau^i-)) = \sigma^i(\tau^i-) - M < \sigma^i(\tau^i+) - M \leq \tau^i = \overline{\phi}^i(\sigma^i(\tau^i-)),$$

where the first equality follows from Step 2, the first inequality from the discontinuity of σ^i at τ^i, and the second inequality from Lemma 1(i). Thus, by (38), $R(\sigma^i(\tau^i-)) > 0$ and, therefore, as R is strictly decreasing over $(\tilde{\sigma}, \sigma^i(\tau^i-))$, $R(\tilde{\sigma}+) > 0$. This, however, is ruled out by the fact that, as shown in Step 3, $\sigma^i(\tilde{\tau}+) = \sigma^j(\tilde{\tau}+) = \tilde{\sigma}$, so that $\overline{\phi}^i(\tilde{\sigma}+) = \phi^j(\tilde{\sigma}+) = \tilde{\tau}$ and thus, by (38) again, $R(\tilde{\sigma}+) = 0$. This contradiction establishes that σ^i is continuous for all i. The result follows.

(ii) It follows from Lemma 2(i) that ϕ^a and ϕ^b are defined and absolutely continuous over $[\sigma(0), \infty)$, and that they satisfy (8) almost everywhere in $(\sigma(0), \infty)$. As a result,

$$\phi^i(t) = \phi^i(t_0) + \int_{t_0}^{t} \frac{1 - F(\phi^i(s))}{\dot{F}(\phi^i(s))} \frac{\dot{L}(s - \phi^j(s))}{L(s - \phi^j(s))} \, ds, \quad t > t_0 > \sigma(0), \quad i = a, b. \tag{39}$$

But because, for each i, $s - \phi^j(s)$ is bounded away from zero over any compact subinterval of $(\sigma(0), \infty)$, the integrand in (39) is continuous in s over any such interval. Thus, by the fundamental theorem of calculus, one may differentiate (39) everywhere with respect to t to get that (8) holds for all $t > \sigma(0)$. To conclude the proof, observe that for each i, ϕ^i is continuous at $\sigma(0)$ by Lemma 1(ii), and that so is the integrand in (39) and thus $\dot{\phi}^i$ as $\sigma(0) > 0$ by Lemma 1(i). This implies that $\dot{\phi}^i$ can be continuously extended at $\sigma(0)$. The result follows. ■

Proof of Lemma 3. Suppose by way of contradiction that $\phi^i(t) > \phi^j(t)$ for some $t > \sigma(0)$. As $\phi^i(\sigma(0)) = \phi^j(\sigma(0))$, $t_0 \equiv \sup\{s \in [\sigma(0), t) : \phi^i(s) = \phi^j(s)\}$ is well defined and strictly less than t. Moreover, $\phi^i(t_0) = \phi^j(t_0)$ and $\phi^i > \phi^j$ over (t_0, t). Defining the function R as in (38), one has

$$R(t) = \int_{t_0}^{t} \left[\frac{\dot{L}(s - \phi^j(s))}{L(s - \phi^j(s))} - \frac{\dot{L}(s - \phi^i(s))}{L(s - \phi^i(s))} \right] ds. \tag{40}$$

Because $\phi^i(t) > \phi^j(t)$, the left-hand side of (40) is positive by (38). However, because $\phi^i > \phi^j$ over (t_0, t) and L is strictly concave over $[0, M]$, the right-hand side of (40) is negative. This contradiction establishes that $\phi^i(t) = \phi^j(t)$ for all $t > \sigma(0)$. The result follows. ■

Proof of Proposition 1. We only need to show that there exists at least one value of $\sigma_0 \in (0, M)$ such that the solution ϕ to the ODE (9) initiated at $(t, \tau) = (\sigma_0, 0)$ remains in $\mathcal{D}$. It is helpful for the purpose of this proof to consider the differential equation for the inverse σ of ϕ. Specifically, for each $\sigma_0 \in (0, M)$, consider the following initial value problem:

$$\dot{\sigma}(\tau) = [f(\sigma(\tau), \tau)]^{-1}, \quad \tau \geq 0, \tag{41}$$

$$\sigma(0) = \sigma_0. \tag{42}$$

It is easy to check from the definition (10) of f that over the interior $\operatorname{Int}\mathcal{D}'$ of the domain $\mathcal{D}' \equiv \{(\tau, \sigma) : 0 \leq \tau < \sigma \leq \tau + M\}$, the mapping $(\tau, \sigma) \mapsto [f(\sigma, \tau)]^{-1}$ is continuous and locally Lipschitz in σ. Hence, by the Cauchy–Lipschitz theorem, for each $\sigma_0 \in (0, M)$, the problem (41)–(42) has a unique maximal solution $\sigma(\tau, 0, \sigma_0)$ in $\mathcal{D}'$ (see for instance Perko (2001, Section 2.2, Theorem, and Section 2.4, Theorem 1)). Also define the degenerate solutions $\sigma(\cdot, 0, 0) \equiv \{(0, 0)\}$ and $\sigma(\cdot, 0, M) \equiv \{(0, M)\}$ for $\sigma_0 = 0$ and $\sigma_0 = M$, respectively. For each $\sigma_0 \in [0, M]$, define accordingly

$$\tau(\sigma_0) \equiv \sup\{\tau \geq 0 : (\tau', \sigma(\tau', 0, \sigma_0)) \in \operatorname{Int}\mathcal{D}' \text{ for all } \tau' \in (0, \tau)\},$$

with $\sup \emptyset = 0$ by convention, so that $\tau(0) = \tau(M) = 0$. The proof is complete if we show that $\tau(\sigma_0) = \infty$ for some $\sigma_0 \in (0, M)$. To this end, define

$$L_0 \equiv \{\sigma_0 \in [0, M] : \tau(\sigma_0) < \infty \text{ and } \sigma(\tau(\sigma_0), 0, \sigma_0) = \tau(\sigma_0)\},$$
$$U_0 \equiv \{\sigma_0 \in [0, M] : \tau(\sigma_0) < \infty \text{ and } \sigma(\tau(\sigma_0), 0, \sigma_0) = \tau(\sigma_0) + M\}.$$

Clearly $L_0 \neq \emptyset$ as $0 \in L_0$, $U_0 \neq \emptyset$ as $M \in U_0$, and $L_0 \cap U_0 = \emptyset$. If we knew that both L_0 and U_0 were relatively open in $[0, M]$ then, because $[0, M]$ is connected and thus cannot be the union of two disjoint open sets, we could argue that there must exist some $\sigma_0 \in [0, M]$ such that $\sigma_0 \notin L_0 \cup U_0$. Given the definitions of $\tau(\sigma_0)$, L_0, and U_0, it would follow that $\tau(\sigma_0) = \infty$. In fact, using the observation that $[f(\sigma, \tau)]^{-1}$ is strictly increasing in σ, one would get the stronger result that $L_0 = [0, \underline{\sigma})$ and $U_0 = (\overline{\sigma}, M]$, where $0 < \underline{\sigma} \leq \overline{\sigma} < M$, so that $\tau(\sigma_0) = \infty$ if and only if $\sigma_0 \in \Sigma_0 \equiv [\underline{\sigma}, \overline{\sigma}]$. The proof that this is indeed the case relies on the following claim, the proof of which can be found below.

Claim 2 *For each $(\tau_1, \sigma_1) \in \operatorname{Int}\mathcal{D}'$, the terminal value problem*

$$\dot{\sigma}(\tau) = [f(\sigma(\tau), \tau)]^{-1}, \quad \tau \leq \tau_1, \tag{43}$$

$$\sigma(\tau_1) = \sigma_1 \tag{44}$$

has a unique solution $\sigma(\cdot, \tau_1, \sigma_1)$ over $[0, \tau_1]$ in $\mathcal{D}'$.

We now show that L_0 is relatively open. Given that L_0 is an interval that contains 0, we only need to show that if $\sigma_0 \in L_0$, then $\sigma_0' \in L_0$ for some $\sigma_0' > \sigma_0$. For each $\varepsilon > 0$, consider the solution $\sigma(\cdot, \tau(\sigma_0), \tau(\sigma_0) + \varepsilon)$ to problem (43)–(44) with terminal condition $\sigma(\tau(\sigma_0)) = \tau(\sigma_0) + \varepsilon$. By Claim 2, this solution can be maximally extended to 0, and $\sigma_0'(\varepsilon) \equiv \sigma(0, \tau(\sigma_0), \tau(\sigma_0) + \varepsilon) \in (\sigma_0, M)$. Now, consider the solution $\sigma(\cdot, 0, \sigma_0'(\varepsilon))$ to problem (41)–(42) with initial condition $\sigma(0) = \sigma_0'(\varepsilon)$, so that $\sigma(\tau(\sigma_0), 0, \sigma_0'(\varepsilon)) = \tau(\sigma_0) + \varepsilon$, and suppose that $\sigma_0'(\varepsilon) \notin L_0$ for all $\varepsilon > 0$, so that $\sigma(\cdot, 0, \sigma_0'(\varepsilon))$ never leaves $\mathcal{D}'$ through its lower boundary. Notice that there exist $k < 1$ and an open ball B with radius $\eta > 0$ centered at $(\tau(\sigma_0), \tau(\sigma_0))$ such that $[f(\sigma, \tau)]^{-1} \leq k$ for all $(\sigma, \tau) \in B \cap \operatorname{Int} \mathcal{D}'$. Because the slope of the lower boundary $\sigma = \tau$ of $\mathcal{D}'$ is $1 > k$, for each $\varepsilon \in (0, (1-k)\eta\sqrt{2}/2)$, the segment of slope k connecting $(\tau(\sigma_0), \tau(\sigma_0) + \varepsilon)$ to $(\tau(\sigma_0) + \varepsilon/(1-k), \tau(\sigma_0) + \varepsilon/(1-k))$ is contained in B. As $\sigma(\cdot, 0, \sigma_0'(\varepsilon))$ does not leave $\mathcal{D}'$ through its lower boundary, it must eventually leave B; but, because of the above observation, and because $\sigma(\tau(\sigma_0), 0, \sigma_0'(\varepsilon)) = \tau(\sigma_0) + \varepsilon$, it cannot do so before time $\tau(\sigma_0) + \varepsilon/(1-k)$. In particular,

$$\sigma\left(\tau(\sigma_0) + \frac{\varepsilon}{1-k}, 0, \sigma_0'(\varepsilon)\right) \leq \tau(\sigma_0) + \varepsilon + k\left[\tau(\sigma_0) + \frac{\varepsilon}{1-k} - \tau(\sigma_0)\right] = \tau(\sigma_0) + \frac{\varepsilon}{1-k},$$

so that $\sigma(\cdot, 0, \sigma_0'(\varepsilon))$ must actually leave $\mathcal{D}'$ through its lower boundary at some time $\tau \leq \tau(\sigma_0) + \varepsilon/(1-k)$, a contradiction. It follows that $\sigma_0'(\varepsilon) \in L_0$ for all $\varepsilon > 0$ close enough to zero, which proves the claim as $\sigma_0'(\varepsilon) > \sigma_0$ for any such ε. The proof that U_0 is relatively open is similar, and is therefore omitted. Hence the result.

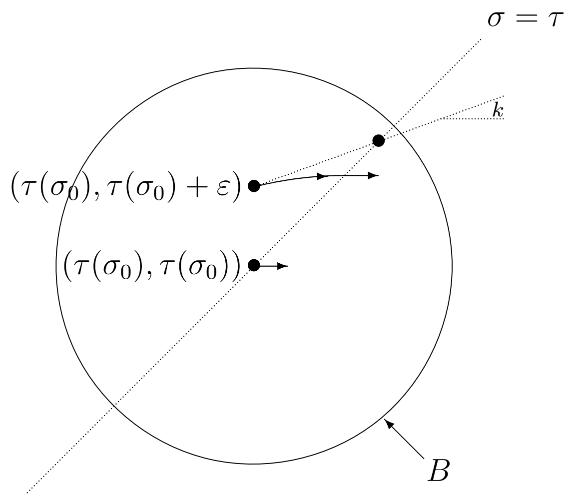

Figure 3 Illustration of the last step of the proof.

To complete the proof of Proposition 1, it remains to prove Claim 2. By the Cauchy–Lipschitz theorem, the terminal value problem (43)–(44) has a unique maximal solution

$\sigma(\cdot, \tau_1, \sigma_1)$ in $\mathcal{D}'$. Because $(\tau_1, \sigma_1) \in \operatorname{Int} \mathcal{D}'$,

$$\tau_0 \equiv \inf \{\tau \leq \tau_1 : (\sigma(\tau', \tau_1, \sigma_1), \tau') \in \operatorname{Int} \mathcal{D}' \text{ for all } \tau' \in (\tau, \tau_1)\} < \tau_1.$$

We now show that $\tau_0 = 0$, which concludes the proof. Suppose by way of contradiction that $\tau_0 > 0$. Then either $\sigma(\cdot, \tau_1, \sigma_1)$ leaves $\mathcal{D}'$ through its lower boundary, so that $\sigma(\tau_0, \tau_1, \sigma_1) = \tau_0$, or $\sigma(\cdot, \tau_1, \sigma_1)$ leaves $\mathcal{D}'$ through its upper boundary, so that $\sigma(\tau_0, \tau_1, \sigma_1) = \tau_0 + M$. In the first case, there exist $k < 1$ and $\varepsilon_0 \in (0, \tau_1 - \tau_0)$ such that $(\partial\sigma/\partial\tau)(\tau_0 + \varepsilon, \tau_1, \sigma_1) < k$ for all $\varepsilon \in (0, \varepsilon_0)$. For any such ε, we have

$$\sigma(\tau_0 + \varepsilon, \tau_1, \sigma_1) < \sigma(\tau_0, \tau_1, \sigma_1) + k\varepsilon = \tau_0 + k\varepsilon < \tau_0 + \varepsilon,$$

so that $(\tau_0 + \varepsilon, \sigma(\tau_0 + \varepsilon, \tau_1, \sigma_1)) \notin \operatorname{Int} \mathcal{D}'$, in contradiction with the definition of τ_0. In the second case, there exist $K > 1$ and $\varepsilon_0 \in (0, \tau_1 - \tau_0)$ such that $(\partial\sigma/\partial\tau)(\tau_0 + \varepsilon, \tau_1, \sigma_1) > K$ for all $\varepsilon \in (0, \varepsilon_0)$. For any such ε, we have

$$\sigma(\tau_0 + \varepsilon, \tau_1, \sigma_1) < \sigma(\tau_0 + \varepsilon_0, \tau_1, \sigma_1) - K(\varepsilon_0 - \varepsilon) < \tau_0 + M - (K - 1)\varepsilon_0 + K\varepsilon,$$

where the second inequality follows from the definition of τ_0. Letting ε go to zero, we obtain that $\sigma(\tau_0, \tau_1, \sigma_1) \leq \tau_0 + M - (K - 1)\varepsilon_0 < \tau_0 + M$, a contradiction. Hence $\tau_0 = 0$, and moreover $\sigma(0, \tau_1, \sigma_1) \in (0, M)$. The result follows. ■

Proof of Proposition 2. The logic is similar to Hubbard and West (1991, Exercise 4.7#3). Consider two symmetric equilibria σ_1 and σ_2 with $\sigma_1(0) \geq \sigma_2(0)$, and introduce the following gap function:

$$g(\tau) \equiv \sigma_1(\tau) - \sigma_2(\tau), \quad \tau \geq 0. \tag{45}$$

We have by construction

$$M > \sigma_1(\tau) - \tau \geq \sigma_2(\tau) - \tau > 0 \tag{46}$$

for all $\tau \geq 0$, where the middle inequality follows from the uniqueness part of the Cauchy–Lipschitz theorem along with the assumption that $\sigma_1(0) \geq \sigma_2(0)$. It follows in particular that g is bounded above by M. We have

$$\begin{aligned} \dot{g}(\tau) &= \frac{\dot{F}(\tau)}{1 - F(\tau)} \left[\frac{L(\sigma_1(\tau) - \tau)}{\dot{L}(\sigma_1(\tau) - \tau)} - \frac{L(\sigma_2(\tau) - \tau)}{\dot{L}(\sigma_2(\tau) - \tau)} \right] \\ &\geq \frac{\dot{F}(\tau)}{1 - F(\tau)} \left\| \frac{d}{dm} \left(\frac{L}{\dot{L}} \right) \right\|_{[\sigma_2(\tau) - \tau, \sigma_1(\tau) - \tau]} g(\tau) \end{aligned} \tag{47}$$

$$\geq \frac{\dot{F}(\tau)}{1-F(\tau)}\, g(\tau)$$

for all $\tau \geq 0$, where the equality follows from (10) and (41), the first inequality from the definition of g, and the second inequality from the fact that $(d/dm)(L/\dot{L})(m) = 1 - [L\ddot{L}/(\dot{L})^2](m) \geq 1$ for all $m \in [0, M)$. Integrating (47) yields

$$g(\overline{\tau}) \geq g(0) \exp\left(\int_0^{\overline{\tau}} \frac{\dot{F}(\tau)}{1-F(\tau)}\, d\tau \right) \tag{48}$$

and thus

$$g(0) \leq M[1 - F(\overline{\tau})] \tag{49}$$

for all $\overline{\tau} \geq 0$. Because $g(0) \geq 0$, letting $\overline{\tau}$ go to infinity in (49) shows that $g(0) = 0$ and thus $\sigma_1(0) = \sigma_2(0)$. From the uniqueness part of the Cauchy–Lipschitz theorem, we have $\sigma_1 = \sigma_2$. Hence the result. ■

Proof of Proposition 3. Suppose first that F has a differentiably strictly increasing breakthrough rate. We prove equivalently that the mapping $t \rightarrow t - \phi(t)$ is strictly decreasing over $[\sigma(0), \infty)$. The proof consists of two steps.

Step 1 We first show that one cannot have $\dot{\phi} \leq 1$ over an interval $[t_0, \infty)$. Suppose the contrary holds. Then the mapping $t \mapsto t - \phi(t)$ is nondecreasing over $[t_0, \infty)$. Because (i) $\dot{L}/L$ is strictly decreasing over $[0, M]$, (ii) $(1-F)/\dot{F}$ is strictly decreasing over $[0, \infty)$, and (iii) ϕ is strictly increasing over $[\sigma(0), \infty)$, it follows from (9)–(10) that $\dot{\phi}$ is strictly decreasing over $[t_0, \infty)$. Because $\dot{\phi} \geq 0$, it follows that $\dot{\phi}(t)$ has a well-defined limit $\dot{\phi}(\infty)$ as t goes to infinity. Clearly, one must have $\dot{\phi}(\infty) < 1$ as $\dot{\phi} \leq 1$ on $[t_0, \infty)$ and $\dot{\phi}$ is strictly decreasing over this interval. But then ϕ would ultimately leave $\mathcal{D}$, a contradiction. The claim follows.

Step 2 By Step 1, there are arbitrarily large times t_0 such that $\dot{\phi}(t_0) > 1$. Fix one of them. We now show that $\dot{\phi} > 1$ over $[\sigma(0), t_0]$. Suppose the contrary holds, and let $t_1 \equiv \sup\{t < t_0 : \dot{\phi}(t) \leq 1\}$. Then $\dot{\phi}(t_1) = 1$ and $\ddot{\phi}(t_1) \geq 0$. Differentiating (9) and using the fact that $\dot{\phi}(t_1) = 1$ yields

$$\begin{aligned}
\ddot{\phi}(t_1) &= \frac{d}{dt}\left(\frac{1-F}{\dot{F}}\right)(\phi(t_1))\, \dot{\phi}(t_1)\, \frac{\dot{L}(t_1 - \phi(t_1))}{L(t_1 - \phi(t_1))} \\
&\qquad\qquad + \frac{1 - F(\phi(t_1))}{\dot{F}(\phi(t_1))}\, \frac{d}{dm}\left(\frac{\dot{L}}{L}\right)(t_1 - \phi(t_1))\, [1 - \dot{\phi}(t_1)] \\
&= \frac{d}{dt}\left(\frac{1-F}{\dot{F}}\right)(\phi(t_1))\, \frac{\dot{L}(t_1 - \phi(t_1))}{L(t_1 - \phi(t_1))}.
\end{aligned}$$

Now, $(d/dt)\,[(1-F)/\dot{F}](\phi(t_1)) < 0$ as F has a differentiably strictly increasing breakthrough rate, and $(\dot{L}/L)(t_1 - \phi(t_1)) > 0$ from $\dot{\phi}(t_1) = 1$ along with (9)–(10). Thus $\ddot{\phi}(t_1) < 0$, a contradiction. Hence $\dot{\phi} > 1$ over $[\sigma(0), t_0]$, as claimed. Because t_0 can be arbitrarily large, it follows that $\dot{\phi} > 1$ over $[\sigma(0), \infty)$. Hence the result.

If F has a differentiably strictly decreasing breakthrough rate, the proof is similar. The only modification to Step 1 consists in observing that if $\dot{\phi}$ is strictly increasing over some interval $[t_0, \infty)$, then $\dot{\phi}(t)$ has a well-defined limit $\dot{\phi}(\infty) \in [0, \infty]$. ∎

Proof of Corollary 1. Four cases must be distinguished.

Case 1 Suppose first that F has a differentiably strictly increasing breakthrough rate with $\lim_{t\to\infty}[\dot{F}/(1-F)](t) \equiv \lambda_\infty < \infty$. Then, according to Proposition 3, the maturation delay $t - \phi(t)$ decreases to a limit m as t goes to infinity. It then follows from (9)–(10) that $\dot{\phi}(t)$ converges to $\dot{\phi}(\infty) \equiv (1/\lambda_\infty)(\dot{L}/L)(m)$ as t goes to infinity. But one must then have $\dot{\phi}(\infty) = 1$, for, otherwise, ϕ would eventually leave $\mathcal{D}$. Thus $m = (\dot{L}/L)^{-1}(\lambda_\infty) = M_{\lambda_\infty}$, as claimed.

Case 2 Suppose next that F has a differentiably strictly increasing breakthrough rate with $\lim_{t\to\infty}[\dot{F}/(1-F)](t) = \infty$. Then, according to Proposition 3, the maturation delay $t - \phi(t)$ decreases to a limit m as t goes to infinity. One cannot have $m > 0$, for, otherwise, according to (9)–(10), $\dot{\phi}$ would converge to zero and ϕ would eventually leave $\mathcal{D}$. Thus $m = 0$, as claimed.

Case 3 Suppose next that F has a differentiably strictly decreasing breakthrough rate with $\lim_{t\to\infty}[\dot{F}/(1-F)](t) \equiv \lambda_\infty > 0$. Then, according to Proposition 3, the maturation delay $t - \phi(t)$ increases to a limit m as t goes to infinity. It then follows from (9)–(10) that $\dot{\phi}(t)$ converges to $\dot{\phi}(\infty) \equiv (1/\lambda_\infty)(\dot{L}/L)(m)$ as t goes to infinity. But one must then have $\dot{\phi}(\infty) = 1$, for, otherwise, ϕ would eventually leave $\mathcal{D}$. Thus $m = (\dot{L}/L)^{-1}(\lambda_\infty) = M_{\lambda_\infty}$, as claimed.

Case 4 Suppose finally that F has a differentiably strictly decreasing breakthrough rate with $\lim_{t\to\infty}[\dot{F}/(1-F)](t) = 0$. Then, according to Proposition 3, the maturation delay $t - \phi(t)$ increases to a limit m as t goes to infinity. One cannot have $m < M$, for, otherwise, according to (9)–(10), $\dot{\phi}$ would diverge to infinity and ϕ would eventually leave $\mathcal{D}$. Thus $m = M = (\dot{L}/L)^{-1}(0) = M_0$, as claimed. Hence the result. ∎

Proof of Corollary 2. We show that if ϕ is a solution to (9) that does not leave $\mathcal{D}$, then

$\dot{\phi}(t)$ goes to 1 as t goes to infinity. Observe first that

$$\limsup_{t\to\infty} \dot{\phi}(t) \geq 1 \geq \liminf_{t\to\infty} \dot{\phi}(t),$$

for, otherwise, ϕ would eventually leave $\mathcal{D}$. Then suppose by way of contradiction that $\limsup_{t\to\infty} \dot{\phi}(t) > 1$. This implies that there exist $\eta > 0$ and an increasing and divergent sequence $\{t_n\}_{n\geq 0}$ such that $\dot{\phi}(t_n) = 1+\eta$ for all $n \geq 0$. Fix $\lambda > \lambda_\infty$ such that $1+\eta > \lambda/\lambda_\infty$. For each $(t,\tau) \in \mathcal{D}$ such that $\tau < t - M_\lambda$, if follows from (10) along with the fact that the breakthrough rate of F eventually converges to λ_∞ that

$$f(t,\tau) = \frac{1-F(\tau)}{\dot{F}(\tau)} \frac{\dot{L}(t-\tau)}{L(t-\tau)} < \frac{1}{\lambda_\infty + \zeta(\tau)} \frac{\dot{L}(M_\lambda)}{L(M_\lambda)} = \frac{\lambda}{\lambda_\infty + \zeta(\tau)} < 1+\eta$$

as long as τ is large enough, where $\zeta \equiv \dot{F}/(1-F) - \lambda_\infty$ asymptotically vanishes. Thus, as $f(t_n, \phi(t_n)) = \dot{\phi}(t_n) = 1+\eta$ for all n, we get in particular that $\phi(t_n) \geq t_n - M_\lambda$ for n large enough. But observe that for each $(t,\tau) \in \mathcal{D}$ such that $\tau \geq t - M_\lambda$,

$$f(t,\tau) \geq \frac{\lambda}{\lambda_\infty + \zeta(\tau)} > 1$$

as long as τ is large enough. It follows that any solution to (9) starting at some $(t,\tau) \in \mathcal{D}$ such that $\tau \geq t - M_\lambda$ will eventually leave $\mathcal{D}$ through its upper boundary $\tau = t$. As such is the case of ϕ, we get a contradiction. Hence $\limsup_{t\to\infty} \dot{\phi}(t) = 1$. The proof that $\liminf_{t\to\infty} \dot{\phi}(t) = 1$ is similar, and uses the fact that if there exist $\eta > 0$ and an increasing and divergent sequence $\{t_n\}_{n\geq 0}$ such that $\dot{\phi}(t_n) = 1 - \eta$ for all $n \geq 0$, one can then fix $\lambda < \lambda_\infty$ such that $1 - \eta < \lambda/\lambda_\infty$ to show that ϕ will eventually leave $\mathcal{D}$ through its lower boundary $\tau = t - M$. The claim follows. From (9)–(10), a direct implication of this claim is that $(\dot{L}/L)(t - \phi(t))$ converges to λ_∞ and $t - \phi(t)$ converges to $M_{\lambda_\infty} \in (0, M)$ as t goes to infinity. Hence the result. ■

Proof of Corollary 3. Suppose by way of contradiction that $\sigma_1(\tau_0) \geq \sigma_2(\tau_0)$ for some $\tau_0 \geq 0$. Then, because $L/\dot{L}$ is strictly increasing over $[0, M)$, it follows from (10), (17), and (41) that

$$\dot{\sigma}_1(\tau_0) = \frac{\dot{F}_1(\tau_0)}{1 - F_1(\tau_0)} \frac{L(\sigma_1(\tau_0) - \tau_0)}{\dot{L}(\sigma_1(\tau_0) - \tau_0)} > \frac{\dot{F}_2(\tau_0)}{1 - F_2(\tau_0)} \frac{L(\sigma_2(\tau_0) - \tau_0)}{\dot{L}(\sigma_2(\tau_0) - \tau_0)} = \dot{\sigma}_2(\tau_0), \tag{50}$$

so that $\sigma_1(\tau) > \sigma_2(\tau)$ for all $\tau > \tau_0$ close enough to τ_0. We now show that $\sigma_1 > \sigma_2$ over (τ_0, ∞). Suppose the contrary holds, and let $\tau_1 \equiv \inf\{\tau > \tau_0 : \sigma_1(\tau) \leq \sigma_2(\tau)\}$. Then $\sigma_1(\tau_1) = \sigma_2(\tau_1)$ and $\dot{\sigma}_1(\tau_1) \leq \dot{\sigma}_2(\tau_1)$. Proceeding as for (50) shows however that $\sigma_1(\tau_1) = \sigma_2(\tau_1)$ implies that $\dot{\sigma}_1(\tau_1) > \dot{\sigma}_2(\tau_1)$, a contradiction. Hence the claim. Now consider the gap function g introduced in (45), restricted to $[\tau_0, \infty)$. By the above reasoning, (46) holds with a strict inequality for all $\tau > \tau_0$, and, as in the proof of Proposition 2, g is bounded above by M. We have

$$\begin{aligned}
\dot{g}(\tau) &= \frac{\dot{F}_1(\tau)}{1-F_1(\tau)} \frac{L(\sigma_1(\tau)-\tau)}{\dot{L}(\sigma_1(\tau)-\tau)} - \frac{\dot{F}_2(\tau)}{1-F_2(\tau)} \frac{L(\sigma_2(\tau)-\tau)}{\dot{L}(\sigma_2(\tau)-\tau)} \\
&> \frac{\dot{F}_1(\tau)}{1-F_1(\tau)} \left[\frac{L(\sigma_1(\tau)-\tau)}{\dot{L}(\sigma_1(\tau)-\tau)} - \frac{L(\sigma_2(\tau)-\tau)}{\dot{L}(\sigma_2(\tau)-\tau)} \right] \\
&\geq \frac{\dot{F}_1(\tau)}{1-F_1(\tau)} g(\tau)
\end{aligned} \tag{51}$$

for all $\tau \geq \tau_0$, where the first inequality follows from the fact that $\dot{F}_1/(1-F_1) > \dot{F}_2/(1-F_2)$, and the second inequality follows along the same lines as (47). Fixing some $\varepsilon > 0$ and integrating (51) yields

$$g(\overline{\tau}) \geq g(\tau_0+\varepsilon) \exp\left(\int_{\tau_0+\varepsilon}^{\overline{\tau}} \frac{\dot{F}_1(\tau)}{1-F_1(\tau)} \, d\tau \right)$$

and thus

$$g(\tau_0+\varepsilon) \leq M \left[\frac{1-F_1(\overline{\tau})}{1-F_1(\tau_0+\varepsilon)} \right] \tag{52}$$

for all $\overline{\tau} \geq \tau_0 + \varepsilon$. Because $g(\tau_0+\varepsilon) > 0$ as $\sigma_1 > \sigma_2$ over (τ_0, ∞), letting $\overline{\tau}$ go to infinity in (52) shows that $g(\tau_0+\varepsilon) = 0$. This contradiction establishes that $\sigma_1(\tau_0) < \sigma_2(\tau_0)$ for all $\tau_0 > 0$. Hence the result. ■

Proof of Corollary 4. Suppose that F_1 has a differentiably strictly decreasing or constant breakthrough rate. Then, by Proposition 3, the maturation delay $\sigma_1(\tau) - \tau$ under F_1 is nondecreasing with respect to the breakthrough time τ. Then

$$\begin{aligned}
2\int_0^\infty [\sigma_2(\tau)-\tau][1-F_2(\tau)]\dot{F}_2(\tau)\, d\tau &> 2\int_0^\infty [\sigma_1(\tau)-\tau][1-F_2(\tau)]\dot{F}_2(\tau)\, d\tau \\
&\geq 2\int_0^\infty [\sigma_1(\tau)-\tau][1-F_1(\tau)]\dot{F}_1(\tau)\, d\tau,
\end{aligned}$$

where the first inequality follows from the fact that $\sigma_2 > \sigma_1$ over $(0, \infty)$ by Corollary 3, and the second inequality follows from the fact that the distribution $F_{(1/2),2} \equiv 1-(1-F_2)^2$ first-order stochastically dominates the distribution $F_{(1/2),1} \equiv 1-(1-F_1)^2$ along with the fact that $\sigma_1(\tau) - \tau$ is nondecreasing in τ. The proof for the case in which F_2 has a differentiably strictly decreasing or constant breakthrough rate is similar, and is therefore omitted. Hence the result. ■

Proof of Corollary 5. Suppose by way of contradiction that $\sigma_1(\tau_0) \geq \sigma_2(\tau_0)$ for some $\tau_0 \geq 0$. Then, because $L_1/\dot{L}_1$ is strictly increasing over $[0, M)$, it follows from (10), (18), and (41) that

$$\dot{\sigma}_1(\tau_0) = \frac{\dot{F}(\tau_0)}{1 - F(\tau_0)} \frac{L_1(\sigma_1(\tau_0) - \tau_0)}{\dot{L}_1(\sigma_1(\tau_0) - \tau_0)} > \frac{\dot{F}(\tau_0)}{1 - F(\tau_0)} \frac{L_2(\sigma_2(\tau_0) - \tau_0)}{\dot{L}_2(\sigma_2(\tau_0) - \tau_0)} = \dot{\sigma}_2(\tau_0),$$

so that $\sigma_1(\tau) > \sigma_2(\tau)$ for all $\tau > \tau_0$ close enough to τ_0. As in Corollary 3, we get that $\sigma_1 > \sigma_2$ over (τ_0, ∞). Now consider the gap function g introduced in (45), restricted to $[\tau_0, \infty)$. By the above reasoning, (46) holds with a strict inequality for all $\tau > \tau_0$ and, as in the proof of Proposition 2, g is bounded above by M. We have

$$\begin{aligned}
\dot{g}(\tau) &= \frac{\dot{F}(\tau)}{1 - F(\tau)} \left[\frac{L_1(\sigma_1(\tau) - \tau)}{\dot{L}_1(\sigma_1(\tau) - \tau)} - \frac{L_2(\sigma_2(\tau) - \tau)}{\dot{L}_2(\sigma_2(\tau) - \tau)} \right] \\
&> \frac{\dot{F}(\tau)}{1 - F(\tau)} \left[\frac{L_2(\sigma_1(\tau) - \tau)}{\dot{L}_2(\sigma_1(\tau) - \tau)} - \frac{L_2(\sigma_2(\tau) - \tau)}{\dot{L}_2(\sigma_2(\tau) - \tau)} \right] \\
&\geq \frac{\dot{F}(\tau)}{1 - F(\tau)} g(\tau)
\end{aligned}$$

for all $\tau \geq \tau_0$, where the first inequality follows from (18), and the second inequality follows along the same lines as (47). The remainder of the proof is as in Corollary 3. Hence the result. ■

Proof of Lemma 4. We prove equivalently that $t - \phi^a(t) \geq t - \phi^b(t)$ for all $t \in (\sigma(0), \infty)$. From (22) along with the fact that $\sigma^a(0) = \sigma^b(0) = \sigma(0)$, one has $\dot{\phi}^a(\sigma(0)) < \dot{\phi}^b(\sigma(0))$ as $\lambda^a > \lambda^b$. Hence $\phi^a(t) < \phi^b(t)$ for all t close to but strictly greater than $\sigma(0)$. We now show that $\phi^a < \phi^b$ over $(\sigma(0), \infty)$. Suppose the contrary holds, and let $t_0 \equiv \inf\{t > 0 : \phi^a(t) \geq \phi^b(t)\}$. But then $\phi^a(t_0) = \phi^b(t_0)$ and $\dot{\phi}^a(t_0) \geq \dot{\phi}^b(t_0)$, in contradiction with (22). The result follows. ■

Proof of Lemma 5. We prove equivalently that $t - \phi^a(t) > M_{\lambda^a}$ and, symmetrically, that $t - \phi^b(t) < M_{\lambda^b}$ for all $t \geq \sigma(0)$. Suppose by way of contradiction that $t - \phi^a(t) \leq M_{\lambda^a}$, and start with the case in which $t - \phi^a(t) < M_{\lambda^a}$. Then, by definition of M_{λ^a}, $(\dot{L}/L)(t - \phi^a(t)) > \lambda^a$ and thus, by (22), $\dot{\phi}^b(t) > \lambda^a/\lambda^b > 1$. Hence $t^a \equiv \inf\{s > t : s - \phi^a(s) \geq M_{\lambda^a}\}$ must be finite, for, otherwise, ϕ^b would eventually leave $\mathcal{D}$. Then $t^a - \phi^a(t^a) = M_{\lambda^a}$ and $\dot{\phi}^a(t^a) \leq 1$. By (22), this implies that $(\dot{L}/L)(t^a - \phi^b(t^a)) \leq \lambda^a$, so that $t^a - \phi^b(t^a) \geq M_{\lambda^a}$. It follows that $t^a - \phi^b(t^a) \geq M_{\lambda^a} = t^a - \phi^a(t^a)$, which contradicts Lemma 4 as $t^a > 0$. Therefore, $t - \phi^a(t) \geq M_{\lambda^a}$ for all $t \geq \sigma(0)$. To complete the proof, we must rule out the case $t - \phi^a(t) = M_{\lambda^a}$. Suppose by way of contradiction that this equality holds for some $t > \sigma(0)$. Then, by (22) along with Lemma 4, one has $\dot{\phi}^a(t) = (\dot{L}/L)(t - \phi^b(t))/\lambda^a > (\dot{L}/L)(t - \phi^a(t))/\lambda^a = 1$ so that $s - \phi^a(s) < M_{\lambda^a}$ for all s close to but strictly greater than t, and the first part of the proof applies. Finally, suppose that $\sigma(0) = M_{\lambda^a}$. Then, by (22), $\dot{\phi}^a(\sigma(0)) = 1$, so that $\dot{\phi}^b(\sigma(0)) > 1$ as shown in the proof of Lemma 4. Differentiating (22) then yields

$$\ddot{\phi}^a(\sigma(0)) = \frac{1}{\lambda^a} \frac{d}{dm}\left(\frac{\dot{L}}{L}\right)(\sigma(0))\,[1 - \dot{\phi}^b(\sigma(0))] > 0.$$

Because $\sigma(0) = M_{\lambda^a}$ and $\dot{\phi}^a(\sigma(0)) = 1$, we get that $t - \phi^a(t) < M_{\lambda^a}$ for all t close to but strictly greater than $\sigma(0)$, and the first part of the proof applies again. The proof that $t - \phi^b(t) < M_{\lambda^b}$ for all $t \geq \sigma(0)$ is similar, and is therefore omitted. This proves (23), and thus (24) by setting $\tau = 0$. The result follows. ■

Proof of Proposition 4. By Lemma 5, the solution to the system (25) starting at (M, M) does not correspond to an equilibrium. Hence, as $\mu^a(0) = \mu^b(0)$ in any equilibrium, we can restrict the study of (25) to the open square $\mathcal{M} \equiv (0, M) \times (0, M)$. Given a point $\boldsymbol{m} \equiv (m^a, m^b)$ in $\mathcal{M}$, we denote by $\boldsymbol{\mu}(\cdot, \boldsymbol{m}) : t \mapsto (\mu^a(t, \boldsymbol{m}), \mu^b(t, \boldsymbol{m}))$ the solution to (25) passing through $\boldsymbol{m}$ at $t = 0$. This solution is defined over a maximal interval $[0, t_{\max}(\boldsymbol{m}))$, where $t_{\max}(\boldsymbol{m}) \in (0, \infty]$. We need to establish that there exists some $\boldsymbol{m} \equiv (m, m)$ with $m \in (M_{\lambda^a}, M_{\lambda^b})$ such that $t_{\max}(\boldsymbol{m}) = \infty$. We will use the following notation: $\mathcal{I}$ is the segment of the diagonal in $\mathcal{M}$ joining $(M_{\lambda^a}, M_{\lambda^a})$ to $(M_{\lambda^b}, M_{\lambda^b})$; $\mathcal{J}^a$ is the segment in $\mathcal{M}$ joining $(M_{\lambda^a}, M_{\lambda^a})$ to $(M_{\lambda^b}, M_{\lambda^a})$; $\mathcal{J}^b$ is the segment in $\mathcal{M}$ joining $(M_{\lambda^b}, M_{\lambda^b})$ to $(M_{\lambda^b}, M_{\lambda^a})$; finally $\mathcal{J} \equiv \mathcal{J}^a \cup \mathcal{J}^b$. The proof consists of three steps.

Step 1 As a preliminary remark, note that any solution $\boldsymbol{\mu}(\cdot, \boldsymbol{m})$ of the system (25) starting at some point $\boldsymbol{m} \in \mathcal{R}^a \equiv ((0, M_{\lambda^b}] \times (0, M_{\lambda^a}]) \setminus \{(M_{\lambda^b}, M_{\lambda^a})\}$ is such that $t_{\max}(\boldsymbol{m}) < \infty$. Indeed, for any such $\boldsymbol{m}$, we have, according to (25), $(\partial \mu^a/\partial t)(t, \boldsymbol{m}) \leq 0$, $(\partial \mu^b/\partial t)(t, \boldsymbol{m}) \leq 0$, and $(\partial \boldsymbol{\mu}/\partial t)(t, \boldsymbol{m}) \neq (0, 0)$ for all $t \in [0, t_{\max}(\boldsymbol{m}))$. This shows that $\boldsymbol{\mu}(t, \boldsymbol{m})$ converges as t goes to $t_{\max}(\boldsymbol{m})$ to a point in the closure $\mathrm{Cl}\,\mathcal{R}^a \setminus \mathcal{J}^a$ of $\mathcal{R}^a \setminus \mathcal{J}^a$. As there is no critical point for the system (25) in $\mathrm{Cl}\,\mathcal{R}^a \setminus \mathcal{J}^a$, $t_{\max}(\boldsymbol{m})$ must be finite. Similarly, any solution $\boldsymbol{\mu}(\cdot, \boldsymbol{m})$ of the system (25) starting at some point $\boldsymbol{m} \in \mathcal{R}^b \equiv ([M_{\lambda^b}, M) \times [M_{\lambda^a}, M)) \setminus \{(M_{\lambda^b}, M_{\lambda^a})\}$ is such that $t_{\max}(\boldsymbol{m}) < \infty$. Observe, incidentally, that this provides an alternative proof of (24). Observe also, from the above proof, that any solution to the system (25) starting in $\mathcal{J} \setminus \{(M_{\lambda^b}, M_{\lambda^a})\}$ meets this set only once, at time zero.

Step 2 For each $\boldsymbol{m} \in \mathcal{I}$, set

$$t_{\mathcal{J}}(\boldsymbol{m}) \equiv \sup\{t \geq 0 : \mu^a(s, \boldsymbol{m}) \leq M_{\lambda^b} \text{ and } \mu^b(s, \boldsymbol{m}) \geq M_{\lambda^a} \text{ for all } s \in [0, t]\} \in [0, \infty].$$

Thus $t_{\mathcal{J}}(\boldsymbol{m})$ is the first time at which the trajectory $\boldsymbol{\mu}(\cdot, \boldsymbol{m})$ starting at $\boldsymbol{m} \in \mathcal{I}$ reaches $\mathcal{J}$. The case $t_{\mathcal{J}}(\boldsymbol{m}) = \infty$ corresponds to a trajectory which remains in the triangle formed by the points $(M_{\lambda^a}, M_{\lambda^a})$, $(M_{\lambda^b}, M_{\lambda^a})$, and $(M_{\lambda^b}, M_{\lambda^b})$, and thus to an equilibrium. Then suppose by way of contradiction that $t_{\mathcal{J}}(\boldsymbol{m}) < \infty$ for all $\boldsymbol{m} \in \mathcal{I}$. We must either have

$$\mu^a(t_{\mathcal{J}}(\boldsymbol{m}), \boldsymbol{m}) = M_{\lambda^b} \text{ and } \mu^b(t_{\mathcal{J}}(\boldsymbol{m}), \boldsymbol{m}) > M_{\lambda^a}$$

or

$$\mu^a(t_{\mathcal{J}}(\boldsymbol{m}), \boldsymbol{m}) < M_{\lambda^b} \text{ and } \mu^b(t_{\mathcal{J}}(\boldsymbol{m}), \boldsymbol{m}) = M_{\lambda^a},$$

because the point $(M_{\lambda^b}, M_{\lambda^a})$ is critical for the system (25) and thus cannot be reached in a finite time $t_{\mathcal{J}}(\boldsymbol{m})$ from any point $\boldsymbol{m} \in \mathcal{I}$. From Step 1, this implies in turn that $t_{\mathcal{J}}(\boldsymbol{m})$ is the unique solution of the equation $\text{Dist}\,(\boldsymbol{\mu}(t, \boldsymbol{m}), \mathcal{J}) = 0$. In other words, if $\mu^i(t, \boldsymbol{m}) = M_{\lambda^j}$ for some i, then $t = t_{\mathcal{J}}(\boldsymbol{m})$. We now prove that the function $t_{\mathcal{J}}$ is continuous over $\mathcal{I}$. Fix $\boldsymbol{m} \in \mathcal{I}$ and assume for instance that $\mu^a(t_{\mathcal{J}}(\boldsymbol{m}), \boldsymbol{m}) < M_{\lambda^b}$ and $\mu^b(t_{\mathcal{J}}(\boldsymbol{m}), \boldsymbol{m}) = M_{\lambda^a}$. Then

$$\frac{\partial(\mu^b - M_{\lambda^a})}{\partial t}(t_{\mathcal{J}}(\boldsymbol{m}), \boldsymbol{m}) = 1 - \frac{1}{\lambda^b} \frac{\dot{L}(\mu^a(t_{\mathcal{J}}(\boldsymbol{m}), \boldsymbol{m}))}{L(\mu^a(t_{\mathcal{J}}(\boldsymbol{m}), \boldsymbol{m}))} \neq 0. \tag{53}$$

Because L is twice continuously differentiable, the flow $(t, \boldsymbol{m}) \mapsto \boldsymbol{\mu}(t, \boldsymbol{m})$ associated to the system (25) is a continuously differentiable mapping (Perko (2001, Section 2.5, Theorem 1, Remark)). From (53), we can thus invoke the implicit function theorem to obtain the continuity of $t_{\mathcal{J}}$. Introduce now the mapping $\Psi : \mathcal{I} \to \mathcal{J} : \boldsymbol{m} \mapsto \boldsymbol{\mu}(t_{\mathcal{J}}(\boldsymbol{m}), \boldsymbol{m})$. Now, by construction, this mapping is continuous. Therefore, as $\mathcal{I}$ is connected, $\Psi(\mathcal{I})$ must be connected in $\mathcal{J}$. Because $\Psi(M_{\lambda^i}, M_{\lambda^i}) = (M_{\lambda^i}, M_{\lambda^i})$ for each i, this implies, given the structure of $\mathcal{J}$, that $(M_{\lambda^b}, M_{\lambda^a}) \in \Psi(\mathcal{I})$. This, however, is impossible, because, as observed above, $(M_{\lambda^b}, M_{\lambda^a})$ is critical for the system (25). This contradiction establishes that there exists $\boldsymbol{m} \in \mathcal{I}$ such that $t_{\mathcal{J}}(\boldsymbol{m}) = \infty$ and thus $t_{\max}(\boldsymbol{m}) = \infty$, and hence that a continuous equilibrium exists.

Step 3 To conclude the proof, observe that by Step 1, any continuous equilibrium must be such that the associated trajectory $\boldsymbol{\mu}(\cdot, \boldsymbol{m})$ of the system (25) remains in the interior of the triangle formed by the points $(M_{\lambda^a}, M_{\lambda^a})$, $(M_{\lambda^b}, M_{\lambda^a})$, and $(M_{\lambda^b}, M_{\lambda^b})$. Thus, according to (25), $(\partial\mu^a/\partial t)(t, \boldsymbol{m}) > 0$ and $(\partial\mu^b/\partial t)(t, \boldsymbol{m}) < 0$, so that $\boldsymbol{\mu}(t, \boldsymbol{m})$ has a limit as t goes to infinity. In turn, this limit must be $(M_{\lambda^b}, M_{\lambda^a})$, the unique critical point of the triangle. Taken together, these observations imply (i)–(iii). Hence the result. ■

Proof of Proposition 5. Suppose by way of contradiction that there exist two continuous equilibria. According to Proposition 4(ii)–(iii), this implies that there exist two distinct points $\boldsymbol{m}_1 \equiv (m_1, m_1)$ and $\boldsymbol{m}_2 \equiv (m_2, m_2)$ in $\mathcal{I}$ such that both $\boldsymbol{\mu}(t, \boldsymbol{m}_1)$ and $\boldsymbol{\mu}(t, \boldsymbol{m}_2)$ converge to $(M_{\lambda^b}, M_{\lambda^a})$ as t goes to infinity. With no loss of generality, assume that $m_1 > m_2$. Let us first observe that

$$\mu^i(t, \boldsymbol{m}_1) > \mu^i(t, \boldsymbol{m}_2), \quad t \geq 0, \quad i = a, b. \tag{54}$$

Indeed, if this were not the case, there would for instance exist some $t > 0$ such that $\mu^a(t, \boldsymbol{m}_1) = \mu^a(t, \boldsymbol{m}_2)$ and $\mu^i(s, \boldsymbol{m}_1) > \mu^i(s, \boldsymbol{m}_2)$ for all $s \in [0, t)$ and $i = a, b$. But then, because $\dot{L}/L$ is strictly decreasing over $(0, M]$,

$$\begin{aligned}
\mu^a(t, \boldsymbol{m}_1) &= m_1 + \int_0^t \frac{\partial \mu^a}{\partial t}(s, \boldsymbol{m}_1)\, ds \\
&= m_1 + \int_0^t \left[1 - \frac{1}{\lambda^a} \frac{\dot{L}(\mu^b(s, \boldsymbol{m}_1))}{L(\mu^b(s, \boldsymbol{m}_1))}\right] ds \\
&> m_2 + \int_0^t \left[1 - \frac{1}{\lambda^a} \frac{\dot{L}(\mu^b(s, \boldsymbol{m}_2))}{L(\mu^b(s, \boldsymbol{m}_2))}\right] ds \\
&= m_2 + \int_0^t \frac{\partial \mu^a}{\partial t}(s, \boldsymbol{m}_2)\, ds \\
&= \mu^a(t, \boldsymbol{m}_2),
\end{aligned}$$

which is ruled out by assumption. This contradiction establishes (54). Consider now the gap function

$$g^{a,b}(t) \equiv \frac{1}{2} \left\| \boldsymbol{\mu}(t, \boldsymbol{m}_1) - \boldsymbol{\mu}(t, \boldsymbol{m}_2) \right\|^2, \quad t \geq 0.$$

Then, for each $t \geq 0$, we have by (25)

$$\begin{aligned}
\dot{g}^{a,b}(t) &= \sum_{i=a,b} \left[\mu^i(t, \boldsymbol{m}_1) - \mu^i(t, \boldsymbol{m}_2)\right] \left[\frac{\partial \mu^i}{\partial t}(t, \boldsymbol{m}_1) - \frac{\partial \mu^i}{\partial t}(t, \boldsymbol{m}_2)\right] \\
&= \sum_{i=a,b} \frac{1}{\lambda^i} \left[\mu^i(t, \boldsymbol{m}_1) - \mu^i(t, \boldsymbol{m}_2)\right] \left[\frac{\dot{L}(\mu^j(s, \boldsymbol{m}_2))}{L(\mu^j(s, \boldsymbol{m}_2))} - \frac{\dot{L}(\mu^j(s, \boldsymbol{m}_1))}{L(\mu^j(s, \boldsymbol{m}_1))}\right],
\end{aligned}$$

which is strictly positive according to (54) and the monotonicity of $\dot{L}/L$. This proves that $g^{a,b}$ is strictly increasing and in particular that $g^{a,b}(t) > g^{a,b}(0) = m_1 - m_2 > 0$ for all $t \geq 0$. This, however, is impossible, because both $\boldsymbol{\mu}(t, \boldsymbol{m}_1)$ and $\boldsymbol{\mu}(t, \boldsymbol{m}_2)$ converge to $(M_{\lambda^b}, M_{\lambda^a})$ as t goes to infinity, and thus $g^{a,b}(t)$ converges to zero as t goes to infinity. This contradiction establishes that there exists a unique continuous equilibrium. Hence the result. ■

Proof of (26). The vector field corresponding to the system (25) is given by

$$\boldsymbol{f}(\boldsymbol{m}) = \begin{pmatrix} 1 - (1/\lambda^a)(\dot{L}/L)(m^b) \\ 1 - (1/\lambda^b)(\dot{L}/L)(m^a) \end{pmatrix}$$

at any point $\boldsymbol{m} \equiv (m^a, m^b)$ in $\mathcal{M}$. For each i, set

$$\rho^i \equiv -\frac{d}{dt}\left(\frac{\dot{L}}{L}\right)(M_{\lambda^i}) = \left(\frac{\dot{L}^2 - L\ddot{L}}{L^2}\right)(M_{\lambda^i}) \tag{55}$$

and

$$\delta \equiv \sqrt{\frac{\rho^a \rho^b}{\lambda^a \lambda^b}}. \tag{56}$$

By Assumption 1, $\delta > 0$. The Jacobian of $\boldsymbol{f}$ at its critical point $(M_{\lambda^b}, M_{\lambda^a})$ is

$$D\boldsymbol{f}(M_{\lambda^b}, M_{\lambda^a}) = \begin{pmatrix} 0 & \rho^b/\lambda^a \\ \rho^a/\lambda^b & 0 \end{pmatrix}$$

and its eigenvalues are therefore δ and $-\delta$. As $\delta > 0$, this shows that the critical point $(M_{\lambda^b}, M_{\lambda^a})$ of $\boldsymbol{f}$ is hyperbolic (Perko (2001, Section 2.6, Definition 1)). Suppose from now on that L is thrice continuously differentiable, so that $\boldsymbol{f}$ is twice continuously differentiable in the neighborhood of $(M_{\lambda^b}, M_{\lambda^a})$. Then, according to Hartman (1960, Theorem (IV)), there exists a C^1-diffeomorphism H from a neighborhood U of $(M_{\lambda^b}, M_{\lambda^a})$ onto an open set containing the origin such that H linearizes the system $\dot{\boldsymbol{m}} = \boldsymbol{f}(\boldsymbol{m})$, locally transforming it into the linear system $\dot{\boldsymbol{m}} = D\boldsymbol{f}(M_{\lambda^b}, M_{\lambda^a})\boldsymbol{m}$. Thus, for each $\boldsymbol{m_0} \in U$, one can locally write

$$H(\boldsymbol{\mu}(t, \boldsymbol{m_0})) = e^{D\boldsymbol{f}(M_{\lambda^b}, M_{\lambda^a})t} H(\boldsymbol{m_0}).$$

Now, let $\mathcal{S}$ be the stable manifold of the nonlinear system $\dot{\boldsymbol{m}} = \boldsymbol{f}(\boldsymbol{m})$ (Perko (2001, Section 2.7, Theorem)), the upper branch of which corresponds to the equilibrium trajectory (μ^a, μ^b). Then, according to Proposition 4, there exists $t_0 \geq 0$ such that $(\mu^a(t), \mu^b(t)) \in \mathcal{S} \cap U$ for all $t \geq t_0$. As H maps $\mathcal{S}$ onto the stable subspace $\{C\boldsymbol{\xi}_{-\delta} : C \in \mathbb{R}\}$ of the linear system $\dot{\boldsymbol{m}} = D\boldsymbol{f}(M_{\lambda^b}, M_{\lambda^a})\boldsymbol{m}$ associated to the eigenvalue $-\delta$, we get that there exists a nonzero scalar constant C such that for any large enough t,

$$(\mu^a(t), \mu^b(t)) = H^{-1}(\exp(-\delta t)C\boldsymbol{\xi}_{-\delta}).$$

From Proposition 4 again, along with the fact that the derivative of H^{-1} at the origin is the identity, it follows in turn that

$$\|(\mu^a(t), \mu^b(t)) - (M_{\lambda^b}, M_{\lambda^a}) - \exp(-\delta t)C\boldsymbol{\xi}_{-\delta})\| = o(\exp(-\delta t)),$$

which implies (26) upon multiplying by $\exp(\delta t)$. Hence the result. ■

Proof of Corollary 6. Observe first that, according to Proposition 4,

$$\lim_{\tau\to\infty} \sigma_2^b(\tau) - \tau = M_{\lambda_2^a} < M_{\lambda_1^a} = \lim_{\tau\to\infty} \sigma_1^b(\tau) - \tau \tag{57}$$

as $\lambda_2^a > \lambda_1^a$, so that the result holds for τ large enough. Then suppose by way of contradiction that $\sigma_1^b(\tau_0) = \sigma_2^b(\tau_0)$ for some time τ_0 or, equivalently, that

$$\phi_1^b(t_0) = \phi_2^b(t_0) \tag{58}$$

for some time t_0, where, by (22),

$$\dot{\phi}_k^j(t) = \frac{1}{\lambda_k^j} \frac{\dot{L}(t - \phi_k^i(t))}{L(t - \phi_k^i(t))}, \quad t \geq \sigma_k(0), \quad i = a, b, \quad k = 1, 2 \tag{59}$$

with $\lambda_2^a > \lambda_1^a$ and $\lambda_2^b = \lambda_1^b$. The proof consists of two steps.

Step 1 Suppose first that

$$\dot{\phi}_1^b(t_0) \geq \dot{\phi}_2^b(t_0). \tag{60}$$

Now, observe that

$$\dot{\phi}_1^a(t_0) > \dot{\phi}_2^a(t_0) \tag{61}$$

by (58)–(59), and that

$$\phi_1^a(t_0) \geq \phi_2^a(t_0) \tag{62}$$

by (59)–(60). Combining (61) with (62) yields that for some $\varepsilon^a > 0$

$$\phi_1^a(t) > \phi_2^a(t) \tag{63}$$

for all $t \in (t_0, t_0 + \varepsilon^a)$. Together with (59), this implies that

$$\dot{\phi}_1^b(t) > \dot{\phi}_2^b(t) \tag{64}$$

for all $t \in (t_0, t_0 + \varepsilon^a)$. Similarly, combining (58) with (64) yields that for some $\varepsilon^b > 0$

$$\phi_1^b(t) > \phi_2^b(t) \tag{65}$$

for all $t \in (t_0, t_0 + \varepsilon^b)$. Together with (59), this implies that

$$\dot{\phi}_1^a(t) > \dot{\phi}_2^a(t) \tag{66}$$

for all $t \in (t_0, t_0 + \varepsilon^b)$. More generally, as long as (63) and (65) hold, so do (64) and (66). As a result, (63) and (65) hold for all $t > t_0$. This, however, is impossible, because, according to (57), $t - \phi_1^b(t) > t - \phi_2^b(t)$ for t large enough.

Step 2 The contradiction obtained in Step 1 implies that if (58) holds, then

$$\dot{\phi}_1^b(t_0) < \dot{\phi}_2^b(t_0). \tag{67}$$

It is easily checked that $\phi_1^b(t_0) = \phi_2^b(t_0) > 0$. Indeed, if one had $\phi_1^b(t_0) = \phi_2^b(t_0) = 0$, then it would follow from Lemma 1(iii) that $\phi_1^a(t_0) = \phi_2^a(t_0) = 0$, and thus by (59) that $\dot{\phi}_1^b(t_0) = \dot{\phi}_2^b(t_0)$, which is ruled out by (67). Thus, in particular, $t_0 > \sigma_1(0) \vee \sigma_2(0)$. A further implication of Step 1 is that ϕ_1^b and ϕ_2^b cannot cross over $[\sigma_1(0) \vee \sigma_2(0), t_0)$, for, given (67), this would imply the existence of some time t_1 in this interval such that $\phi_1^b(t_1) = \phi_2^b(t_1)$ and $\dot{\phi}_1^b(t_1) \geq \dot{\phi}_2^b(t_1)$, which is ruled out by Step 1. As a result, $\sigma_1(0) < \sigma_2(0)$ and

$$\phi_2^b(t) < \phi_1^b(t) \tag{68}$$

for all $t \in (\sigma_1(0), t_0)$, where $\phi_2^b \equiv 0$ over $(\sigma_1(0), \sigma_2(0))$. Now, observe that

$$\phi_1^a(t_0) < \phi_2^a(t_0) \tag{69}$$

by (59) and (67). Because $\sigma_1(0) < \sigma_2(0)$, however,

$$\phi_2^a(t) < \phi_1^a(t) \tag{70}$$

for all $t \in (\sigma_1(0), \sigma_2(0))$, where $\phi_2^a \equiv 0$ over $(\sigma_1(0), \sigma_2(0))$. It follows from (69)–(70) that there must exist some time t_1 in $(\sigma_1(0), t_0)$ such that $\phi_1^a(t_1) = \phi_2^a(t_1)$ and $\dot{\phi}_1^a(t_1) \leq \dot{\phi}_2^a(t_1)$. Together with (59), this last inequality however implies that $\phi_1^b(t_1) < \phi_2^b(t_1)$, which is ruled out by (68). This contradiction establishes that there is no time t_0 such that (58) holds. Hence the result. ■

Proof of Corollary 7. Observe first by Proposition 4(ii) that a change in λ^a does not affect the limit M_{λ^b} of $\mu^a(t)$ as t goes to infinity. It thus follows from (26) that it is sufficient to study how the eigenvalue δ defined in (56) varies with λ^a: a higher value of δ translates into a faster convergence of $\mu^a(t)$ to M_{λ^b} as t goes to infinity and thus, as $\mu^a(t) < M_{\lambda^b}$ for all t by Proposition 4(i), into asymptotically longer maturation delays. From (55)–(56) and the definition of M_λ, we just need to study the variations of the mapping

$$\lambda \mapsto -\frac{G'(G^{-1}(\lambda))}{\lambda} = -\frac{1}{\lambda(G^{-1})'(\lambda)},$$

where $G \equiv \dot{L}/L$. The derivative of this mapping has the same sign as

$$(G^{-1})'(\lambda) + \lambda(G^{-1})''(\lambda).$$

When L is given by (1), $G^{-1}(\lambda) = M_\lambda$ is given by (15), so that

$$\begin{aligned}
\operatorname{sgn}\left((G^{-1})'(\lambda) + \lambda (G^{-1})''(\lambda)\right) &= \operatorname{sgn}\left(\frac{1}{\lambda + r} - \frac{1}{\lambda + r - \mu} + \lambda\left[\frac{1}{(\lambda + r - \mu)^2} - \frac{1}{(\lambda + r)^2}\right]\right) \\
&= \operatorname{sgn}\left(\lambda\left(\frac{1}{\lambda + r - \mu} + \frac{1}{\lambda + r}\right) - 1\right) \\
&= \operatorname{sgn}\left(\lambda^2 - r(r - \mu)\right).
\end{aligned}$$

When L is given by (3), $G^{-1}(\lambda) = M_\lambda$ is given by (16), so that

$$\begin{aligned}
\operatorname{sgn}\left((G^{-1})'(\lambda) + \lambda (G^{-1})''(\lambda)\right) &= \operatorname{sgn}\left(-\frac{1}{(\lambda + r)^2} + \frac{2\lambda}{(\lambda + r)^3}\right) \\
&= \operatorname{sgn}(\lambda - r).
\end{aligned}$$

Hence the result. $\blacksquare$

References

[1] Abreu, D., and M.K. Brunnermeier (2003): "Bubbles and Crashes," *Econometrica*, 71(1), 173–204.

[2] Anderson, S.T., D. Friedman, and R. Oprea (2010): "Preemption Games: Theory and Experiment," *American Economic Review*, 100(4), 1778–1803.

[3] Arozamena, L., and E. Cantillon (2004): "Investment Incentives in Procurement Auctions," *Review of Economic Studies*, 2004, 71(1), 1–18.

[4] Athey, S. (2001): "Single Crossing Properties and the Existence of Pure Strategy Equilibria in Games of Incomplete Information," *Econometrica*, 69(4), 861–889.

[5] Aumann, R.J., and M. Kurz (1977): "Power and Taxes," *Econometrica*, 45(5), 1137–1161.

[6] Bobtcheff, C., and T. Mariotti (2012): "Potential Competition in Preemption Games," *Games and Economic Behavior*, 75(1), 53–66.

[7] Borjas, G.J., and K.B. Doran (2012): "The Collapse of the Soviet Union and the Productivity of American Mathematicians," *Quarterly Journal of Economics*, 127(3), 1143–1203.

[8] Brunnermeier, M.K., and J. Morgan (2010): "Clock Games: Theory and Experiments," *Games and Economic Behavior*, 68(2), 532–550.

[9] Card, D., and S. DellaVigna (2013): "Nine Facts about Top Journals in Economics," *Journal of Economic Literature*, forthcoming.

[10] Darwin, F. (1887): *The Life and Letters of Charles Darwin, Including an Autobiographical Chapter*, London: John Murray.

[11] Décamps, J.-P., and T. Mariotti (2004): "Investment Timing and Learning Externalities," *Journal of Economic Theory*, 118(1), 80–102.

[12] Desmond, A., and J. Moore (1991): *Darwin*, London: Michael Joseph.

[13] Dubois, P., J.-C. Rochet, and J.-M. Schlenker (2012): "Productivity and Mobility in Academic Research: Evidence from Mathematicians," Working Paper 12–307, Toulouse School of Economics.

[14] Ellison, G. (2002a): "The Slowdown of the Economics Publishing Process," *Journal of Political Economy*, 110(5), 947–993.

[15] Ellison, G. (2002b): "Evolving Standards for Academic Publishing: A *q-r* Theory," *Journal of Political Economy*, 110(5), 994–1034.

[16] Evans, J.A. (2007): "Start-ups in Science: Entrepreneurs, Diverse Backing, and Novelty Outside Business," in *The Sociology of Entrepreneurship* (*Research in the Sociology of Organizations*, *Volume 25*), ed. by M. Ruef and M. Lounsbury. Bingley, UK: Emerald Group Publishing Limited, 261–305.

[17] Foncel, J., and N. Treich (2005): "Fear of Ruin," *Journal of Risk and Uncertainty*, 31(3), 289–300.

[18] Fudenberg, D., and J. Tirole (1985): "Preemption and Rent Equalization in the Adoption of New Technology," *Review of Economic Studies*, 52(3), 383–401.

[19] Fudenberg, D., and J. Tirole (1986): "A Theory of Exit in Duopoly," *Econometrica*, 54(4), 943–960.

[20] Gilbert, R.J., and R.G. Harris (1984): "Competition with Lumpy Investment," *RAND Journal of Economics*, 15(2), 197–212.

[21] Giorgi, G., and S. Komlósi (1992): "Dini Derivatives in Optimization–Part I," *Decisions in Economics and Finance*, 15(1), 3–30.

[22] Hartman, P. (1960): "On Local Homeomorphisms of Euclidean Spaces," *Boletín de la Sociedad Matemática Mexicana*, 5(2), 220–241.

[23] Hartman, P. (1964): *Ordinary Differential Equations*, New York: John Wiley & Sons.

[24] Hendricks, K. (1992): "Reputations in the Adoption of a New Technology," *International Journal of Industrial Organization*, 10(4), 663–677.

[25] Hollingsworth, J.R., and E.J. Hollingsworth (2000): "Major Discoveries and Biomedical Research Organizations: Perspectives on Interdisciplinarity, Nurturing Leadership, and Integrated Structure and Cultures," in *Practising Interdisciplinarity*, ed. by P. Weingart and N. Stehr. Toronto: University of Toronto Press, 215–244.

[26] Hopenhayn, H.A., and F. Squintani (2010): "Patent Rights and Innovation Disclosure," Unpublished Manuscrit, University of Warwick.

[27] Hopenhayn, H.A., and F. Squintani (2011): "Preemption Games with Private Information," *Review of Economic Studies*, 78(2), 667–692.

[28] Hubbard, J.H., and B.H. West (1991): *Differential Equations: A Dynamical Systems Approach. Ordinary Differential Equations*, New York: Springer.

[29] Jones, B.F. (2009): "The Burden of Knowledge and the "Death of the Renaissance Man": Is Innovation Getting Harder?" *Review of Economic Studies*, 76(1), 283–317.

[30] Jones, B.F. (2010): "Age and Great Invention," *Review of Economics and Statistics*, 92(1), 1–14.

[31] Jones, B.F., and B.A. Weinberg (2011): "Age Dynamics in Scientific Creativity," *Proceedings of the National Academy of Sciences*, 108(47), 18910–18914.

[32] Lambrecht, B., and W. Perraudin (2003): "Real Options and Preemption under Incomplete Information," *Journal of Economic Dynamics and Control*, 27(4), 619–643.

[33] Lebrun, B. (1998): "Comparative Statics in First Price Auctions," *Games and Economic Behavior*, 25(1), 97–110.

[34] Levin, S.G., and P.E. Stephan (1991): "Research Productivity Over the Life Cycle: Evidence for Academic Scientists," *American Economic Review*, 81(1), 114–132.

[35] March, J.G. (1991): "Exploration and Exploitation in Organizational Learning," *Organization Science*, 2(1), 71–87.

[36] Maskin, E., and J. Riley (2000): "Asymmetric Auctions," *Review of Economic Studies*, 67(3), 413-438.

[37] Perko, L. (2001): *Differential Equations and Dynamical Systems*, New York: Springer.

[38] Reinganum, J.F. (1981): "On the Diffusion of New Technology: A Game Theoretic Approach," *Review of Economic Studies*, 48(3), 395–405.

[39] Riley, J. (1980): "Strong Evolutionary Equilibrium and the War of Attrition," *Journal of Theoretical Biology*, 82(3), 383–400.

[40] Simon, L.K., and M.B. Stinchcombe (1989): "Extensive Form Games in Continuous Time: Pure Strategies," *Econometrica*, 57(5), 1171–1214.

[41] Simonton, D.K. (1988): *Scientific Genius: A Psychology of Science*, Cambridge, UK: Cambridge University Press.

[42] Stigler, G.J. (1963): *The Intellectual and the Market Place, and Other Essays*, Glencoe, IL: The Free Press of Glencoe.

[43] Zuckerman, H., and R.K. Merton (1972): "Age, Aging, and Age Structure in Science," in *A Sociology of Age Stratification* (*Aging and Society*, *Volume 3*), ed. by M.W. Riley, M. Johnson, and A. Foner. New York: Russell Sage Foundation, 292–356.