Available online at www.sciencedirect.com

ScienceDirect

Journal of Economic Theory 185 (2020) 104945

JOURNAL OF Economic Theory

www.elsevier.com/locate/jet

Competitive real options under private information ☆

Leandro Gorno *, Felipe S. Iachan

FGV EPGE, Praia de Botafogo 190/1100, Rio de Janeiro, RJ 22250-900, Brazil

Received 9 February 2018; final version received 2 September 2019; accepted 21 September 2019
Available online 27 September 2019

Abstract

We study a research and development race by extending the standard investment under uncertainty framework. Each firm observes the stochastic evolution of a new product's expected profitability and chooses the optimal time to release it. Firms are imperfectly informed about the state of their opponents, who could move first and capture the market. We characterize a family of priors for which the game admits a stationary equilibrium. In this case, the equilibrium is unique and can be explicitly constructed. Across games with priors in this family, there is a maximal intensity of competition that can be supported, which is a simple function of the environment's parameters. Away from this family, we offer sufficient conditions for convergence of a non-stationary equilibrium. When these hold, the intensity of competition tends to the maximal possible value. Furthermore, we develop methods that can be useful for other applications, including a modified Kolmogorov forward equation for tracking posterior beliefs and an algorithm for computing non-stationary equilibria.

JEL classification: C73; D92; G31

Keywords: Real options; Uncertainty; Investment; Learning; Competition; Private information

☆ This paper has been previously circulated as "Competition and learning in real options". We would like to thank two editors, four anonymous reviewers, Lucas Maestri, Bernard Salanié, and seminar participants at FGV EPGE, the 2015 World Congress of the Econometric Society, the 2015 SBE Meeting, the 2017 LAMES, Universidad de San Andrés, EIEF, USP, PUC, EESP, Universidad de los Andes (CHL), the Barcelona GSE Summer Forum 2019, and JEI Madrid 2019. This study was financed in part by the Coordenação de Aperfeiçoamento de Pessoal de Nível Superior - Brasil (CAPES) - Finance Code 001 - and Brazil's National Research Council (CNPq) through Chamada Universal 421313/2016-1.

* Corresponding author.
E-mail addresses: leandro.gorno@fgv.br (L. Gorno), felipe.iachan@fgv.br (F.S. Iachan).

https://doi.org/10.1016/j.jet.2019.104945

1. Introduction

Real options, such as the option to interrupt a product's development and schedule its release, lack clear contractual terms. For instance, they typically do not expire on a proper deadline, but lose a significant part of their value if a competitor moves first.

Consider a race between several firms to develop, produce, and market an autonomous car. The first marketed product gets the possibility to set-up a new industry standard, lock in key suppliers, and obtain significantly higher profits than any follower. Although technical knowledge can only accumulate and contribute to a better product, the same unambiguous evolution does not apply to expected profits. Prototyping often evidences problems in implementation. Marketing studies convey a combination of good and bad news about consumer perceptions. Suppliers might be lost and financing dry up. These issues can be addressed with additional expenses and further delay. But waiting is risky, as a competitor might move first.

The conditions of these competitors are typically only imperfectly known to each other. First, a firm does not observe the private technological achievements of opponents. Second, even for shocks that are publicly observed, as when new regulatory standards are applied to the industry, a given firm does not know how badly compromised the specific designs of competitors are. Moreover, the final decision to produce and market a product depends on several other financial assessments which are, at best, imperfectly anticipated by opponents, such as projections of the marginal impact of the new product on previous business lines.

We study this situation by extending the continuous-time real option framework. Our model features both competition and incomplete information. Each player is privately informed about the evolution of his or her expected payoffs. He or she also continuously faces the choice between exercising the option (entry) or delaying this decision. The benefit of delay originates from increments to expected profits, which involve some randomness.[1] In addition to deferred revenues, the cost of delay includes the possibility that an opponent might enter the market first and wipe out the player's profit opportunities. Beliefs about the likelihood of an opponent's entry in the future are central determinants of optimal exercise strategies.

Our main results are the following. First, we characterize the class of prior beliefs for which a stationary equilibrium exists. For each prior within this class, we show that the associated stationary equilibrium is unique and explicitly construct it. Moreover, a particular, *canonical prior* leads to the stationary equilibrium with the highest sustainable intensity of competition. We provide an explicit formula for this maximal equilibrium intensity in terms of primitives, namely the drift and volatility of each opponent's expected payoff of entry.

Second, we track the evolution of beliefs about opponents' states for priors that lead to non-stationary equilibria and provide a partial analytical characterization of these equilibria. In particular, we give conditions for convergence toward the stationary equilibrium of the game associated with the canonical prior. The analytic methods we use to obtain these results are likely to be of interest beyond competitive real options.

Last, we compute non-stationary equilibria. The algorithm we develop for this purpose jointly iterates on the forward-looking differential equations that characterize value functions and a backward-looking integral equations for beliefs. This approach allows the study of asymmetric competition and comparative dynamics across different industries, but it can also be useful in other contexts. In our setting, we illustrate how meaningful changes in the competitive environ-

[1] See Dixit and Pindyck (1994) for a canonical reference.

ment, such as providing a firm with an initial advantage, have both mechanical effects (that firm is closer to any exercise threshold) and strategic ones (opponents initially see stronger competition and respond more aggressively).

The strategic effects vary over time, often non-monotonically. The intuition is that if one's opponent is more aggressive in the initial months, one should respond more aggressively during that period because the risk of preemption is higher; however, once that initial phase passes without any entry, this constitutes evidence that the opponent was never in a particularly strong position. As such, competition weakens. Transitions can be extremely long-lived and have meaningful effects on firm value and optimal strategies. We conclude that accounting for the time varying nature of competition can be important for applied researchers and financial managers alike.

To introduce some of the main ideas in this paper, we start with an important particular case of the model. Two symmetric players compete in a race to develop a product and first enter a market. We seek to construct a symmetric stationary equilibrium. In the recursive formulation below, two objects are key for the equilibrium characterization: the value function and the beliefs about opponents' conditions.

For clarity, we look at the problem from the perspective of Player 1, who does not observe the actual level of development of Player 2 and only holds a prior F about it.

At the same time, Player 1 privately observes the evolution of his or her own expected profitability, summarized by a payoff state $X_1(t)$, and discounts the future at a rate $r > 0$. The cost of the product's introduction into the market is $K > 0$, so that $X_n(t) - K$ is the net payoff from exercise at time t, for $n = 1, 2$. If Player 2 enters the market first, the game ends and Player 1 obtains a payoff of zero. This winner-take-all feature of the game simplifies the exposition.

The state $X_n(t)$ follows

$$dX_n(t) = \mu dt + \sigma dZ_n(t),$$

where $Z_n(t)$ for $n = 1, 2$ are two standard independent Brownian motions. We assume that $\mu > 0$, focusing on the case in which longer product development processes generate, on average, higher profits. Actual increments to profitability, however, are random and can be negative, with $\sigma > 0$ representing their volatility.

In a stationary equilibrium, Player 1 conjectures a constant defeat rate, $\lambda \geq 0$. A simple extension of well-known results[2] implies that the value function, $V(x)$, satisfies the following stationary Hamilton-Jacobi-Bellman (HJB) equation:

$$rV(x) = \max\left\{\mu \frac{dV(x)}{dx} + \frac{1}{2}\sigma^2 \frac{d^2V(x)}{dx^2} - \lambda V(x),\ r(x - K)\right\}. \tag{1}$$

The maximization above is between continuation or immediate exercise, in this order. The evolution of the continuation value is the combination of the instantaneous deterministic product improvement, uncertain innovations to profitability, and the possible arrival of a defeat.

The solution features a constant threshold, $\beta > K$, so that exercise is optimal if and only if $X_1(t) \geq \beta$.[3]

Static net present value (NPV) maximization would lead to investment whenever $X_1(t) \geq K$. The optimal threshold β displays a positive wedge relative to this static criterion, due to the

2 See, for example, Mcdonald and Siegel (1986) and Dixit and Pindyck (1994). For a recent and formal treatment of one-dimensional stochastic control and optimal stopping problems in economics, see Strulovici and Szydlowski (2015).

3 This threshold satisfies $\beta = K + 1/\xi$, where $\xi = \sigma^{-2}\left(\sqrt{\mu^2 + 2\sigma^2(r + \lambda)} - \mu\right)$ is the positive root associated with the characteristic polynomial of Equation (1) when continuation is optimal.

option value of delayed entry. The defeat and the discount rates play analogous roles: an increase in either decreases the wedge by the same amount. This is consistent with a literature devoted to investment practitioners that suggests the use of an increased discount rate to account for competition.[4] By varying λ from zero to infinity, one can span degrees of competition between monopoly and full profit dissipation. One of this paper's contributions is to offer a game-theoretic foundation for that rate. Another contribution is to show that optimal exercise thresholds, even non-stationary ones, are bounded by the monopolist's and zero-NPV policies.

In equilibrium, exercise thresholds and perceived defeat rates must be mutually consistent. In particular, in a stationary equilibrium, the belief distribution about Player 2's payoff state needs to satisfy

$$-\mu \frac{dF(x)}{dx} + \frac{1}{2}\sigma^2 \frac{d^2F(x)}{dx^2} + \lambda F(x) = 0, \tag{2}$$

with support in $(-\infty, \beta)$ and boundary conditions $F(\beta) = 1$ and $dF(x)/dx|_{x=\beta} = 0$.[5]

We derive this modified Kolmogorov forward equation for (stationary) conditional beliefs in Section 3.2 and offer for now only a preview of its intuition. The interpretation of the first two terms is standard: a positive drift makes it less likely that the state is below any given value as time passes, while the diffusion component leads to a smoothing of the distribution over time. The novelty lies in the last term, which originates from conditioning on the absence of defeat. As time passes and Player 2 is expected to cross the exercise threshold at a rate λ, the conditional probability of his or her state being below any $x < \beta$ (given that defeat was not observed) increases proportionately at that same rate. Intuitively, the absence of defeat is good news for Player 1: had Player 2 been close to the threshold, he or she would have been relatively more likely to enter the market. In this game, survival is indicative of a relatively weaker opponent than previously thought.

We show that Equation (2) admits a single (prior) probability distribution as a solution for any $\lambda \in \left(0, \lambda^*\right]$, where $\lambda^* \equiv \frac{1}{2}\frac{\mu^2}{\sigma^2}$ is the highest level of perceived competition that can occur in a stationary equilibrium. A key consequence is that, for each $\lambda \in \left(0, \lambda^*\right]$, the game in which the prior marginal distribution about the opponent's condition satisfies Equation (2) has a stationary equilibrium with the value function determined by Equation (1). Also, if the prior marginal distribution does not satisfy Equation (2) for any $\lambda \in \left(0, \lambda^*\right]$, no stationary equilibrium exists and a more general approach is required.

In the rest of the paper, we go beyond the stationary case and consider a flexible model, which allows for multiple asymmetric players and arbitrary priors.

2. Model

2.1. Description of the game

Time is continuous and the horizon is infinite. Players are indexed by $n \in \boldsymbol{N} \equiv \{1, 2, ..., N\}$. The discount rate is $r > 0$ for every player. Each player, $n \in \boldsymbol{N}$, privately observes the evolution

[4] It is also well known that the wedge increases in the volatility of the state (Dixit and Pindyck, 1994). For an application featuring an ad hoc discount rate increase, see Trigeorgis (1995, Chapter 9).

[5] The boundary conditions reflect the fact that stationarity is inconsistent with any positive mass above (or a non-vanishing density at) the exercise threshold (see Sections 3.2 and 3.4).

of a position $X_n(t)$, where $X_n \equiv \{X_n(t)\}_{t\geq 0}$ is a stochastic process with initial condition $X_n(0) = x_n^0$. We denote by F^0 the (common) prior distribution over the players' initial conditions. We assume that initial conditions are independent across players and denote by F_n^0 the prior marginal distribution for Player n. The evolution of the stochastic process X_n satisfies

$$dX_n(t) = \mu_n dt + \sigma_n dZ_n(t),$$

where Z_n is a Wiener process and $\mu_n > 0$ and $\sigma_n > 0$ represent constant player-specific drift and volatility. The processes $Z_1, ..., Z_N$ are independent and all parameters are common knowledge.[6]

The positions $X_1(t), ..., X_N(t)$ represent the development state of different projects, measured as a gross expected payoff from current exercise. Their evolution is private information, so each player knows his or her own progress, but does not know the progress of the opponents. While we are restricting attention to stochastic increments in the states that are independent across agents, the drift term can incorporate common deterministic trends in the exercise payoffs.

Each player decides at every instant whether to exercise the option or wait for more information. If Player n exercises when $X_n(t) = x_n$, the game ends at time t and the player obtains a payoff of $x_n - K_n$, while the opponents get 0. We assume that the exercise cost is positive, common knowledge and that there is no running cost for staying in the game, so that waiting is optimal whenever x_n is sufficiently low. To rule out situations in which the game ends at date $t = 0$ with probability one, we introduce the following condition, which we assume throughout the paper.

Assumption. For all $n \in \boldsymbol{N}$, the prior marginal distribution satisfies $\lim_{x_n \uparrow K_n} F_n^0(x_n) > 0$.

2.2. *Information, strategies, and payoffs*

For each $n \in \boldsymbol{N}$, let $\mathcal{F}_n \equiv \{\mathcal{F}_n(t)\}_{t\geq 0}$ be the filtration generated by X_n. A strategy for Player n is a $\mathcal{F}_n-$stopping time, generically denoted τ_n. We allow stopping times to be infinite when a player never exercises (receiving a payoff of 0).

Let $\mathcal{F}$ be the product filtration jointly generated by $X_1, ..., X_N$. Notice that $\mathcal{F}$ contains more information than observed by each player individually. The game ends as soon as any player exercises, that is, at the $\mathcal{F}-$stopping time $\min_{n\in \boldsymbol{N}} \tau_n$. Player n can only observe the passage of time, the absence of any opponent's exercise, and the evolution of their own position $\{X_n(t)\}_{t\geq 0}$. If a strategy for Player n is the first-passage time of X_n through a lower-semicontinuous threshold, we call it a *threshold* strategy. We say that τ_n is a *stationary* strategy if it is a time-invariant threshold strategy and satisfies $\Pr\{\tau_n = 0\} = 0$. That is, stationary strategies are first-passage times through some constant threshold.

Let $\mathcal{S}_n$ and $\mathcal{T}_n$ be the set of strategies and threshold strategies, respectively, for Player n. We also define $\tau_{[-n]} \equiv \min_{m\in \boldsymbol{N}\setminus\{n\}} \tau_m$, the minimal stopping time among Player n's opponents. As usual, the subscript $-n$ denotes strategy or strategy set profiles for the opponents of Player n.

[6] We choose an arithmetic Brownian process for both analytical tractability and the assumption that each firm's research and development (R&D) efforts generate a constant flow of expected profit innovations. The geometric Brownian case requires a simple change of variables and is discussed in Section C.1 of the appendix. Under that specification, profit innovations from further delay are proportional to current expected profits, an assumption that we find less appealing for R&D applications.

Player n's expected discounted payoff at time $t \geq 0$ of using strategy $\tau_n \geq t$ when opponents use τ_{-n} is given by

$$J_n(\tau_n, \tau_{-n}|t) \equiv \begin{cases} \mathbb{E}\left\{e^{-r(\tau_n - t)} 1_{\tau_n < \tau_{[-n]}} \left(X_n\left(\tau_n\right) - K_n\right) \middle| \mathcal{F}_n(t), \tau_{[-n]} \geq t\right\} & \text{if } \tau_{[-n]} \geq t, \\ 0 & \text{if } \tau_{[-n]} < t. \end{cases}$$

There are three features of the expected discounted payoffs worthy of attention. First, if two players ever exercise at exactly the same time, they both collect a payoff of zero. The implicit assumption is that $X_n\left(\tau_n\right) - K_n$ represents the payoff that a monopolist would obtain and any other arrangement, with multiple players competing to sell their products, leads to complete dissipation of market power.

Second, notice that, besides the information from the filtration generated by X_n, Player n at any particular moment also knows whether the game has not yet ended with her defeat.

Third, notice that, for any profile of strategies of opponents, the value process

$$\sup_{\tau_n \in \mathcal{S}_n | \tau_n \geq t} J_n(\tau_n, \tau_{-n}|t)$$

is a Markov process with a private state that contains both $X_n\left(t\right)$ and the knowledge of whether any of the opponents has stopped before the current date t.

2.3. Equilibrium

The following definition introduces the equilibrium notions employed in the rest of the paper.

Definition 1. A (Nash) *equilibrium* is a strategy profile $\hat{\tau} = (\hat{\tau}_1, ..., \hat{\tau}_N) \in \prod_{n=1}^{N} \mathcal{S}_n$ such that $J_n(\hat{\tau}_n, \hat{\tau}_{-n}|0) \geq J_n(\tau_n, \hat{\tau}_{-n}|0)$ for all $\tau_n \in \mathcal{S}_n$ and $n \in \boldsymbol{N}$. A *stationary equilibrium* is an equilibrium in stationary strategies.

In equilibrium, each strategy $\hat{\tau}_n$ maximizes the expected discounted payoffs of Player n, holding strategies $\hat{\tau}_{-n}$ fixed for all other players.

Note that, from the viewpoint of Player n, the behavior of all opponents is effectively summarized by the distribution of the time of Player n's defeat, which is determined by $\tau_{[-n]}$. Moreover, the optimal stopping problem arising from any such distribution is solved by a threshold strategy. This means that threshold strategies are enough for each player to best respond, even to opponents playing in arbitrary ways. More formally, despite the strict inclusion $\mathcal{T}_n \subset \mathcal{S}_n$, we have

$$\max_{\tau_n \in \mathcal{T}_n} J_n(\tau_n, \tau_{-n}|0) = \sup_{\tau_n \in \mathcal{S}_n} J_n(\tau_n, \tau_{-n}|0)$$

for all $\tau_{-n} \in \mathcal{S}_{-n}$ and $n \in \boldsymbol{N}$.[7] The bottom line is that, for the purposes of equilibrium analysis, we can restrict attention to threshold strategies without loss of generality.

7 For details, please refer to Section S4 (Lemma 31, in particular) in the supplementary material.

2.4. A recursive representation

Fix an equilibrium $\hat{\tau} \equiv \left(\hat{\tau}_1, ..., \hat{\tau}_N\right)$. Let $V_n(x_n, t)$ be the equilibrium payoff of Player n at state $X_n(t) = x_n$ conditional on the knowledge that opponents have not stopped before $t \geq 0$, that is,

$$V_n(x_n, t) \equiv \sup_{\tau_n \in \mathcal{S}_n | \tau_n \geq t} \mathbb{E}\left\{ e^{-r(\tau_n - t)} 1_{\tau_n < \hat{\tau}_{[-n]}} (X_n(\tau_n) - K_n) \middle| X_n(t) = x_n, \hat{\tau}_{[-n]} \geq t \right\}. \tag{3}$$

Standard arguments show that $V_n(x_n, t)$ is increasing and convex in x_n. Moreover, since the option to stop is always available, the value function must satisfy $V_n(x_n, t) \geq x_n - K_n$ for all $x_n \in \mathbb{R}$. These properties imply that the value function induces an optimal exercise threshold

$$\beta_n(t) \equiv \sup\left\{ x_n \in \mathbb{R} | V_n(x_n, t) > x_n - K_n \right\}. \tag{4}$$

For notation simplicity, let us leave implicit the dependence on the state (x_n, t) and write V_n to represent $V_n(x_n, t)$. Whenever the distribution of $\hat{\tau}_{[-n]}$ is absolutely continuous, its hazard rate, λ_n, defines the *equilibrium defeat rate* of Player n and the associated Hamilton-Jacobi-Bellman (HJB) equation is

$$rV_n = \max\left\{ \mu_n \frac{\partial V_n}{\partial x_n} + \frac{1}{2}\sigma_n^2 \frac{\partial^2 V_n}{\partial x_n^2} + \frac{\partial V_n}{\partial t} + \lambda_n(t)(0 - V_n),\ r(x_n - K_n) \right\}. \tag{5}$$

In other words, $\lambda_n(t)$ is the arrival rate of the end of the game induced by the equilibrium exercise from any of the opponents of Player n, conditional on the game not having ended. The first term inside the maximization is the value of continuation and the second one represents the value from current exercise. On the former, one can notice, in order, the effects from the drift in the process $X_n(t)$, the volatility, the time dependence, and the possibility of the game ending with defeat, which induces a instantaneous jump to zero in the continuation value. Notice that all the information about opponents that is necessary to solve one's optimization problem is summarized by the function λ_n. Also, the time dependence of the value function originates exclusively from the defeat rate: whenever λ_n is constant, the value function is stationary.

Note that, in order for the HJB to be well-defined in a classic sense, the value function V_n must be smooth enough. If these conditions hold, the HJB equation is solved as a free-boundary problem of the partial differential equation (PDE)

$$[r + \lambda_n(t)]\, V_n = \mu_n \frac{\partial V_n}{\partial x} + \frac{1}{2}\sigma_n^2 \frac{\partial^2 V_n}{\partial x^2} + \frac{\partial V_n}{\partial t}, \tag{6}$$

on the region $x_n < \beta_n(t)$, with free-boundary conditions given by

$$V_n(\beta_n(t), t) = \beta_n(t) - K_n \tag{7}$$

and

$$\left. \frac{\partial V_n(x_n, t)}{\partial x_n} \right|_{x_n = \beta_n(t)} = 1, \tag{8}$$

where $\beta_n(t)$ is a free-boundary, which might depend on t. Equation (7) represents the value-matching condition at the boundary, and Equation (8) is the smooth-pasting condition. To provide a formal representation result, let us say that a value-threshold pair (V_n, β_n) is *smooth* if $V_n : \mathbb{R} \times [0, \infty) \to \mathbb{R}$ and $\beta_n : [0, \infty) \to \mathbb{R}$ are continuously differentiable functions everywhere, and V_n is twice continuously differentiable in space whenever $x_n \neq \beta_n(t)$. Then, we have the following:

Proposition 1. *For each $n \in \boldsymbol{N}$, let (V_n, β_n) be a smooth value-threshold pair and let $\hat{\tau}_n$ be a $\mathcal{F}_n$−stopping time.*

i) *Suppose that $(\hat{\tau}_1, ..., \hat{\tau}_N)$ is an equilibrium that induces $(V_n, \beta_n)_{n\in \boldsymbol{N}}$ through Equations (3) and (4). Then, for each $n \in \boldsymbol{N}$, the distribution of $\hat{\tau}_{[-n]}$ has a continuous hazard rate λ_n, and (V_n, β_n) solves the free-boundary problem posed by Equations (6), (7), and (8) given λ_n.*
ii) *Suppose that $\hat{\tau}_n$ is the first-passage time of X_n through β_n. Then, the random time $\hat{\tau}_{[-n]}$ has a continuous hazard rate λ_n. Moreover, if the pair (V_n, β_n) solves the free-boundary problem posed by Equations (6), (7), and (8) given λ_n, for each $n \in \boldsymbol{N}$, then $(\hat{\tau}_1, ..., \hat{\tau}_N)$ is an equilibrium.*

Note that Proposition 1 only concerns equilibria displaying enough smoothness. As we will see in Section 3.4, the class of such equilibria includes all stationary equilibria. It is currently an open question whether there exists an equilibrium that induces a value-threshold pair that fails to be smooth. The key step to establish the second part of the proposition is the verification argument provided by Lemma 2 in the appendix.

3. Main results

3.1. Bounds on exercise thresholds

It is natural to expect the optimal behavior of a competitive player to lie somewhere between the behavior of a monopolist, who does not face the threat of any possible preemption, and the behavior under the most extreme form of competition, in which any positive NPV option is instantly exercised. These intuitive bounds imply direct restrictions on equilibrium exercise thresholds and exercise times. Proposition 2 below establishes these bounds in any equilibrium in threshold strategies by eliminating dominated strategies.

To formally state the result, define individual specific constant thresholds $\underline{\beta}_n \equiv K_n$ and $\overline{\beta}_n \equiv K_n + 1/\xi_n$, where

$$\xi_n \equiv \frac{1}{\sigma_n^2}\left(\sqrt{\mu_n^2 + 2\sigma_n^2 r} - \mu_n\right).$$

Here, $\underline{\beta}_n$ represents the perfectly competitive zero NPV threshold and $\overline{\beta}_n$ the stationary threshold that prevails for the optimal exercise of a monopolist. The number ξ_n is the positive root of $(1/2)\sigma_n^2\xi^2 + \mu_n\xi - r = 0$, the characteristic polynomial associated with the ordinary differential equation that describes the monopolist's value function in the continuation region.

Using these thresholds, we define stopping times $\underline{\tau}_n \equiv \inf\left\{t > 0 \middle| X_n(t) \geq \underline{\beta}_n\right\}$ and $\overline{\tau}_n \equiv \inf\left\{t > 0 \middle| X_n(t) \geq \overline{\beta}_n\right\}$, which represent the random times for the first crossing of the lowest (most aggressive) zero-NPV threshold and the (least aggressive) monopolistic threshold. The next result shows that the ranking of the two constant thresholds is translated to these stopping times and, more importantly, that these stopping times bound threshold strategies.

Proposition 2. *Let $(\hat{\tau}_1, ..., \hat{\tau}_N)$ be an equilibrium with associated exercise thresholds $(\beta_1, ..., \beta_N)$, following Equation (4). Then, $\underline{\tau}_n \leq \hat{\tau}_n \leq \overline{\tau}_n$ and $\underline{\beta}_n \leq \beta_n \leq \overline{\beta}_n$ for every player $n \in \boldsymbol{N}$.*

Proposition 2 is important for constraining possible equilibrium exercise thresholds and stopping times. It is especially useful in describing the long-run properties of the game, as the limited amount of rationality imposed by the bounds above is sufficient to pin down the asymptotic behavior of the rate of arrival of defeat. In fact, we provide a convergence result in Section 3.5. However, before focusing on the limit, we study how conditional belief distributions and the dynamics of competition evolve in this setting.

3.2. Equilibrium exercise densities and belief evolution

To characterize equilibria, we first resort to an intermediate result that describes the evolution of a Brownian motion density when subject to a given absorbing boundary, β_n. This result is directly related to the distribution of players' stopping times and is important for characterizing equilibrium beliefs about conditions of opponents and the likelihood of their exercise.

We denote the density of the current state for paths that have not previously hit the boundary by $f_n : \mathbb{R} \times \mathbb{R}_+ \to \mathbb{R}_+$, so that $f_n(x_n, t)$ is the density at payoff state x_n and time t. The evolution of this density is described by the following standard Kolmogorov forward equation

$$\frac{\partial f_n}{\partial t} = -\mu_n \frac{\partial f_n}{\partial x_n} + \frac{1}{2}\sigma_n^2 \frac{\partial^2 f_n}{\partial x_n^2}, \;\; for\ x_n < \beta_n(t). \tag{9}$$

On the left-hand side, we have the time evolution of the density at a state (x_n, t). The first term on the right-hand side describes how a drift imposes a lateral shift in the density: whenever $\partial f_n/\partial x_n > 0$ ($\partial f_n/\partial x_n < 0$), a given state x_n loses (gains) density in proportion to the drift μ_n. The second term originates from the volatility in process X_n, which diffuses mass over neighboring payoff states as time passes.

Importantly, this density does not integrate to one, but only to the probability that the state has not yet crossed the boundary β_n up to time t. That is,

$$\int_{-\infty}^{\beta_n(t)} f_n(x_n, t) dx_n = \Pr\{X_n(s) < \beta_n(s), \forall s \leq t\} = 1 - \Gamma_n(t),$$

where $\Gamma_n(t) \equiv \Pr\{\exists s \leq t,\ X_n(s) \geq \beta_n(s)\}$ is the *cumulative distribution of the exercise* by Player n, that is, the distribution of the first-passage time of X_n through the boundary β_n. Additionally, let γ_n be the *exercise density* of Player n (*i.e.* the density of the first-arrival time of the process X_n at the boundary β_n). It is well-known that this density exists whenever the boundary is continuously differentiable.[8]

Agents share independent common priors over their initial conditions. Let $f_n^0(x_n)$ denote the prior's generalized density over the starting point of player n (accommodating any mass points using Dirac's delta function). This density serves as the initial condition for Equation (9), so

$$f_n(x_n, 0) = f_n^0(x_n). \tag{10}$$

Given that β_n works as an absorbing boundary, the density vanishes at that boundary, implying the following boundary condition for the PDE in Equation (9):

$$f_n(\beta_n(t), t) = 0. \tag{11}$$

[8] See Lehmann (2002) for general results relating the degree of smoothness of the absorbing boundary, β_n, with that of the absorbing density, γ_n.

We use Equations (9) through (11) to characterize the probability distribution of the state, $X_n(t)$, and the exercise density γ_n. Indeed, Equations (9) and (11) imply that the following auxiliary condition is satisfied[9] at the boundary,

$$\gamma_n(t) = -\frac{1}{2}\sigma_n^2 \frac{\partial f_n(\beta_n(t), t)}{\partial x_n}. \tag{12}$$

This shows that the instantaneous absorption intensity at time t is governed by the strength of the diffusion effect and also by the slope of the density at the boundary. The intuition for this is the following: The more mass is present near the boundary (which increases with the absolute value of the slope of the density), the more mass hits it in the immediate future; also, the more randomness (higher σ_n^2) in the environment, the more movement this mass experiences and the larger is the induced absorption. In Appendix B.1, we obtain and interpret an integral representation to this backward-looking system. We use it later in the algorithm that computes non-stationary equilibria in Section 4.

Before proceeding, let us use this system for characterizing the evolution of beliefs about a player's state, conditional on absence of exercise by this player. These conditional beliefs are central to the construction of stationary equilibria of Section 3.4.

For that purpose, notice first that, while opponents do not observe the private information of Player n, they learn something from the absence of previous exercise. For instance, had a path ever been close to the boundary in the past, it would have been likely to cross it. So, the absence of a previous defeat conveys information about the relative likelihood of different paths and, consequently, about current positions.

Formally, let $\hat{f}_n : \mathbb{R} \times \mathbb{R}_+ \to \mathbb{R}_+$, defined as

$$\hat{f}_n(x_n, t) \equiv \frac{f_n(x_n, t)}{1 - \Gamma_n(t)},$$

represent the conditional belief density that opponents hold over Player n's position, $X_n(t) \leq \beta_n(t)$. We call $\hat{F}_n(\cdot, t)$ its cumulative distribution function.

From the evolution of the unconditional belief distribution (Equation (10) and (11)), it follows that

$$\frac{\partial \hat{f}_n}{\partial t} = -\mu_n \frac{\partial \hat{f}_n}{\partial x_n} + \frac{1}{2}\sigma_n^2 \frac{\partial^2 \hat{f}_n}{\partial x_n^2} + \eta_n(t)\hat{f}_n, \;\; for\ x_n < \beta_n(t), \tag{13}$$

with boundary condition $\hat{f}_n(\beta_n(t), t) = 0$ and probability preservation condition $\int_{-\infty}^{\beta_n(t)} \hat{f}_n(x, t)dx = 1$.

Here, the rescaling coefficient $\eta_n(t)$ is the instantaneous arrival rate of Player n's state to his or her boundary β_n, the *exercise rate* of that player, that can be written as

$$\eta_n(t) \equiv \frac{\gamma_n(t)}{1 - \Gamma_n(t)} = -\frac{1}{2}\sigma_n^2 \frac{\partial \hat{f}_n(\beta_n(t), t)}{\partial x_n}. \tag{14}$$

[9] A heuristic derivation is the following. Integrate Equation (9) over x_n in the region below the boundary. Then, use $F_n(\beta_n(t), t) = 1 - \Gamma_n(t)$ and $f_n(\beta_n(t), t) = 0$ to obtain

$$\frac{d(1 - \Gamma_n(t))}{dt} = \frac{1}{2}\sigma_n^2 \frac{\partial f_n(\beta_n(t), t)}{\partial x_n}.$$

.

Equation (14) illustrates an important linkage between the conditional belief distribution and the exercise rate.[10] The behavior of this conditional belief near the boundary explains the perceived threat of entry. The intuition for the effects of the density's slope and the volatility of the innovations are the same as before. Also importantly, while the unconditional exercise density, γ_n, tends to vanish as time passes, we show in Section 3.5 that η_n tends to a strictly positive limit. As a consequence, perceived competition does not vanish.

The evolution of these conditional beliefs is common knowledge. At any moment in time, as long as no option has been exercised, one can define a new game, starting from a common prior defined over initial positions, $\left\{x_n^0\right\}_{n\in \boldsymbol{N}}$, given by $\left\{F_n^0 = \hat{F}_n(\cdot, t)\right\}_{n\in \boldsymbol{N}}$. The equilibrium of this game coincides with the continuation equilibrium of the original game. That is, the environment is time homogeneous once these conditional beliefs are explicitly accounted for. We refrain from this time-homogeneous formulation, since it requires an infinite dimensional state-space encoding players' beliefs. We work instead with the non-stationary problem, by either bounding or fully characterizing the effect of time on player's payoffs and strategies.

In the next section, we relate the local intensity of defeat every player induces on his or her opponents back to the overall intensity of competition perceived by each player, which is the single input necessary for the characterization of the value function and optimal exercise strategies.

3.3. Defeat rates and optimal policy

A key ingredient in the decision problem of Player n is the perceived arrival rate of his or her defeat. In equilibrium, this perception must coincide with the conditional arrival rate of the end of the game effectively induced by the opponents of Player n. Note that, since the game is over the first time a player exercises an option, we need to find the distribution of the earliest stopping time among the opponents of Player n, that is, $\hat{\tau}_{[-n]} \equiv \min_{m\neq n} \hat{\tau}_m$. This random variable is characterized by the cumulative distribution function

$$G_{[-n]}(t) \equiv \Pr\left\{\hat{\tau}_{[-n]} \leq t\right\} = 1 - \prod_{m\neq n} (1 - \Gamma_m(t)),$$

with the associated density function given by $g_{[-n]}(t)$. The equilibrium arrival rate to the defeat of Player n, which is essential for the description of Player n's HJB equation, is

$$\lambda_n(t) \equiv \frac{g_{[-n]}(t)}{1 - G_{[-n]}(t)}.$$

Given independence of the innovations across opponents, the defeat rate of Player n is the sum of the hazard rates associated with the conditional distributions of the exercise times of Player n's opponents, that is,[11]

$$\lambda_n(t) = \sum_{m\neq n} \eta_m(t). \tag{15}$$

---

[10] Exactly as in Proposition 6, we can solve (13) and obtain an integral representation for the conditional belief and the associated arrival rate to the boundary.

[11] Notice that $\lambda_n(t) = -\frac{d}{dt}\ln\left(1 - G_{[-n]}(t)\right) = -\frac{d}{dt}\ln\left(\prod_{m\neq n}(1 - \Gamma_m(t))\right) = \sum_{m\neq n}\left(\frac{\gamma_m(t)}{1-\Gamma_m(t)}\right)$.

In loose terms, keeping strategies fixed, if one doubles the number of players, the defeat rate of any of those would double. In equilibrium, however, players' strategies respond to a potential increased competition. Section 3.5 shows that despite that strategic response, a linearity of the defeat rate in the total number of opponents is still true in the limit.

In Appendix B.2, we provide integral expressions for the threshold and the value function. In these formulations, all influence from opponents on each individual problem is summarized by an effective discount factor, which increments the discount rate (r) with the equilibrium defeat rate, following equations (14) and (15).

3.4. Stationary equilibria

In this section, we fully characterize the set of games that admit a stationary equilibrium. As we shall see, the existence of a stationary equilibrium requires very specific priors, which we explicitly parameterize using the exercise rates of the players.

Moreover, we prove uniqueness: each given game (with a fixed prior) may admit at most one stationary equilibrium. The combination of these results allows us to establish a one-to-one correspondence between the set of stationary equilibria (across different games with appropriately parametrized priors) and the set of equilibrium exercise rate profiles.

Proposition 3 below, offers the existence result.

Proposition 3. *For each vector $\overline{\eta} \in \mathbb{R}^N$, satisfying $\overline{\eta}_n \in \left(0, \frac{1}{2}\frac{\mu_n^2}{\sigma_n^2}\right]$ for all $n \in \mathbf{N}$, there exists a prior F^0 and a strategy profile $\tau = (\tau_1, ...\tau_N)$ such that:*

i) *The profile τ is a stationary equilibrium of the game under the prior F^0.*
ii) *For each $n \in \mathbf{N}$, $\overline{\eta}_n$ is the (constant) hazard rate of the distribution of τ_n.*

The proof of the result is constructive and calls attention to the shape of the prior, F^0, that supports this stationary equilibrium and the strategy profile, τ, that implements it. First, given constant exercise rates and Equation (15), defeat rates are also constant and satisfy

$$\lambda_n(t) = \overline{\lambda}_n \equiv \sum_{m \neq n} \overline{\eta}_m. \tag{16}$$

Second, with constant defeat rates, each player faces a textbook optimal stopping problem under a modified discount rate of $r + \overline{\lambda}_n$. The optimal exercise threshold of Player n ensures value matching and smooth pasting and is given by

$$\beta_n(t) = \overline{\beta}_n \equiv K_n + \frac{1}{\xi_n}, \tag{17}$$

while the associated value function is

$$V_n(x_n, t) = \overline{V}_n(x_n) \equiv \begin{cases} x_n - K_n & , \text{for } x_n \geq \overline{\beta}_n \\ \frac{e^{\xi_n(x_n - \overline{\beta}_n)}}{\xi_n} & , \text{for } x_n < \overline{\beta}_n \end{cases}, \tag{18}$$

where $\xi_n \equiv \left(\sqrt{\mu_n^2 + 2\sigma_n^2\left(r + \overline{\lambda}_n\right)} - \mu_n\right)/\sigma_n^2$.[12]

Constant exercise rates impose that the cumulative distribution of exercise is of the particular form $\Gamma_n(t) = 1 - e^{-\overline{\eta}_n t}$. In the stationary equilibrium, Γ_n is also the distribution of the first-passage time of Player n's state through the constant threshold from Equation (17). These two pieces together impose restrictions on F_n^0 and lead to the following question: given the exercise threshold $\overline{\beta}_n$, is there a prior marginal distribution over the initial state of Player n that sustains the particular first-passage distribution Γ_n? We provide an explicit positive answer in the following lemma.

Lemma 1. *For each $\overline{\eta}_n \in \left(0, \frac{1}{2}\frac{\mu_n^2}{\sigma_n^2}\right]$ and $\overline{\beta}_n$ there exists a unique prior marginal distribution F_n^0 (over the initial state $X_n(0)$) that induces $1 - \Gamma_n(t) = e^{-\overline{\eta}_n t}$. The support of F_n^0 is $(-\infty, \overline{\beta}_n]$, with its density given by*

$$f_n^0(x) = \overline{f}_n(x) = \begin{cases} 2\overline{\eta}_n e^{-\frac{\mu_n(\overline{\beta}_n - x)}{\sigma_n^2}} \dfrac{\sinh\left(\frac{(\overline{\beta}_n - x)\sqrt{\mu_n^2 - 2\overline{\eta}_n\sigma_n^2}}{\sigma_n^2}\right)}{\sqrt{\mu_n^2 - 2\overline{\eta}_n\sigma_n^2}} & , \textit{if } \overline{\eta}_n < \frac{1}{2}\frac{\mu_n^2}{\sigma_n^2}, \\ 2\overline{\eta}_n e^{-\frac{\mu_n(\overline{\beta}_n - x)}{\sigma_n^2}} \dfrac{\overline{\beta}_n - x}{\sigma_n^2} & , \textit{if } \overline{\eta}_n = \frac{1}{2}\frac{\mu_n^2}{\sigma_n^2} \end{cases} \tag{19}$$

and is the unique solution of the differential equation

$$0 = -\mu_n \frac{d\overline{f}_n}{dx_n} + \frac{1}{2}\sigma_n^2 \frac{d^2\overline{f}_n}{dx_n^2} + \overline{\eta}_n \overline{f}_n, \tag{20}$$

that satisfies the boundary condition $\overline{f}_n\left(\overline{\beta}_n\right) = 0$ and the probability preservation constraint $\int_{-\infty}^{\overline{\beta}_n} \overline{f}_n(x)\,dx = 1$.

Lemma 1 consists of two parts. Its first part shows that there is a unique distribution that ensures a given constant exercise rate against the constant threshold. Furthermore, its density is given in Equation (19).

The second part proves a modified Kolmogorov forward equation that has a straightforward economic interpretation and can be useful in other contexts. Equation (20) shows that the distribution characterized in Equation (19), for a given exercise rate $\overline{\eta}_n$, is also the stationary solution of the evolution of conditional beliefs (Eq. (13), holding that rate fixed).

There are two consequences. First, the shape of the distribution F_n^0 is such that the uninformed opponents expect Player n to exercise exactly at the constant rate $\overline{\eta}_n$. Second, after any interval of time for which exercise does not occur, the posterior opponents hold over the private state of Player n is identical to the prior. Equation (20) offers an alternative characterization of F_n^0 that sustains the constant exercise rate: one can solve the ordinary differential equation in Eq. (20), with the appropriate boundary conditions, and obtain the density of that unique distribution.

So far, our characterization of games admitting stationary equilibria is partial: given an admissible profile of exercise rates, we can specify a game and a stationary equilibrium of this

[12] It is easy to check that the (stationary) HJB equation, $r\overline{V}_n = \mu_n \frac{d\overline{V}_n}{dx_n} + \frac{1}{2}\sigma_n^2 \frac{d^2\overline{V}_n}{dx_n^2} - \overline{\lambda}_n \overline{V}_n$, holds in the continuation region and that ξ_n is the single positive root of its characteristic polynomial. Since $\overline{V}_n$ is continuously differentiable, by standard verification arguments (or the more general Proposition 1), Equation (18) is the value function for Player n.

game that implements the prescribed rates. To obtain a complete characterization, we need to determine whether there are any games that have stationary equilibria with exercise rates outside the range studied. Moreover, ruling out multiple stationary equilibria (for a given game) can also strengthen the characterization. The following proposition accomplishes both tasks.

Proposition 4. *Suppose that the strategy profile τ is a stationary equilibrium of a game (with a fixed prior F^0). Then,*

i) *τ is the unique stationary equilibrium of the game.*
ii) *The hazard rate of the distribution of each τ_n is a constant $\overline{\eta}_n \in \left(0, \frac{1}{2}\frac{\mu_n^2}{\sigma_n^2}\right]$.*
iii) *Each defeat rate is a constant $\overline{\lambda}_n$, given by Equation (16) for $\overline{\eta}_{[-n]}$ above.*
iv) *Each exercise threshold, $\overline{\beta}_n$, and value function, $\overline{V}_n$, follows Equations (17) and (18), for $\overline{\lambda}_n$ above.*
v) *Each prior marginal F_n^0 admits density in Eq. (19) with $\overline{\eta}_n$ and $\overline{\beta}_n$ given above.*

Proposition 4 concludes our characterization. There is a limited range of exercise rates that can occur in a stationary equilibrium of some game. Additionally, stationarity imposes a severe consistency requirement on priors. Since priors are predetermined and part of the description of any game, only a narrow set of games admits a stationary equilibrium. Uniqueness of the stationary equilibrium in any particular game is ensured.

It is possible to take instead an alternative perspective on the previous results. Consider an outside observer who knows all the environment of the game, except the prior. From this observer's perspective, the parameters η can be used to index exercise thresholds in Equation (17), then priors with Equation (19) and, as a consequence, fully describe a family of games and their associated stationary equilibria. Without knowledge of the prior, multiple equilibrium exercise rates can be rationalized for each player.

In this sense, the strongest prediction this observer can make is the existence of an upper bound on possible stationary exercise rates of Player n, given by $\eta_n^* = \frac{1}{2}\frac{\mu_n^2}{\sigma_n^2}$. We call these maximal rates *canonical*. In the next section, we show that the significance of canonical exercise rates extends beyond stationary equilibria: they are the long-run limit equilibrium exercise rates of a large and economically relevant set of games.

3.5. The long-run equilibrium behavior

In this section, we analytically characterize the long-run properties of equilibrium dynamics. Our main result shows that, under a differentiability assumption, equilibrium behavior and underlying beliefs converge toward a very particular steady state.

We say that the distribution of a random variable is *canonical (for Player n)* if it satisfies Equation (19) for the canonical rate ($\eta_n^* = \frac{1}{2}\frac{\mu_n^2}{\sigma_n^2}$) and the location β_n^* that characterizes the best reply to opponents' canonical exercise rates (according to Eqs. (16) and (17)). We denote this distribution by F_n^*. The canonical prior F^* is the joint distribution of the N independent random variables defined in this way, each describing the initial position of a player. Let also $\left\{V_n^*, \beta_n^*, \lambda_n^*\right\}_{n \in N}$ denote the recursive representation of the unique stationary equilibrium associated with this prior. Notice that among all stationary equilibria (across different games, induced by the particular priors characterized in the previous section), this equilibrium features the highest possible exercise rates. Also, among all distributions consistent with stationary beliefs (i.e., distributions that satisfy that Eq. (19) for some $\eta_n \in \left(0, \eta_n^*\right]$), the canonical distribution for Player n has the fastest decay in its left tail.

In what follows, we define a distribution $H:\mathbb{R} \to [0,1]$ to have *fast decay (for Player n)*, if

$$\int_{-\infty}^{0} e^{\frac{\mu_n}{\sigma_n^2}|x|} |x| \, H(dx) < +\infty.$$

Every distribution with a left tail that vanishes strictly faster than the canonical distribution (of Player n) satisfies this requirement. Important examples include degenerate distributions representing mass points (i.e., a commonly known initial conditional $X_n(0) = x_n^0$), any distribution with bounded support, and normal distributions.

To obtain our main convergence result, we restrict the prior beliefs in the following way:

Assumption 1. For every $n \in \boldsymbol{N}$, the prior marginal distribution F_n^0 is a (not necessarily strict) convex combination of the canonical and some fast-decay distribution.

We also impose the following smoothness requirement on equilibrium defeat rates.

Assumption 2. For every $n \in \boldsymbol{N}$, the defeat rate λ_n is continuously differentiable on $(0, +\infty)$ with a uniformly bounded derivative.

This assumption is trivially satisfied for stationary equilibria. The simulations in the next section suggest a wider validity. However, formally establishing sufficient smoothness of the distribution of equilibrium stopping times or finding a weaker alternative are open issues for future research.[13]

We are then able to provide an explicit description of asymptotic equilibrium behavior in terms of the exogenous parameters of the model.

Proposition 5. *Let $\{V_n, \beta_n, \lambda_n\}_{n \in N}$ be a recursive representation of an equilibrium satisfying Assumptions 1 and 2. Then, for every player $n \in \boldsymbol{N}$, we have*

i) *Values converge uniformly:* $\lim_{t\to+\infty} \sup_{x\in\mathbb{R}} |V_n(x,t) - V_n^*(x)| = 0$.
ii) *Exercise thresholds converge:* $\lim_{t\to+\infty} \beta_n(t) = \beta_n^*$.
iii) *Defeat rates converge:* $\lim_{t\to+\infty} \lambda_n(t) = \lambda_n^*$.
iv) *Conditional beliefs converge:* $\lim_{t\to+\infty} \hat{F}_n(x,t) = F_n^*(x)$ *for all* $x \in \mathbb{R}$ *and* $n \in \boldsymbol{N}$.

Proposition 5 establishes convergence and reveals the long-run determinants of equilibrium strategies and beliefs. It shows that, for a large set of priors, the importance of initial conditions vanishes and the equilibrium of the game converges to the stationary equilibrium associated with the canonical prior. Importantly, given that conditional beliefs fully summarize all public

[13] On the one hand, the distribution of any equilibrium stopping time is continuous (this result is available upon request). On the other hand, the key difficulty for a general smoothness proof is that the best reply induces distributions that fail Assumption 2. This is the case, for example, when an opponent exercises with positive probability at a given $t > 0$, a behavior that is itself inconsistent with equilibrium.

information about the past, we can say that their convergence is driving the convergence of the exercise rates and value functions.[14]

Proposition 5 has three important consequences. First, it illustrates the particular importance of the canonical prior. In the previous section, we characterized a large family of games and their stationary equilibria. The priors that supported each of these equilibria were all very particular and there was no guidance on their relative importance. Proposition 5 shows that the canonical case is the attractor of a large class of economically important games. This class plausibly exhausts all cases of interest for applied work, since priors that do not satisfy Assumption 1 require large probabilities of extremely negative initial conditions.

Second, numerical approaches to equilibrium characterization, as we implement in the next section, typically require a finite grid and the use of an artificial boundary condition after a sufficiently large horizon. Proposition 5 obtains the infinite horizon limit, which offers a natural terminal condition for an approximation.[15]

Last, given that the steady state admits a closed form, we can establish the following set of comparative statics.

Corollary 1. *An increase in μ_n or a decrease in σ_n leads to:*

i) *A decrease in limit values for all opponents $m \in \boldsymbol{N} \setminus \{n\}$ and a corresponding decrease in their limit optimal exercise thresholds β^*_m.*
ii) *A first-order stochastic dominance increase in the limit conditional beliefs about position $X_n(t)$.*
iii) *No change in the shape of limit beliefs about any $X_m(t)$ for $m \neq n$, but a first-order stochastic dominance decrease, due to the location change of β^*_m.*
iv) *An increase in the limit arrival rate of the end of the game, with an increase in the relative likelihood of exercise by player n.*

Additionally, either an increase in μ_n or in σ_n leads to an increase in Player n's own limit value function and exercise threshold, without any change in defeat rate.

Corollary 2. *The inclusion of an opponent $N+1$, with payoff drift $\mu_{N+1} > 0$ and volatility σ_{N+1} leads to*

i) *A decrease in limit values for all players $n \in \{1, ..., N\}$ and a corresponding decrease in their limit optimal exercise thresholds.*
ii) *An increase in the limit hazard rate for the end of the game of $\frac{1}{2}\left(\frac{\mu_{N+1}}{\sigma_{N+1}}\right)^2$.*

These results have consequences for the industry-wide limit dynamics. Consider, for instance, two industries with different innovation processes. The industry defined by the faster innovation processes is represented by a higher μ_n for all players. This industry becomes more competitive

[14] As discussed previously, conditional beliefs can be used as the public state in a time-independent recursive representation of equilibria.

[15] The quality of the numerical approximation depends on the choice of the artificial terminal horizon and the speed of convergence. Our results in the next section illustrate the importance of the use of a long horizon, as transitional dynamics are slow.

in the long run, effective discount rates are higher, and products are brought to market under lower profit expectations. As the value functions are forward looking, that increased competition is also propagated toward the transition phase, as we will study in the next section. A similar conclusion follows from comparing industries with different number of participants, as identified in Corollary 2.

The consequences of increased volatility of a given player n are more subtle. Higher volatility increases the option value of waiting, raising exercise thresholds and payoffs for that player. The consequences over opponents tend to be ambiguous. In principle, payoff innovation is less predictable. From the interior of the region in which player n is willing to wait, larger volatility makes him or her more likely to obtain a large sudden improvement in expected profits, leading to exercise. More formally, Equation (14) shows that for a given conditional belief about the state of this player and boundary, exercise rates increase when volatility increases. On the other hand, however, there are two forces. First, the agent becomes less aggressive in exercise thresholds. Second, the belief updating process changes. Absence of exercise informs opponents that high payoff states were unlikely, as they could have easily led to the counterfactual end of competition. In the limit, the dominant force is this, as more volatility decreases the stationary belief that opponents hold about Player n's position in a first-order stochastic dominance sense.

Indeed, increases in the uncertainty about payoff innovations tend to stir competition in the short run, while discouraging it in the long run. This is due to the offsetting nature of the effects of the increased likelihood of breakthroughs, in one direction, dominating in the short run, and information updating about the state of opponents, in the opposing direction, which dominates in the long run. We further extend this analysis and study with additional dynamic aspects of competition in the next section.

4. Simulations

In this section, we present results from simulations and comparative dynamics. First, we compute the equilibrium for a simple symmetric two-player set-up. We normalize the payoff units to set the exercise cost to unity, that is, $K_n = 1$, and the initial condition to $x_n^0 = 0$ for all players. To provide a clear meaning to time, we set the reference time unit to a year and the interest rate $r = 2\%$. We then choose the values of the drift and volatility parameters of the stochastic payoff process to match two moment conditions. The first condition is that in half of the possible histories, the firm should cross the zero NPV threshold $(X_n(t) = K_n)$ within the first two years. The second condition is that out of the remaining histories, half should cross it within the next four years. We obtain $\mu_n = 0.04$ and $\sigma_n = 0.96$.[16]

Fig. 1 plots the symmetric equilibrium exercise thresholds and the exercise rates. The dotted lines indicate the asymptotic limit of the variable on display, while the arrow on the right-hand axis marks the distance to that limit at a long eighty-year horizon. A few features are noticeable.

First, both objects display economically meaningful dynamics. At its peak, competition induces a defeat rate of almost 2.5%, which means that the effective instantaneous discount rate can be more than doubled relative to the baseline case in which competition is absent. Notice that this magnitude should get significantly larger in the presence of more opponents, a fact we ex-

[16] The evolution of the logarithm of the value function, which is comparable to an asset return, has an exposure to innovations of $\frac{\partial V_n(x,t)/\partial x}{V_n(x,t)}\sigma_n dZ_n(t)$. Near the exercise threshold, that value is approximately $\frac{1}{4}\sigma_n$.

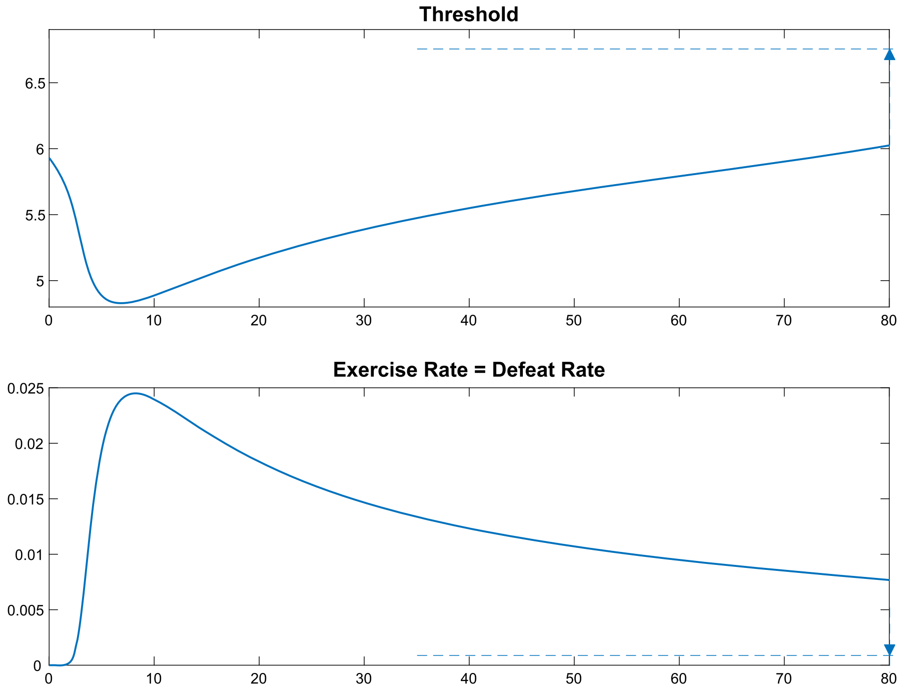

Fig. 1. Baseline Equilibrium Characterization. Symmetric parameters set to $K_n = 1$, $x_n^0 = 0$, $\mu_n = 0.04$, and $\sigma_n = 0.96$. The arrows and dotted lines mark asymptotic limits.

plore soon. The limit value of the defeat rate is to the order of 10^{-3}, so a pure study of the steady state would have concluded that competition is irrelevant quantitatively. While this depends on the drift and volatility of the calibration, it holds true for any choice that delivers projects with a significant probability of not succeeding within a window of 5 or 10 years.

Second, as the value function is forward-looking, the exercise threshold anticipates changes in the defeat rate, hitting its most aggressive point of approximately $\beta_n(t) = 4.8$ before the defeat rate reaches its peak. It then recedes toward the steady-state value of $\lim_{t\to+\infty} \beta(t) = 6.75$. For these baseline parameter values, the zero-NPV threshold is given by $\underline{\beta} = 1$, while the monopoly boundary is $\overline{\beta} = 6.9$. We can see then that the variation in the equilibrium exercise thresholds over time covers almost a third of that range. Therefore, while it is well-known that uncertainty can create a large distance between zero-NPV rules and optimal exercise, this simulation exercise shows this gap can be greatly reduced in the presence of short-term competition, while still converging close to its maximum in the long run.

Third, another striking feature of the simulation is that convergence toward the steady state is very slow. In the later phase, defeat rates display half-lives that are more than decades long. While the speed of convergence varies with parameters of the environment, this conclusion appears robust in additional explorations. There is still a meaningful effect of competition decades after its peak of intensity.

Next, we investigate and discuss comparative statics on the simulated model, with particular emphasis on heterogeneity and distinctions between partial effects, when opponents strategies are kept fixed, and the full equilibrium characterization.

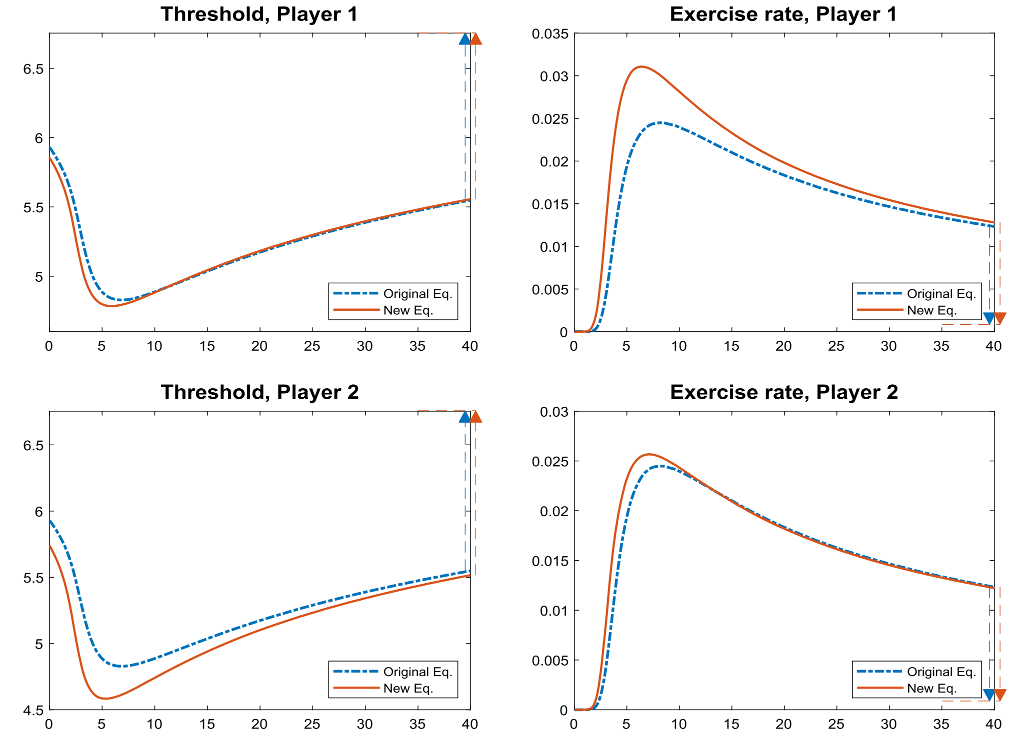

Fig. 2. Equilibrium comparison with an initial lead for Player 1. The arrows and dotted lines indicate asymptotic limits.

4.1. An initial lead

We now study the case in which Player 1 has a technological lead. She starts at $x_1^0 = 0.5$, half the original distance from zero net present value. The opponent, Player 2, still starts at $x_2^0 = 0$. The initial lead of Player 1 is common knowledge to both players, and all other parameters are kept the same as in the previous section. Fig. 2 plots the results.

A lead for Player 1 would, all else held constant, increase the defeat rate imposed on Player 2. If Player 2 did not change his or her exercise threshold, Player 1 would still be subject to the same defeat rates and would not have any incentives to change her exercise threshold, which does not depend on the initial condition. Nevertheless, as a consequence of the improved initial condition, he or she would still be more likely to hit that same threshold earlier. In the presence of a more likely early defeat, Player 2 has incentives to become more aggressive in the short run, increasing the likelihood of an early exercise. Player 1 replies to this with a more aggressive (lower) exercise threshold.

The overall consequences for the equilibrium under the new initial conditions can be seen in Fig. 2. In the equilibrium with a initial lead for Player 1, both agents behave more aggressively early on. Exercise rates increase and make the immediate end of the game more likely. Interestingly, most of the quantitative response of the equilibrium thresholds is concentrated on Player 2, since his or her defeat rate respond more strongly. The effects of the initial lead eventually vanish for both players, since the steady state does not depend on this particular initial condition.

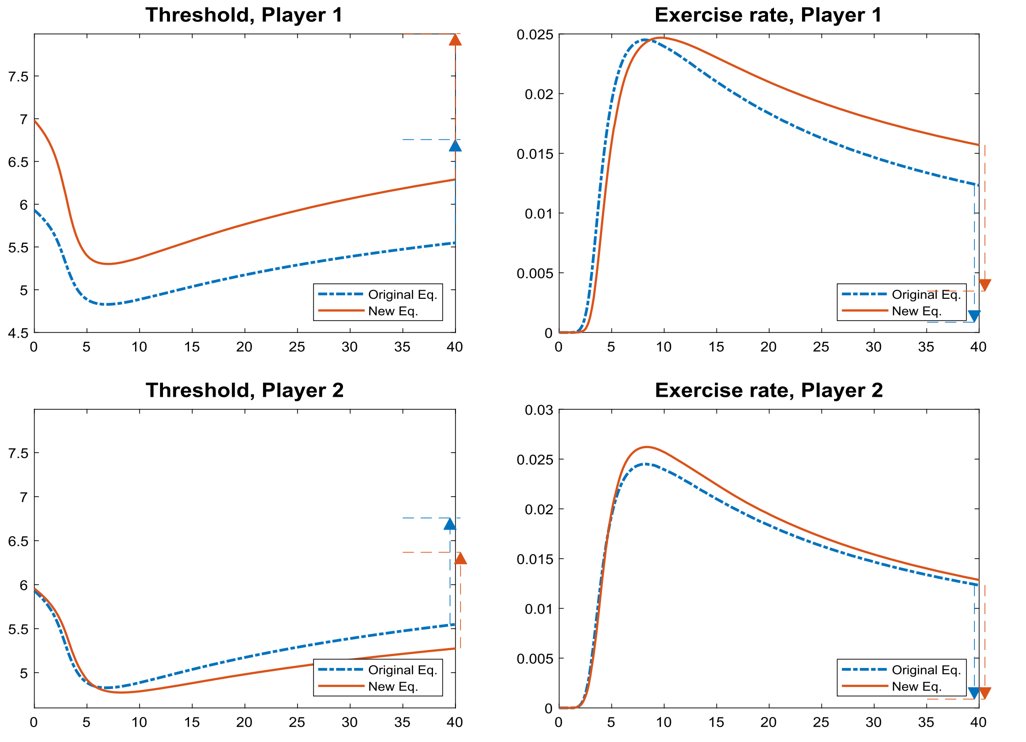

Fig. 3. Equilibrium comparison when Player 1 is subject to larger expected payoff increments. The arrows and dotted lines indicate asymptotic limits.

4.2. Faster product development

We now suppose that one player, Player 1, has faster payoff improvements than Player 2. In particular, $\mu_1 = 0.08$ is twice the benchmark rate, while $\mu_2 = 0.04$. This represents the case in which a leader is expected to reach any given level of development faster.

Given that Player 1 is subject to faster payoff improvements, he or she always has weakly higher incentives to wait instead of exercising earlier. As a consequence, we can see in the top-left panel of Fig. 3 that his or her optimal exercise threshold becomes uniformly less aggressive (higher). Two opposing forces are at play: Faster improvements increase the option value and induce the firm to be more conservative in the entry decision, but they also make sure any possible exercise trigger is reached earlier. Which of the two forces dominates depends on the horizon which is studied. As the top-right panel in Fig. 3 illustrates, in the short run, the consequences of a less aggressive exercise behavior dominate. The exercise rate lies below the symmetric original equilibrium for about the first ten years. In the long run, however, the effect of faster technological progress dominates and Player 1 imposes a more intense competition on Player 2, despite the less aggressive exercise policy.

Given this, Player 2 has incentives to behave less aggressively in the short run and more aggressively in the future. The first effect is quantitatively very small, while the second is more pronounced, as seen in Fig. 3. The equilibrium reduction of his or her threshold, after around year 7, helps partially offset the weaker deterrence incentives that a higher drift creates for Player 1.

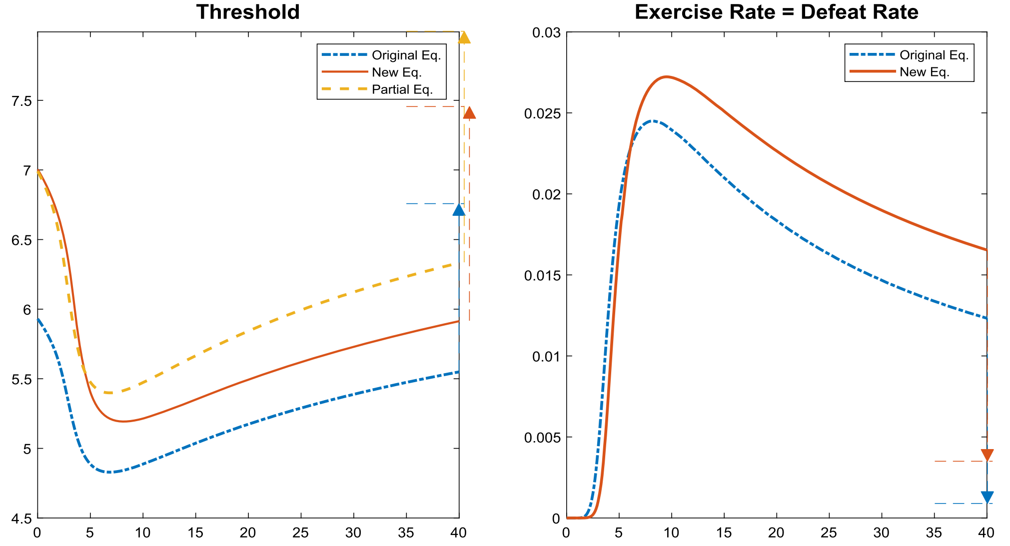

Fig. 4. Consequences of symmetric doubling of drift in the payoff process, from $\mu_n = 0.04$ (original equilibrium) to $\mu_n = 0.08$ (new equilibrium). Partial equilibrium refers to a situation in which beliefs about opponents exercise rates are kept fixed at the original equilibrium, but the new level for one's own drift is taken into account. The arrows and dotted lines indicate asymptotic limits.

In this case, unlike in the case of a simple initial lead, there are asymptotic effects. The higher drift means that, in the limit, Player 1 is more intensely pushed against her threshold. Although Player 2 replies with a threshold that converges to a higher value as a response, that has no consequences on the defeat rate that she imposes on Player 1 in the limit, which only depends on Player 2's own drift and volatility, not on the level of the asymptotic threshold, as indicated by Equation (16).

A similar logic follows if we analyze a situation in which both players have higher drifts. This comparative exercise can be used to contrast industries with different innovation dynamics. Fig. 4 illustrates this. The line labeled as partial equilibrium on the left panel studies the consequences on a firm's behavior from taking into account its own higher drift, while not internalizing the change in competition. That is, for Player 1, it keeps λ_1 (the defeat rate imposed by Player 2) fixed. Notice that an increased drift would make this firm less aggressive, as illustrated by the upward displacement of the threshold relative to the baseline (lower drift) situation.

In equilibrium, however, despite this less aggressive threshold, the higher rate of innovation increases the perceived intensity of competition. This effect, also present in the previous exercise, dampens the tendency for less aggressive behavior. The line labeled new equilibrium illustrates that industries with higher rates of innovation face higher entry cutoffs.

Section S5, in the Online Supplement, compares industries where product development is subject to different levels of risk. Again, the dynamics of competition respond in nontrivial ways: a riskier environment corresponds to an enhanced entry threat in the short run, which partially offsets the increase in option values generated by the additional uncertainty, but also to a concern for preemption that vanishes faster in the long run.

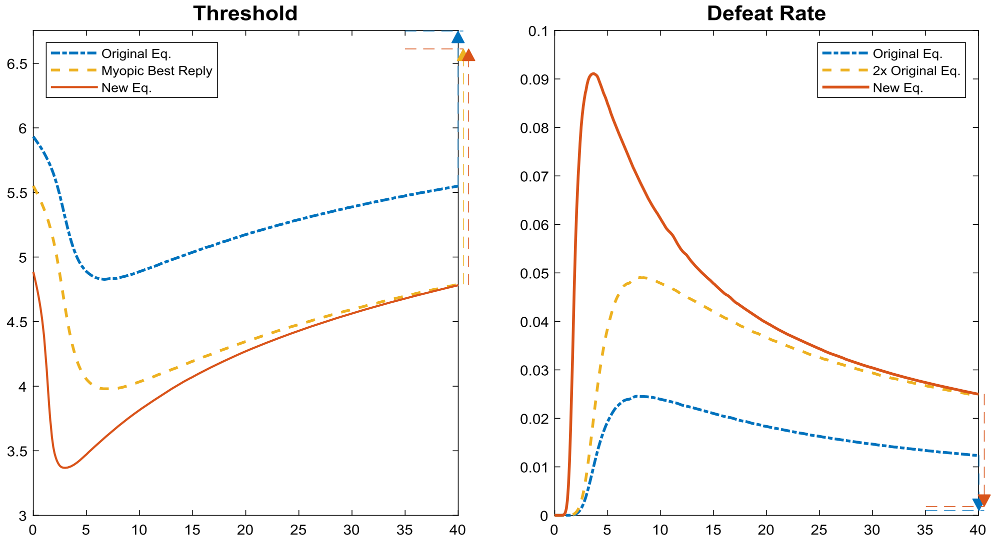

Fig. 5. The consequences from the increased number of competing players from $N = 2$ to $N = 3$. Partial equilibrium refers to a situation in which beliefs about the opponent's exercise policies are kept fixed at the original equilibrium, but the increase in the number of competitors is taken into account. The arrows and dotted lines indicate asymptotic limits.

4.3. *Increase in number of opponents*

Here, we study the consequences of increasing the number of competitors from $N = 2$ to $N = 3$.

The dashed line in the left panel of Fig. 5 illustrates a myopic approach. In this artificial situation, a player disregards the change in the strategic exercise behavior of his or her opponents, but takes into account that the presence of more players directly implies that the first passage through this fixed exercise threshold occurs earlier. Given the independence assumption regarding the payoff increments, the defeat rate for this counterfactual exercise is simply twice the original one, as each player now faces twice as many individual opponents. The best reply to that belief is to decrease exercise thresholds. Its magnitude is much larger in the long run than in the short run, as defeat rates are initially low.

The full equilibrium response is illustrated by the solid lines in Fig. 5. Notice that two effects come into play during the transition phase: amplification and anticipation. As players expect more intense competition in the future, they respond more aggressively in the present. This effect in itself increases further current exercise rates, but also propagates back to the previous dates. Amplification is noticeable from the fact that the new equilibrium threshold lies below the myopic approach, while defeat rates always lie above. Anticipation can be better noticed by looking at the troughs in the thresholds and the peaks in the new equilibrium, which occur significantly earlier than their myopic counterparts.

5. Additional discussion

In this section, we discuss important extensions and the paper's connection with a broad literature on investment in the presence of uncertainty and competition.

5.1. Relationship with the literature

This paper is related to a growing literature on dynamic contests, competitive real options, and R&D studies. In particular, the game we study belongs to the class of optimal-stopping games, as initially laid out by Dutta and Rustichini (1993), and the subclass of preemption games, notably studied in Fudenberg and Tirole (1985). Our approach can be also applied to closely related war-of-attrition and other exit games, once private information is introduced. Laraki et al. (2005) contains both a review of applications and equilibrium existence results under complete information and continuous time.

Another strand of literature applies game-theoretical insights into a real options framework. An early example is Grenadier (1996), who studies real estate market dynamics in a model with a single state variable, which all players observe.[17] We introduce two novel features into that framework. First, each firm is subject to a particular state describing its payoffs if the option is exercised. This is a natural assumption for the study of research and product development processes, but makes the problem multidimensional. Second, each firm is privately informed about the evolution of its own expected payoff, while other firms can only draw some noisy inference about that variable.[18]

The closest paper to this set-up is Hopenhayn and Squintani (2011). As ours, the model they study has both private information and one state variable for the payoff of each firm. The key distinction lies in the stochastic process driving payoffs. Hopenhayn and Squintani (2011) assume a nondecreasing process, so that exercise can only become more valuable and, due to increasing perceived competition, also more likely as time passes. Our paper is a more direct descendant of the traditional investment under uncertainty framework (Mcdonald and Siegel, 1986; Dixit and Pindyck, 1994): Payoffs follow a Brownian motion with drift, allowing also for reductions in expected profitability.

Importantly, the choice of the stochastic process driving the exercise payoffs is critical for the results and has intrinsic economic content. Hopenhayn and Squintani (2011) obtain a degree of competition that monotonically increases toward an implicit limit. Intuitively, in a set-up in which opponents constantly accumulate discrete breakthroughs, it becomes increasingly more likely that the next innovation (even if only marginal) is sufficient to lead to exercise. In the setting we study, the equilibrium threat of a competitor's entry is typically time varying and non-monotonic.

As we discussed in the introduction, allowing for bad news about profitability is natural for many economic applications. It is also essential for this non-monotonicity. The differences between the two models are particularly clear when we examine their long-run limits. In Hopenhayn and Squintani (2011), a firm that has been engaged in R&D for a sufficiently long period of time without releasing a product tends to be perceived by its competitors to be in the strongest possible position: any new breakthrough leads to an immediate launch. In the set-up we have studied, such significant delays are instead rationally interpreted as the consequence of a combination

[17] Similar environments are present in Grenadier (2002) and Weeds (2002). Grenadier (2000) provides a good review of prior work.

[18] Thijssen (2010) considers multidimensionality without private information. Lambrecht and Perraudin (2003) study an environment with a common randomly evolving payoff state and private information regarding a static exercise cost. Quah and Strulovici (2013) study an individual optimal stopping problem in the presence of non-stationary discounting. Seel and Strack (2013) consider competition in an optimal stopping problem under private information without strategic deterrence, i.e., the timing of exercise is not relevant.

of negative shocks. As a result, firms entertain the possibility that competing products long in development are actually far away from profitable release in the near future.

While we contribute to a growing literature on R&D competition, there is a complementary literature that focuses on R&D efforts within firms. For instance, Bonatti and Hörner (2011) study moral hazard in teams, with belief updates about a project's profitability, while Guo and Roesler (2018) introduce endogenous exit and the associated threat of an informed collaborator leaving the firm.[19]

Methodologically, our approach relies on a coupled system of differential equations: a forward-looking value function (or equivalently an exercise threshold) and a backward-looking evolution of beliefs about opponents. Similar coupled systems, with forward-looking value functions and backward-looking population dynamics, are studied in the growing mean-field games literature.[20] In particular, Bayraktar et al. (2018) study a R&D tournament with a continuum of players and costly efforts. The payoffs depend on the order of completion of a project, where completion occurs when the state reaches a fixed level. We see our approach as complementary, since we allow firms to choose when to market a product, creating a tension between option values and deterrence, while Bayraktar et al. (2018) focus on the intensive margin of R&D efforts.

5.2. Extensions

In Appendix C, we briefly cover multiple extensions of the model. We start by formalizing how a simple change of variable can be used to deal with an innovation process that follows a geometric Brownian motion, common in many real option applications. We also discuss how some results continue to hold for alternative payoff structures, including a less extreme assumption that followers receive some residual payoff and another assumption in which competitors face running costs. Last, we discuss the technical challenges in dealing with correlated innovations in profitability, which are left for future work.

5.3. Existence, uniqueness, and regularity for arbitrary initial conditions

In Section 3.4, we fully characterize the set of priors which are consistent with stationarity. For each prior in this class, we prove existence and uniqueness of a stationary equilibrium. Section 3.5 builds on these results. We show that, for a large class of priors, equilibria that display differentiable exercise rates converge over time to the stationary equilibrium displaying the highest possible intensity of competition. Some open questions remain.

First, existence, uniqueness, and regularity of equilibria remain to be established for arbitrary initial conditions.[21] Second, it is plausible that each initial condition that does not belong to

[19] Bobtcheff and Mariotti (2012) and Bobtcheff et al. (2016) study environments in which opponents come into play at random times, after they are enabled by a seminal technological breakthrough. Whenever active, players decide whether to release or delay a new product. Exercise payoffs evolve deterministically at that stage ("maturation"). Hopenhayn and Squintani (2015) study optimal policy in a related set up, while Dosis and Muthoo (2019) study competitive experimentation in a two-stage R&D race. By bridging the gap between this growing literature and the standard real option approach, where both good and bad news about profitability can be revealed, we facilitate the exploration of a new set of interactions between pricing, competition, information, and policy.

[20] See, for instance, Lasry and Lions (2007) and Bensoussan et al. (2013). For macroeconomic applications, relying on general equilibrium theory interactions, see Achdou et al. (2014).

[21] The main difficulty lies in proving the continuity of the distribution of the optimal stopping times with respect to opponents' strategies. One of the reasons is that establishing enough regularity of the optimal stopping threshold

the class we have considered (of distributions with bounded support) still converges to a given stationary equilibrium within the set we have exhaustively characterized. There is an active literature in applied probability, including Martinez and San Martin (1994); Martinez et al. (1998), that studies this question in non-strategic settings. The complete characterization of the mapping from priors to limit behavior in strategic settings, as ours, is a challenging topic for future research.

5.4. Conclusion

Our model naturally extends the canonical investment under uncertainty setting, incorporating private information and strategic preemption. We explicitly characterize stationary equilibria, with a particular focus on the intensity of competition that players perceive, given by a defeat rate. We also develop methods for describing the dynamics of conditional beliefs about opponents' conditions, optimal exercise strategies, and market-entry rates.

Due to their generality, these methods promise to shed light on a large class of games combining evolving information and belief dynamics. We keep the main set-up particularly simple, abstracting from important issues like price competition, the optimal intensity of R&D efforts, and strategic information revelation. We believe some extensions can fruitfully address questions related to optimal technological development policies and the value of information in technological competition.[22]

We also develop an algorithm and illustrate the applied potential from this framework by performing equilibrium computation and comparative dynamics exercises. For example, from a simple project valuation perspective, as the intensity of competition significantly changes over time and transition dynamics are very long lived, any analysis based on *ad hoc* effective discount rates can lead to large valuation errors.

Appendix A. Proofs omitted from the main text

The following verification argument is used in the proof of Proposition 1:

Lemma 2. *If* (V_n, β_n) *is a smooth value-threshold pair that solves the free-boundary problem given by Equations (6), (7), and (8), then*

$$V_n(x_n, t) = \sup_{\tau_n \in \mathcal{S}_n} J_n(\tau_n, \tau_{-n}|t)$$

for all $\tau_{-n} \in \mathcal{S}_{-n}$ *that induce the defeat rate* $\lambda_n(t)$*. Moreover, the first-passage time through* β_n *is an optimal stopping time.*

for a general non-stationary problem is hard, if not impossible. If, to tackle that issue, restrictions are imposed on the distribution of players' optimal stopping times, then the difficulty lies in establishing that the best reply is consistent with these additional restrictions.

[22] More generally, our model is a particular case in a larger class, where population dynamics and optimal stopping interact. Other instances involve equilibrium price resetting under menu costs, optimal contracting with a population of agents, and industry dynamic models with costly entry and exit. The out-of-steady-state behavior of most of these models remains largely to be explored, for instance.

Proof. The proof is an application of Theorem 1 in Brekke and Øksendall (1991). To apply the result, define $h_n(x_n,t) \equiv e^{-rt-\int_0^t \lambda_n(s)ds} V_n(x_n,t)$. Adopting the shorthand $h_n \equiv h_n(x_n,t)$, it is easy to verify that

$$\begin{aligned}&\mu_n \frac{\partial h_n}{\partial x_n} + \frac{1}{2}\sigma_n^2 \frac{\partial^2 h_n}{\partial x_n^2} + \frac{\partial h_n}{\partial t} \\ &= e^{-rt-\int_0^t \lambda_n(s)ds} \left(\mu_n \frac{\partial V_n}{\partial x_n} + \frac{1}{2}\sigma_n^2 \frac{\partial^2 V_n}{\partial x_n^2} + \frac{\partial V_n}{\partial t} - [r+\lambda_n(t)]\, V_n \right) = 0,\end{aligned}$$

for all $x_n < \beta_n(t)$ and $t > 0$. Moreover, $h_n(\beta_n(t),t) = e^{-rt-\int_0^t \lambda_n(s)ds}(\beta_n(t) - K_n)$ and $\frac{\partial h_n(x_n,t)}{\partial x}|_{x=\beta_n(t)} = e^{-rt-\int_0^t \lambda_n(s)ds}$. Condition 2 for Lemma 1 in Brekke and Øksendall (1991) holds, as X_n is uniformly elliptic and the open set $D \equiv \{(x_n,t) \in \mathbb{R} \times [0,\infty] | x_n < \beta_n(t), t > 0\}$ has a continuously differentiable boundary in $\mathbb{R} \times (0,\infty)$ with a zero Lebesgue measure spatial boundary for each fixed t. Moreover, since $\mu_n > 0$ and $\hat{\tau}_n \leq \overline{\tau}_n$ by Proposition 2, the first-exit time from D is a.s. finite. It thus follows from Theorem 1 in Brekke and Øksendall (1991) that $h_n(x_n,t) = \sup_{\tau_n \in \mathcal{S}_n} e^{-rt-\int_0^t \lambda_n(s)ds} J_n(\tau_n, \tau_{-n}|t)$ and that this value is obtained by the first-passage time through β_n. We conclude that $V_n(x_n,t) = e^{rt+\int_0^t \lambda_n(s)ds} h_n(x_n,t) = \sup_{\tau_n \in \mathcal{S}_n} J_n(\tau_n, \tau_{-n}|t)$. □

Proof of Proposition 1. For Part 1, Theorem 5 in Lehmann (2002) implies that the distribution of $\hat{\tau}_n$ has a continuous density for each $n \in \boldsymbol{N}$. The existence of continuous hazard rates thus follows from Equation (15). The smoothness assumption on V_n directly implies the boundary conditions given by Equations (7) and (8). The validity of the HJB equation in the continuation region is a standard application of Itô's lemma.

As for Part 2, existence and continuity of the hazard rates $(\lambda_1, ...\lambda_N)$ follow from the argument in Part 1. By Lemma 2, each first-passage time $\hat{\tau}_n$ is a best-response to $\hat{\tau}_{-n}$ for player $n \in \boldsymbol{N}$ after any of his or her private histories. This means that $(\hat{\tau}_1, ..., \hat{\tau}_N)$ is an equilibrium. □

Proof of Proposition 2. In the supplementary material, we prove that equilibrium value functions are increasing and convex in the state. These basic properties imply that a value matching condition holds, so that V_n and β_n satisfy $V_n(\overline{\beta}_n, t) = \overline{\beta}_n - K_n$ and

$$\beta_n(t) = \inf\{x_n \in \mathbb{R} | V_n(x_n,t) \leq x_n - K_n\}$$

It follows that $\beta_n(t) \leq \overline{\beta}_n$. Suppose, seeking a contradiction, that $\beta_n(t_0) < \underline{\beta}_n$ for some $t_0 \in \mathbb{R}_+$. Then, $V_n(\beta_n(t_0), t_0) = \beta_n(t_0) - K_n$ by value matching. Because $K_n = \underline{\beta}_n > \beta_n$, we have $V(\beta_n(t_0), t_0) < 0$. This cannot happen in equilibrium since never exercising (i.e. $\hat{\tau}_n = +\infty$) is a feasible strategy which guarantees a zero payoff. Once we have $\underline{\beta}_n \leq \beta_n \leq \overline{\beta}_n$, the inequalities for the stopping times are immediate. □

Proof of Proposition 3. The proof is constructive. Given, $\overline{\eta} = (\overline{\eta}_1, ..., \overline{\eta}_N)$, Equation (17) defines exercise thresholds $\overline{\beta} = (\overline{\beta}_1, ..., \overline{\beta}_N)$. For each $n \in \boldsymbol{N}$, Lemma 1 provides the unique prior marginal distribution F_n^0 that induces $\overline{\eta}_n$ as the hazard rate of the first-passage time of X_n through $\overline{\beta}$. Fix the prior at $F^0 = (F_1^0, ..., F_N^0)$ and let τ_n be the first-passage time of X_n through $\overline{\beta}_n$ (using F_n^0 as the distribution of $X_n(0)$). It remains to verify that $\tau \equiv (\tau_1, ..., \tau_N)$ is a stationary equilibrium. For each $n \in \boldsymbol{N}$, using the value function defined in Equation (18), we can construct a value-threshold pair $(V_n, \overline{\beta}_n)$ satisfying Equations (6), (7), and (8) given the (constant) defeat rate $\overline{\lambda}_n$ defined in Equation (16). By the second part of Proposition 1, τ is an equilibrium in threshold strategies. In fact, since the exercise thresholds used in the construction are constant, τ is a stationary equilibrium. □

Proof of Lemma 1. It is easy to show that the proposed prior marginal distribution, $\overline{F}_n$, induces the desired absorption $\overline{\eta}_n$ and that its density, $\overline{f}_n$, satisfies Equation (20) as well as the boundary condition $\overline{f}_n(\overline{\beta}_n) = 0$. It is also relatively straightforward (albeit a bit tedious) to show that no other probability density over $\left(-\infty, \overline{\beta}_n\right]$ solves Equation (20) as well as the boundary condition $\overline{f}_n(\overline{\beta}_n) = 0$.

It remains to establish that no other prior marginal distribution induces the desired absorption. For this, we adapt Proposition 1 in Jackson et al. (2009). We are interested in a distribution over $\left(-\infty, \overline{\beta}_n\right]$, with density g, such that the absorption probability over the interval $[0, t]$ is $\Gamma_n(t) = 1 - e^{-\overline{\eta}_n t}$. This is equivalent to the absorption density satisfying $\gamma_n(t) = \overline{\eta}_n e^{-\overline{\eta}_n t}$. Notice that the Laplace transform of γ_n is $L\gamma_n(s) \equiv \int_0^\infty e^{-st}\gamma_n(t)\,dt = (\overline{\eta}_n + s)^{-1}\overline{\eta}_n$. With a constant absorption boundary at $\overline{\beta}_n$, drift μ_n, and volatility σ_n^2, the first-passage time for a fixed initial condition x_n^0 has density

$$\gamma_n\left(t|x_n^0\right) = \frac{\left(\overline{\beta}_n - x_n^0\right)}{\sigma_n\sqrt{2\pi t^3}} e^{-\frac{\left(\overline{\beta}_n - x_n^0 - \mu_n t\right)^2}{2\sigma_n^2 t}} \tag{21}$$

and moment generating function

$$M_n\left(s|x_n^0\right) = \int_0^\infty e^{st}\gamma_n\left(t|x_n^0\right)dt = \exp\left(\left(\frac{\mu_n}{\sigma_n} - \sqrt{\frac{\mu_n^2}{\sigma_n^2} - 2s}\right)\left(\frac{\beta_n - x_n^0}{\sigma_n}\right)\right).$$

The first-passage time, given the initial density g, satisfies

$$\gamma_n(t) = \int_{-\infty}^{\overline{\beta}_n} \gamma_n\left(t|x_n^0\right) g\left(x_n^0\right) dx_n^0.$$

Applying the Laplace transform to the RHS, we obtain

$$\begin{aligned}
L\gamma_n(s) &= \int_0^\infty e^{-st}\left[\int_{-\infty}^{\overline{\beta}_n} \gamma_n\left(t|x_n^0\right) g\left(x_n^0\right) dx_n^0\right] dt \\
&= \int_{-\infty}^{\overline{\beta}_n}\left[\int_0^\infty e^{-st}\gamma_n\left(t|x_n^0\right)dt\right] g\left(x_n^0\right) dx_n^0, \\
&= \int_{-\infty}^{\overline{\beta}_n} M_n\left(-s|x_n^0\right) g\left(x_n^0\right) dx_n^0 = \int_{-\infty}^{\overline{\beta}_n} e^{\left(\frac{\mu_n}{\sigma_n} - \sqrt{\frac{\mu_n^2}{\sigma_n^2} + 2s}\right)\left(\frac{\overline{\beta}_n - x_n^0}{\sigma_n}\right)} g\left(x_n^0\right) dx_n^0.
\end{aligned}$$

We change spatial variables, taking $y \equiv \sigma_n^{-1}\left(\overline{\beta}_n - x_n^0\right)$ and defining $\nu(y) \equiv \sigma_n^{-1} g\left(\overline{\beta}_n - \sigma_n y\right)$, so that

$$L\gamma_n(s) = \sigma_n^2 \int_0^\infty e^{-\left(\sqrt{\frac{\mu_n^2}{\sigma_n^2}+2s}-\frac{\mu_n}{\sigma_n}\right)y} \nu(y)\,dy \equiv \sigma_n^2 L\nu(w),$$

where $L\nu$ is the Laplace transform of ν and $w \equiv \sqrt{\mu_n^2/\sigma_n^2 + 2s} - \mu_n/\sigma_n$. Solving for s to invert this last change of variables, we obtain $s = \frac{1}{2}w^2 + \frac{\mu_n}{\sigma_n}w$. Thus, we can write $L\gamma_n\left(\frac{1}{2}w^2 + \frac{\mu_n}{\sigma_n}w\right) = \sigma_n^2 L\nu(w)$ and, therefore, we have

$$L\nu(w) = \frac{1}{\sigma_n^2}\left(\frac{\overline{\eta}_n}{\overline{\eta}_n + \frac{1}{2}w^2 + \frac{\mu_n}{\sigma_n}w}\right). \tag{22}$$

Note that, using $\overline{f}_n$ as our g and defining $\overline{\nu}(y) \equiv \sigma_n^{-1}\overline{f}_n\left(\overline{\beta}_n - \sigma_n y\right)$, we obtain the transform

$$L\overline{\nu}(w) \equiv \int_0^\infty e^{-wy}\overline{\nu}(y)\,dy = \frac{1}{\sigma_n^2}\left(\frac{\overline{\eta}_n}{\overline{\eta}_n + \frac{1}{2}w^2 + \frac{\mu_n}{\sigma_n}w}\right) = L\nu(w).$$

By the invertibility of the Laplace transform, this implies that $\overline{\nu} = \nu$. Undoing the spatial change of variables, we obtain $g = \overline{f}_n$, proving the claim. □

It is easy to verify that, if the defeat rate perceived by a player is constant, then his or her optimal exercise threshold is constant. The following lemma establishes that the converse is also true.

Lemma 3. *If the optimal exercise threshold for Player n is a constant $\overline{\beta}_n$ when the defeat rate is λ_n, then $\lambda_n(t) = \overline{\lambda}_n \equiv \mu_n\left(\overline{\beta}_n - K_n\right)^{-1} + \frac{1}{2}\sigma_n^2\left(\overline{\beta}_n - K_n\right)^{-2} - r$ for all $t \geq 0$.*

Proof. Assume that the constant $\overline{\beta}_n$ is the optimal exercise threshold for Player n when his or her perceived defeat rate is λ_n. We claim that $\lambda_n(t) = \overline{\lambda}_n$ for all $t \geq 0$. Since the exercise threshold is constant, the value function can be written as:

$$V_n(x_n, t) = \left(\overline{\beta}_n - K_n\right)\int_0^\infty e^{-\rho_n(t+s,s)}\gamma_n\left(s|x_n, 0, \overline{\beta}_n\right)ds,$$

where we define the effective discount factor $\rho_n(t+s,t) \equiv \int_t^{t+s}[r + \lambda_n(h)]\,dh$ and $\gamma_n(s|x_n, 0, \overline{\beta}_n)$ is the density of the first-passage time through $\overline{\beta}_n$ at time s starting from state x_n at time 0. This implies that

$$\frac{\partial V_n(x_n,t)}{\partial x_n} = \left(\overline{\beta}_n - K_n\right)\int_0^\infty e^{-\rho_n(t+s,s)}\frac{\partial \gamma_n\left(s|x_n, 0, \overline{\beta}_n\right)}{\partial x_n}ds.$$

Furthermore, the exercise threshold needs to be optimal against uniform perturbations on $\overline{\beta}_n$. Using the translation invariance $\gamma_n\left(s|x_n, 0, \overline{\beta}_n\right) = \gamma_n\left(s|x_n - \overline{\beta}_n, 0, 0\right)$, we obtain

$$\frac{\partial V_n(x_n,t)}{\partial \overline{\beta}_n} = \int_0^\infty e^{-\rho_n(t+s,s)}\gamma_n\left(s|x_n, 0, \overline{\beta}_n\right)ds$$

$$-\left(\overline{\beta}_n - K_n\right)\int_0^\infty e^{-\rho_n(t+s,s)}\frac{\partial \gamma_n\left(s|x_n,0,\overline{\beta}_n\right)}{\partial x_n}ds = 0.$$

It follows that $V_n\left(x_n,t\right) = \left(\overline{\beta}_n - K_n\right)\partial V_n(x_n,t)/\partial x_n$ for all $x_n < \overline{\beta}_n$. Differentiating further and substituting, we obtain $V_n\left(x_n,t\right) = \left(\overline{\beta}_n - K_n\right)^2 \partial^2 V_n(x_n,t)/\partial x_n^2$. The HJB equation then yields

$$\begin{aligned}\frac{\partial V_n\left(x_n,t\right)}{\partial t} &= \left[r + \lambda_n\left(t\right) - \mu_n\frac{1}{\left(\overline{\beta}_n - K_n\right)} - \frac{1}{2}\sigma_n^2\frac{1}{\left(\overline{\beta}_n - K_n\right)^2}\right]V_n(x_n,t)\\ &= \left[\lambda_n\left(t\right) - \overline{\lambda}_n\right]V_n(x_n,t).\end{aligned}$$

Solving for $\lambda_n(t)$, we obtain $\lambda_n\left(t\right) = \overline{\lambda}_n + \frac{\partial V_n(x_n,t)}{\partial t}\frac{1}{V_n(x_n,t)}$, which is valid for all $x_n < \overline{\beta}_n$. Taking the limit as $x_n \uparrow \overline{\beta}_n$, we obtain

$$\begin{aligned}\lambda_n\left(t\right) &= \lim_{x_n\uparrow\overline{\beta}_n}\left[\overline{\lambda}_n + \frac{\partial V_n\left(x_n,t\right)}{\partial t}\frac{1}{V_n(x_n,t)}\right] = \overline{\lambda}_n + \frac{\partial V_n\left(\overline{\beta}_n,t\right)}{\partial t}\frac{1}{V_n(\overline{\beta}_n,t)}\\ &= \overline{\lambda}_n + \frac{0}{\beta_n - K_n} = \overline{\lambda}_n,\end{aligned}$$

as claimed. □

Proof of Proposition 4. We will first establish Properties ii to v, and then Property i. Let τ be a stationary equilibrium. Then, by definition, each τ_n is the first-passage time through some constant exercise threshold, $\overline{\beta}_n$. By Lemma 3, the defeat rate of Player n must be the constant $\overline{\lambda}_n \equiv \mu_n\left(\overline{\beta}_n - K_n\right)^{-1} + \frac{1}{2}\sigma_n^2\left(\overline{\beta}_n - K_n\right)^{-2} - r$. Recall that, in equilibrium, $\lambda_n(t) = \sum_{m\neq n}\eta_m(t)$ for $n \in \boldsymbol{N}$. Independently of whether the equilibrium is stationary or not, this system of linear equations can be explicitly inverted to yield $\eta_n(t) = (N-1)^{-1}\left[\sum_{m\neq n}\lambda_m(t) - \lambda_n(t)\right]$. It follows that the equilibrium exercise rates must also be constant: $\eta_n(t) = \overline{\eta}_n \equiv (N-1)^{-1}\left[\sum_{m\neq n}\overline{\lambda}_m - \overline{\lambda}_n\right]$. To establish $\overline{\eta}_n \in \left(0, \frac{1}{2}\frac{\mu_n^2}{\sigma_n^2}\right]$, note that Equation (22) in the proof of Lemma 1 can be formally obtained for any $\overline{\eta}_n \in \mathbb{R}$. Inverting the Laplace transform in this expression, we obtain

$$\nu\left(y\right) = 2\overline{\eta}_n e^{-\frac{\mu_n}{\sigma_n}y}\frac{\sinh\left(\frac{\sqrt{\mu_n^2 - 2\overline{\eta}_n\sigma_n^2}}{\sigma_n}y\right)}{\sigma_n\sqrt{\mu_n^2 - 2\overline{\eta}_n\sigma_n^2}},$$

where $y \in [0,+\infty)$. Note that $\overline{\eta}_n < 0$ is inconsistent with equilibrium, as there is no mass infusion in this model, only absorption. Also, if $\overline{\eta}_n = 0$, we have $g\left(x_n^0\right) = \sigma_n\nu\left(\sigma_n^{-1}(\overline{\beta}_n - x_n^0)\right) = 0$ for all $x_n^0 \in (-\infty, \overline{\beta}_n]$, which is not a proper probability density. Moreover, if $\overline{\eta}_n > \frac{1}{2}\frac{\mu_n^2}{\sigma_n^2}$, we can define $A(y) \equiv \frac{2\overline{\eta}_n e^{-\frac{\mu_n}{\sigma_n}y}}{\sigma_n\sqrt{2\overline{\eta}_n\sigma_n^2 - \mu_n^2}} \in [0,+\infty)$, $B \equiv \frac{\sqrt{2\overline{\eta}_n\sigma_n^2 - \mu_n^2}}{\sigma_n} \in [0,+\infty)$ and write the expression above as $\nu(y) = A(y)(1/i)\sinh(Byi) = A(y)\sin(By)$. It follows that $\nu(y)$ is negative whenever $\sin(By)$ is negative. As a result, $g\left(x_n^0\right) = \sigma_n\nu\left(\sigma_n^{-1}(\overline{\beta}_n - x_n^0)\right)$ is negative over a set of $x_n^0 \in (-\infty, \overline{\beta}_n]$ that has positive Lebesgue measure and, thus, cannot be a probability density. We conclude that, if $\overline{\eta}_n \le 0$ or $\overline{\eta}_n > \frac{1}{2}\frac{\mu_n^2}{\sigma_n^2}$, there exists no prior marginal distribution F_n^0 that induces $1-\Gamma_n(t) = e^{-\overline{\eta}_n t}$. Together with Lemma 1, this observation implies Property ii.

Given that defeat rates are constant, Equations (16), (17), and (18) must hold in any stationary equilibrium, so Properties iii and iv are necessarily satisfied. Property v is an immediate consequence of Lemma 1.

Finally, to show that Property i holds, suppose, seeking a contradiction, that there exists another stationary equilibrium $\tau' \neq \tau$. Clearly, there must be at least one player for whom the exercise threshold in equilibrium τ' must differ from the one in equilibrium τ, say $\beta_n' \neq \overline{\beta}_n$. Following the steps of the argument we used to prove Property ii, we can determine the equilibrium exercise rate $\eta_n' \in \left(0, \frac{1}{2}\frac{\mu_n^2}{\sigma_n^2}\right]$. In equilibrium, the prior marginal distribution for Player n should be consistent with inducing a constant first-passage rate η_n' through the threshold β_n'. Using Lemma 1, it is easy to see such prior marginal distribution should have support $(-\infty, \beta_n']$, while F_n^0 has support $(-\infty, \overline{\beta}_n] \neq (-\infty, \beta_n']$. We conclude that τ' cannot be a stationary equilibrium under F^0. □

The following definition and lemma will be used in the proof of Proposition 5. Let

$$\Upsilon_n(t,h) \equiv \log\left(\frac{1-\Gamma_n(t+h)}{1-\Gamma_n(t)}\right).$$

Lemma 4. *Assume that the prior is degenerate at some arbitrary x^0 and consider an equilibrium such that $\beta(0) > x^0$. Then, for every $n \in \boldsymbol{N}$ and $h \in \mathbb{R}_+$, we have*

$$\lim_{t\to+\infty}\left(\frac{\Upsilon_n(t,h)}{h}\right) = \eta_n^*.$$

Proof. According to Proposition 2, equilibrium exercise thresholds must satisfy $\underline{\beta}_n \le \beta_n \le \overline{\beta}_n$ with $\underline{\beta}_n < \overline{\beta}_n$ for every $n = 1, ..., N$. Let $\underline{\Gamma}_n$ and $\overline{\Gamma}_n$ be the absorption probabilities associated with constant exercise thresholds $\underline{\beta}_n$ and $\overline{\beta}_n$. Clearly, $\overline{\Gamma}_n(t) \le \Gamma_n(t) \le \underline{\Gamma}_n(t)$ for all $t \in \mathbb{R}_+$. We will start showing that there exists a constant $A \in [0, +\infty)$ such that, for all $h \in [0, +\infty)$, we have

$$\limsup_{t\to+\infty} \Upsilon_n(t,h) \le \eta_n^* h + A. \tag{23}$$

Clearly, $\overline{\Gamma}_n(t) < \underline{\Gamma}_n(t)$ for all $t > 0$. Hence, for every $t > 0$ and $h \in \mathbb{R}_+$, we have $\frac{1-\Gamma_n(t+h)}{1-\Gamma_n(t)} > \frac{1-\underline{\Gamma}_n(t+h)}{1-\overline{\Gamma}_n(t)}$. Thus, $\Upsilon_n(t,h) < -\ln\left(\frac{1-\Gamma_n^K(t+h)}{1-\Gamma_n^M(t)}\right)$. Using L'Hôpital's rule, we can explicitly compute:

$$\lim_{t\to+\infty}\left(\frac{1-\underline{\Gamma}_n(t+h)}{1-\overline{\Gamma}_n(t)}\right) = e^{-\frac{\mu_n}{\sigma_n^2}\left(\overline{\beta}_n - \underline{\beta}_n + \frac{1}{2}\mu_n h\right)}\left(\frac{\underline{\beta}_n - x_n^0}{\overline{\beta}_n - x_n^0}\right).$$

It follows that

$$\begin{aligned}\limsup_{t\to+\infty} \Upsilon_n(t,h) &\le \lim_{t\to+\infty}\left[-\ln\left(\frac{1-\underline{\Gamma}_n(t+h)}{1-\overline{\Gamma}_n(t)}\right)\right] \\ &= -\ln\left[\lim_{t\to+\infty}\left(\frac{1-\underline{\Gamma}_n(t+h)}{1-\overline{\Gamma}_n(t)}\right)\right] = \eta_n^* h + A\end{aligned}$$

where we define $A \equiv \frac{\mu_n}{\sigma_n^2}\left(\overline{\beta}_n - \underline{\beta}_n\right) + \ln\left(\frac{\overline{\beta}_n - x_n^0}{\underline{\beta}_n - x_n^0}\right)$. Running a symmetric argument, we can obtain a lower bound for the limit inferior: $\liminf_{t\to+\infty} \Upsilon_n(t,h) \geq \eta_n^* h - A$. Next, we will show that, for all $h \in \mathbb{R}_+$, we actually have

$$\lim_{t\to+\infty} \Upsilon_n(t,h) = \eta_n^* h.$$

Fix $h \in \mathbb{R}_+$ and an arbitrary increasing and unbounded sequence of times $\{t_j\}_{j\in\mathbb{N}}$. The claim will be proven if we can show $\lim_{j\to+\infty} \Upsilon_n(t_j,h) = \eta_n^* h$. Since $\{\Upsilon_n(t_j,h)\}_{j\in\mathbb{N}}$ is eventually confined to the compact interval $[0, \eta_n^* h + A + 1]$, there is no loss in assuming that the whole sequence lies in a compact interval. Moreover, it is well-known that a sequence in a compact space X converges to $x \in X$ if and only if every convergent subsequence converges to x. As a result, it suffices to show that $\lim_{j\to+\infty} \Upsilon_n(t_j,h) = \eta_n^* h$ whenever the limit exists. So, assuming that the limit $\lim_{j\to+\infty} \Upsilon_n(t_j,h)$ exists, for every $m \in \mathbb{N}$, we have

$$\begin{aligned}\Upsilon_n(t_j, mh) &= -\ln\left(\frac{1-\Gamma_n(t_j+mh)}{1-\Gamma_n(t_j)}\right) = -\ln\left[\prod_{l=1}^{m}\left(\frac{1-\Gamma_n(t_j+lh)}{1-\Gamma_n(t_j+(l-1)h)}\right)\right] \\ &= \sum_{l=1}^{m}\left[-\ln\left(\frac{1-\Gamma_n(t_j+lh)}{1-\Gamma_n(t_j+(l-1)h)}\right)\right] = \sum_{l=1}^{m}\Upsilon_n(t_j+lh, h).\end{aligned}$$

This formally implies that

$$\begin{aligned}\lim_{j\to+\infty} \Upsilon_n(t_j, mh) &= \lim_{j\to+\infty}\sum_{l=1}^{m}\Upsilon_n(t_j+lh,h) = \sum_{l=1}^{m}\lim_{j\to+\infty}\Upsilon_n(t_j+lh,h) \\ &= m\lim_{j\to+\infty}\Upsilon_n(t_j,h).\end{aligned}$$

Reversing the derivation proves that the limit in the left-hand-side must also exist. It follows that

$$\lim_{j\to+\infty}\Upsilon_n(t_j,mh) = \liminf_{j\to+\infty}\Upsilon_n(t_j,mh) = \limsup_{j\to+\infty}\Upsilon_n(t_j,mh).$$

Then, using the inequalities for the limit inferior and superior, we get

$$\eta_n^* mh - A \leq \lim_{j\to+\infty}\Upsilon_n(t_j,mh) \leq \eta_n^* mh + A.$$

Combined with the additivity obtained above, this implies that $\eta_n^* h - \frac{1}{m}A \leq \lim_{j\to+\infty}\Upsilon_n(t_j,h) \leq \eta_n^* h + \frac{1}{m}A$. Since this inequality holds for every $m \in \mathbb{N}$, we must have $\eta_n^* h \leq \lim_{j\to+\infty}\Upsilon_n(t_j,h) \leq \eta_n^* h$, establishing the desired result. □

Now we can proceed to prove Proposition 5.

Proof of Proposition 5. Since the proof is relatively long, we only sketch the key steps here. A complete proof is available in Part S2 of the supplementary material.

Fix an equilibrium satisfying Assumptions 1 and 2. By Assumption 1, there is positive probability of the game continuing after $t = 0$ (and, in fact, after any $t \geq 0$). As a result, we can safely ignore those paths of play along which the game is stopped at $t = 0$, as they are irrelevant for future equilibrium behavior (and, thus, for asymptotics).

The equilibrium exercise threshold of Player n is constrained between $\underline{\beta}_n$ and $\overline{\beta}_n$. Moreover, Lemma 4 implies that, in the case of a degenerate prior, we have $\lim_{t\to+\infty}\Upsilon_n(t,h) = \eta_n^* h$ for every $h \in \mathbb{R}_+$. A technical argument (see S2.1) shows that this limit also holds when the prior satisfies Assumption 1. This result is important because it pins down the asymptotic behavior of the effective discount factors players use to compute their optimal strategies. More specifically, if we define

$$\Lambda_n(t,h) \equiv \log\left(\frac{1 - G_{[-n]}(t+h)}{1 - G_{[-n]}(t)}\right)$$

the effective discount factor of Player n is $e^{-rh-\Lambda_n(t,h)}$. It is easy to check that $\Lambda_n(t,h) = \sum_{m \neq n} \Upsilon_m(t,h)$, so the limit $\lim_{t\to+\infty} \Upsilon_n(t,h) = \eta_n^* h$ in fact implies that

$$\lim_{t\to+\infty} \Lambda_n(t,h) = \lim_{t\to+\infty} \sum_{m\neq n} \Upsilon_m(t,h) = \sum_{m\neq n} \lim_{t\to+\infty} \Upsilon_m(t,h) = \sum_{m\neq n} \eta_n^* h = \lambda_n^* h,$$

which obviously leads to the convergence of the effective discount factor.

The convergence of effective discount factors implies the uniform convergence of values (Property i). To see this, let $U_n(x_n, t, \beta_n)$ be the payoff that Player n obtains by playing an arbitrary continuation boundary $\beta_n : [0, +\infty) \to \mathbb{R}$ when he or she is at state x_n at time t and has a discount factor $e^{-rh-\Lambda_n(t,h)} \geq 0$ from time t to time $t+h$. Let $U_n^*(x_n, \beta_n)$ be the payoff that a monopolist with discount rate $r + \lambda_n^*$ would obtain at state x_n by playing the same continuation boundary β_n. Define $V_n(x_n, t) \equiv \sup_{\beta_n} U_n(x_n, t, \beta_n)$ and $V_n^*(x_n) \equiv \sup_{\beta_n} U_n^*(x_n, \beta_n)$. Both suprema are attained by thresholds taking values in $[\underline{\beta}_n, \overline{\beta}_n]$. Using $\lim_{t\to+\infty} \Lambda_n(t,h) = \lambda_n^* h$, we prove that, for every $x \in \mathbb{R}$ and $\beta_n : [0, +\infty) \to [\underline{\beta}_n, \overline{\beta}_n]$, we have $\lim_{t\to+\infty} U_n(x_n, t, \beta_n) = U_n^*(x_n, \beta_n)$ (see S2.2). Let $\hat{\beta}_n$ be a threshold that attains $V_n(x_n, t)$ and let β_n^* be the constant threshold that attains $V_n^*(x_n)$. On the one hand, $V_n(x_n, t) = U(x_n, t, \hat{\beta}_n(t)) \geq U(x_n, t, \beta_n^*)$ for all $t \geq 0$. Since $\lim_{t\to+\infty} U(x_n, t, \beta_n^*) = U^*(x_n, \beta_n^*)$ by the argument above, we have

$$\liminf_{t\to+\infty} V_n(x_n, t) \geq \lim_{t\to+\infty} U(x_n, t, \beta_n^*) = U^*(x_n, \beta_n^*) = V_n^*(x_n).$$

On the other hand, dominated convergence can be used to show $\limsup_{t\to+\infty} V_n(x_n, t) \leq V_n^*(x_n)$. This gives pointwise convergence of the value functions. Uniform convergence follows from combining pointwise convergence with the following properties of the value functions: they are non-negative, increasing, continuous, agree on $[\overline{\beta}_n, +\infty)$ and vanish when $x \to -\infty$.

To establish Property iii, note that, under Assumption 2, we have

$$\lim_{t\to+\infty} \lambda_n(t) = \lim_{t\to+\infty} \lim_{h\downarrow 0} \left(\frac{\Lambda_n(t,h)}{h}\right) = \lim_{h\downarrow 0} \lim_{t\to+\infty} \left(\frac{\Lambda_n(t,h)}{h}\right) = \lambda_n^*,$$

where the possibility of exchanging limits can be deduced from the assumption that the derivative $d\lambda(t)/dt$ is uniformly bounded and the Moore-Osgood theorem.

To obtain Property ii, we define $\lambda_n^L(t) \equiv \inf_{h\geq 0} \lambda_n(t+h)$ and $\lambda_n^H(t) \equiv \sup_{h\geq 0} \lambda_n(t+h)$. By construction, $\lambda_n^L(t) \leq \lambda(t+h) \leq \lambda_n^H(t)$ for all $t, h \geq 0$. Let $\beta_n^L(t)$ and $\beta_n^H(t)$ be the optimal exercise threshold of a monopolist with constant discount rates $r + \lambda_n^L(t)$ and $r + \lambda_n^H(t)$, respectively. A simple argument shows that $\beta_n^L(t) \geq \beta_n(t) \geq \beta_n^H(t)$. Property iii implies that $\lim_{t\to\infty} \lambda_n^L(t) = \liminf_{t\to\infty} \lambda_n(t) = \lambda_n^*$ and $\lim_{t\to\infty} \lambda_n^H(t) = \limsup_{t\to\infty} \lambda_n(t) = \lambda_n^*$. Thus, by definition, $\lim_{t\to+\infty} \beta_n^L(t) = \lim_{t\to+\infty} \beta_n^H(t) = \beta_n^*$, as desired.

Finally, it remains to establish convergence of beliefs. The argument proceeds as follows. The characteristic function of the conditional belief $\hat{F}_n(\cdot, t)$ has the following integral representation:

$$\psi_n(\omega,t) = \frac{\psi_n(\omega,0) - \int_0^t e^{M_n(\omega)s + i\omega\beta_n(s)}\Gamma_n(ds)}{e^{M_n(\omega)t}\left[1-\Gamma_n(t)\right]},$$

where $M_n(\omega) \equiv (1/2)\sigma_n^2\omega^2 - \mu_n\omega i$, while the characteristic function of $\hat{F}_n^*$ satisfies

$$\zeta_n(\omega) = \frac{e^{\omega\beta_n^* i}\lambda_n^*}{\lambda_n^* - M_n(\omega)}.$$

The application of an extension of L'Hôpital's rule to the complex function ψ_n proves that there exists $\omega_0 > 0$ such that $\psi_n(\omega,t)$ converges to $\zeta_n(\omega)$ for all $\omega \in (-\omega_0, \omega_0)$. Convergence of characteristic functions in a fixed neighborhood of 0 is enough to guarantee convergence in distribution of the state conditional on the absence of exercise. More precisely, $\lim_{t\to+\infty} \hat{F}_n(x_n,t) = \hat{F}_n^*(x_n)$ for all x_n at which $\hat{F}_n^*(\cdot)$ is continuous (that is, everywhere). The proof of this last claim combines the fact that $\hat{F}_n(\overline{\beta}_n,t) = 1$ for all $t \geq 0$ with a modification of Lévy's continuity theorem for sequences of random variables uniformly bounded above (or below) due to Zygmund (1951). □

Appendix B. Integral representations of beliefs, absorption rates, and value functions

B.1. Integral representation of the distribution over payoff states and the absorption density

In this section, we offer an integral representation of the backward-looking system in Equations (9)-(12). To simplify the exposition, we focus on the case in which the prior marginal distribution for Player $n \in \boldsymbol{N}$ is a point mass at x_n^0, so Equation (10) specializes to $f_n(x_n,0) = \delta\left(x_n - x_n^0\right)$, where δ is the Dirac delta function.

Proposition 6. *Whenever the absorption boundary β_n is continuously differentiable on $(0,+\infty)$, the survival density $f_n(x_n,t|x_n^0)$ admits the following integral representation:*

$$f_n(x_n,t|x_n^0) = \frac{\phi\left(\frac{x_n - x_n^0 - \mu_n t}{\sigma_n\sqrt{t}}\right)}{\sigma_n\sqrt{t}} - \int_0^t \frac{\phi\left(\frac{x_n - \beta_n(h) - \mu_n(t-h)}{\sigma_n\sqrt{t-h}}\right)}{\sigma_n\sqrt{t-h}}\gamma_n(h|x_n^0)dh. \tag{24}$$

In turn, the exercise density γ_n is the unique bounded solution to

$$\gamma_n(t|x_n^0) = \frac{\phi\left(A_n(t|x_n^0)\right)A_n(t|x_n^0)}{t} - \int_0^t \frac{\phi\left(B_n(t,h)\right)B_n(t,h)}{t-h}\gamma_n(h|x_n^0)dh, \tag{25}$$

where

$$A_n(t|x_n^0) \equiv \frac{\beta_n(t) - x_n^0 - \mu_n t}{\sigma_n\sqrt{t}} \quad \textit{and} \quad B_n(t,h) \equiv \frac{\beta_n(t) - \beta_n(h) - \mu_n(t-h)}{\sigma_n\sqrt{t-h}}.$$

Proof. In Part S3 of the supplementary material. □

The interpretation of Equation (24) is as follows. The first term on the right-hand side is always positive and describes the density of a Brownian motion without taking absorption into account. However, some paths that would have reached $X_n(t) = x_n$ have crossed the boundary previously at some time $h < t$ and need to be subtracted. At instant $h < t$, a density $\gamma_n(h|x_n^0)$ of paths is absorbed at state $X_n(h) = \beta_n(h)$. Conditional on being at that state at time h, they would have reached x_n at time t with a probability density given by

$$\frac{\phi\left(\frac{x_n-\beta_n(h)-\mu_n(t-h)}{\sigma_n\sqrt{t-h}}\right)}{\sigma_n\sqrt{t-h}}.$$

Therefore, the last term in Equation (24) integrates over $0 \leq h < t$, thereby effectively subtracting all previously absorbed paths.

Notice, however, that the characterization of the density f_n is incomplete without a description of the absorption density $\gamma_n(t|x_n^0)$. That absorption rate can be obtained as a function of the mass that is near the boundary, β_n, at time t, as indicated by Equation (12). It is also worth noting that Equation (25) is quite convenient for computational purposes,[23] because it has a recursive backward-looking structure and can be easily approximated by a finite sum. We also define the distribution associated with density $\gamma_n(t)$, which is particularly important for describing the arrival rate of the end of the game.

Together, Equations (24) and (25) fully characterize the dynamics of the individual state conditional on any arbitrary boundary. Whenever we restrict attention to the equilibrium threshold, β_n, these equations describe the equilibrium beliefs of the opponents of Player n. As previously discussed, that includes more information than strictly necessary to compute the optimal policies of those players. For that, it is sufficient to describe the defeat rate as perceived by them, which is a sufficient statistic for the individual problem.

So far, Equations (24) and (25) compute the survival and absorption densities when the initial position x_n^0 is commonly known. To generalize them toward any prior marginal distribution F_n^0, one simply needs to integrate these two functions against that distribution.

B.2. Optimal policy

In this section, we provide analytic expressions for optimal exercise thresholds and value functions in smooth equilibria. First, we define Player n's effective discount factor between dates t and $h > t$, $e^{-\rho_n(h,t)}$, by setting

$$\rho_n(h,t) \equiv \int_t^h [r + \lambda_n(s)]\, ds. \tag{26}$$

This effective discount factor summarizes all the strategic information about Player n's competitors and allows us to state the following result.

Proposition 7. *Suppose that, for each $n \in \mathbf{N}$, (V_n, β_n) is an equilibrium smooth value-threshold pair and $\lim_{t\to\infty} V_n(x_n, t)$ exists for every $x_n \in \mathbb{R}$. Then, β_n satisfies the following integro-differential equation:*

$$\beta_n(t) - K_n = \int_t^\infty e^{-\rho_n(h,t)} \frac{\phi\left(\frac{\beta_n(h)-\beta_n(t)-\mu_n(h-t)}{\sigma_n\sqrt{h-t}}\right)}{\sigma_n\sqrt{h-t}} \times \left[\sigma_n^2 + \left(\frac{\beta_n(h)-\beta_n(t)}{h-t} - 2\frac{d\beta_n(h)}{dh} + \mu_n\right)(\beta_n(h) - K_n)\right] dh, \tag{27}$$

[23] Equation (25) belongs to the class of Volterra integral equations of the second kind.

while the corresponding value function V_n is described, in the continuation region, by

$$V_n(x_n,t) = \frac{1}{2}\int_t^{\infty} e^{-\rho_n(h,t)} \frac{\phi\left(\frac{\beta_n(h)-x_n-\mu_n(h-t)}{\sigma_n\sqrt{h-t}}\right)}{\sigma_n\sqrt{h-t}} \times \left[\sigma_n^2 + \left(\frac{\beta_n(h)-x_n}{h-t} - 2\frac{d\beta_n(h)}{dh} + \mu_n\right)(\beta_n(h)-K_n)\right] dh. \tag{28}$$

Proof. In the supplementary material. □

Proposition 7 shows that the equilibrium exercise threshold is a fixed point of the operator on the right-hand side of Equation (27). The existence of the limit for the value function is guaranteed under Assumption 1 by Lemma 10.

Notice that Equation (27) does not require the separate computation of the evolution of the exercise density over future exercise times, which is embedded in the operator. This feature is common to some analytic representations of the value of American call-options, as derived by McKean (1965), Kim (1990), and Jamshidian (1992).[24] Moreover, the value function is fully determined by the behavior of the exercise threshold.

Appendix C. Extensions

In this section, we briefly discuss possible extensions of the model.

C.1. Geometric Brownian motion and alternative stochastic processes for payoffs

The model we have studied assumes that payoff innovations are additive, identically distributed, and independent. In the investment under uncertainty literature, another process is frequently used, the geometric Brownian motion, which features multiplicative innovations. It can be represented by

$$\frac{d\hat{X}_n(t)}{\hat{X}_n(t)} = \hat{\mu}_n dt + \hat{\sigma}_n dZ_n(t),$$

where $\hat{\mu}_n$ represents a geometric drift term and $\hat{\sigma}_n$ a exposure of the growth rate to the innovation in the standard Brownian $Z_n(t)$.

We can do the change of variables $X_n(t) \equiv \log \hat{X}_n(t)$ and obtain

$$X_n(t) = \mu_n dt + \sigma_n dZ_n(t),$$

where $\mu_n = \hat{\mu}_n - \frac{\hat{\sigma}_n^2}{2}$ and $\sigma_n = \hat{\sigma}_n$. In terms of these new variables, we write

$$V_n(x_n,t) = \sup_{\tau_n \geq t} \mathbb{E}\left\{e^{-r(\tau_n - t)} 1_{\tau_n < \hat{\tau}_{[-n]}} \left(e^{X_n(\tau_n)} - K_n\right) \middle| X_n(t) = x_n, \hat{\tau}_{[-n]} \geq t\right\}.$$

The HJB equation in the continuation region is still given by Equation (6). The only relevant changes are in the value-matching and smooth-pasting conditions, which become, respectively,

[24] The integral equation approach to free-boundary problems was pioneered by Kolodner (1956). Peskir and Shiryaev (2006) provide a detailed treatment of the free-boundary approach to optimal stopping. See Chiarella et al. (2004) for a survey of the integral representations for American financial options.

$$V_n(\beta_n(t), t) = e^{\beta_n(t)} - K_n \text{ and } \frac{\partial V_n(\beta_n(t), t)}{\partial x_n} = e^{\beta_n(t)}.$$

In this case, the monopolist problem has a solution as long as $\hat{\mu}_n < r$. Under this assumption and the change in boundary conditions for the value function, the characterization we have in the previous sections applies. In particular, the limit results are valid for the implied arithmetic Brownian motion. Interestingly, the threat of entry by Player n perceived by his or her opponents vanishes in the limit for some cases in which $\hat{\mu}_n > 0$, as it becomes possible that $\mu_n = \hat{\mu}_n - \frac{1}{2}\sigma_n^2 \leq 0$.[25]

The same reasoning, following a change of variables, allows generalizations of all results for processes and terminal payoffs that are increasing functions of an arithmetic Brownian motion. For more general Itô processes, generalizations of the results derived in Section 3.2 can be obtained. The key modification is that probability densities specific to those processes, as opposed to the normal distribution, emerge in the specific version of Proposition 6. Stationary equilibria can be constructed for more general cases following the insights from the literature on Brownian mortality models. However, the corresponding convergence results remain a topic for future research.

C.2. Beyond the winner-take-all case

For simplicity, we have assumed that all players that fail to be the first to exercise obtain a payoff of zero. More generally, we could have assumed that, in the event of defeat, Player n obtains a payoff of $0 \leq L_n(x_n, t) \leq V_n^M(x_n)$, which is convex, smooth, and nondecreasing in x_n, and bounded by the monopolist value function V_n^M. Additionally, let it have a well-defined limit, $\lim_{t\to\infty} L_n(x_n, t) = L_n^*(x_n)$, which also satisfies these assumptions. In this more general case, $L_n(x_n, t)$ could be motivated by another stage of a game, in which late entrants still have actions available.

The HJB would then be given by

$$rV_n = \max\left\{\mu_n \frac{\partial V_n}{\partial x_n} + \frac{1}{2}\sigma_n^2 \frac{\partial^2 V_n}{\partial x_n^2} + \lambda_n(t)\left[L_n(x_n, t) - V_n\right] + \frac{\partial V_n}{\partial t}, r(x_n - K_n)\right\}.$$

In the continuation region, we can rewrite it as

$$[r + \lambda_n(t)]\, V_n = \lambda_n(t) L_n(x_n, t) + \mu_n \frac{\partial V_n}{\partial x_n} + \frac{1}{2}\sigma_n^2 \frac{\partial^2 V_n}{\partial x_n^2} + \frac{\partial V_n}{\partial t}.$$

Notice that, beyond generating a modified discount rate of $r + \lambda_n(t)$, the threat of an opponent's entry generates a flow payoff externality of $\lambda_n(t) L_n(x_n, t)$ on the value of Player n. This flow is now positive, but it was previously normalized to zero. As a consequence, the value function would always be larger than in the case of $L_n(x_n, t) = 0$.

After accounting for this change in the HJB equation, there are no major departures in the characterization. The exercise thresholds are still bounded between a monopolist and perfect competition, and limit behavior is analogous to what has been derived.

[25] In this case, we can characterize a degenerate limit, in which generalized beliefs assign mass points at minus infinity for the position of every opponent.

C.3. Running costs, abandonment options

Again, for simplicity, we have assumed that firms face negligible running costs and a single decision, involving the time of entry. In some applications, researchers can be interested in the case in which running costs are significant and endogenous abandonment occurs.

These setups allow a few variations. Suppose first that exit cannot occur, but a running cost of $c_n > 0$ is present. Then, the HJB equation satisfies

$$r V_n = \max\left\{-c_n + \mu_n \frac{\partial V_n}{\partial x_n} + \frac{1}{2}\sigma_n^2 \frac{\partial^2 V_n}{\partial x_n^2} - \lambda_n(t) V_n + \frac{\partial V_n}{\partial t}, r(x_n - K_n)\right\}$$
$$= \max\left\{\mu_n \frac{\partial V_n}{\partial x_n} + \frac{1}{2}\sigma_n^2 \frac{\partial^2 V_n}{\partial x_n^2} - \lambda_n(t) V_n + \frac{\partial V_n}{\partial t}, r\left(x_n - K_n + \frac{c_n}{r}\right)\right\} - c_n.$$

If we define an auxiliary function, $\overline{V}_n(x_n, t) \equiv V_n(x_n, t) + c_n/r$, the HJB in the continuation region can be written as

$$[r + \lambda_n(t)]\,\overline{V}_n = \lambda_n(t)\frac{c_n}{r} + \mu_n \frac{\partial \overline{V}_n}{\partial x_n} + \frac{1}{2}\sigma_n^2 \frac{\partial^2 \overline{V}_n}{\partial x_n^2} + \frac{\partial \overline{V}_n}{\partial t}.$$

Value matching and smooth pasting then require $\overline{V}_n(\beta_n(t), t) = \beta_n(t) - K_n + c_n/r$ and $\partial \overline{V}_n(\beta_n(t), t)/\partial x_n = 1$. Under this new formulation, the optimal stopping problem is analogous to the previous version, but has a flow payoff externality of $\lambda_n(t) c_n/r$, which has the interpretation of a possible saving of the net present value of all future running costs that occurs with time-varying intensity $\lambda_n(t)$. One can then show that $\beta_n(t) \in \left[K_n - c_n/r, \beta_n^M(t)\right]$, where $\beta_n^M(t)$ is the optimal threshold for Player n in the absence of any competition. The asymptotic results would follow, again, after accounting for the change in the HJB and boundary conditions.

Once an abandonment option is introduced, another endogenous threshold needs to be derived. For sufficiently low states, a player finds it optimal to drop out. Because of the non-stationarity in the intensity of competition, this additional threshold is time varying in general, in the same way as the optimal exercise threshold. Again, we can construct the stationary limit for beliefs, conditional on both no previous exercise and no abandonment by each active player.[26] The methods to study the transitions developed in Sections 3 and 4 can be extended as well. In particular, the equilibrium would again be characterized by a coupled system of differential equations. In this system, backward-looking conditional beliefs take into account the absence of either exercise or abandonment by each of the active players. At the same time, forward-looking value functions take into account the defeat and abandonment rates by each opponent. The key difference in this case is that the list of still-active opponents needs to be incorporated as an additional state variable.

C.4. Correlation and public states

Unlike the previous extensions, allowing for correlation in the evolution of the individual payoffs introduces major difficulties. In the original setting, the defeat rate is a simple function of

[26] Notice that we assume players would observe the abandonment by any opponent. In contrast, if abandonment was not observable and players solely conditioned in the absence of exercise, perceived competition would vanish in the long run. A non-degenerate limit distribution would be recovered if new opponents also entered the competition without being observed. This last feature is present in Bobtcheff and Mariotti (2012).

time. A player's own payoff position and its previous path are not informative about the intensity of opposition she will face in the future. In contrast, correlation creates a linkage between one's own payoff evolution and the expected future competition. In principle, the defeat rate at time t becomes a function of the whole past trajectory of $X(s)$, for $s \leq t$. Extending the current techniques to deal with this non-Markov structure is a challenge left for future work.

Appendix D. Supplementary material

Supplementary material related to this article can be found online at https://doi.org/10.1016/j.jet.2019.104945.

References

Achdou, Y., Buera, F.J., Lasry, J.-M., Lions, P.-L., Moll, B., 2014. Partial differential equation models in macroeconomics. Philos. Trans. R. Soc. A, Math. Phys. Eng. Sci. 372, 20130397.

Bayraktar, E., Cvitanic, J., Zhang, Y., 2018. Large Tournament Games. Working paper.

Bensoussan, A., Frehse, J., Yam, P., 2013. Mean Field Games and Mean Field Type Control Theory. Springer.

Bobtcheff, C., Bolte, J., Mariotti, T., 2016. Researcher's dilemma. Rev. Econ. Stud. 84, 969–1014.

Bobtcheff, C., Mariotti, T., 2012. Potential competition in preemption games. Games Econ. Behav. 75, 53–66.

Bonatti, A., Hörner, J., 2011. Collaborating. Am. Econ. Rev. 101, 632–663.

Brekke, K.A., Øksendall, B., 1991. The high contact principle as a sufficient condition for optimal stopping. In: Lund, D., Øksendall, B. (Eds.), Stochastic Models and Option Values. Applications to Resources, Environment and Investment Problems. North-Holland, pp. 187–208.

Chiarella, C., Ziogas, A., Kucera, A., 2004. A Survey of the Integral Representation of American Option Prices. University of Technology Sydney.

Dixit, A.K., Pindyck, R.S., 1994. Investment Under Uncertainty. Princeton UP, Princeton.

Dosis, A., Muthoo, A., 2019. Experimentation in Dynamic R&D Competition. Working paper.

Dutta, P.K., Rustichini, A., 1993. A theory of stopping time games with applications to product innovations and asset sales. Econ. Theory 3, 743–763.

Fudenberg, D., Tirole, J., 1985. Preemption and rent equalization in the adoption of new technology. Rev. Econ. Stud. 52, 383–401.

Grenadier, S.R., 1996. The strategic exercise of options: development cascades and overbuilding in real estate markets. J. Finance 51, 1653–1679.

Grenadier, S.R., 2000. Option exercise games: the intersection of real options and game theory. J. Appl. Corp. Finance 13, 99–107.

Grenadier, S.R., 2002. Option exercise games: an application to the equilibrium investment strategies of firms. Rev. Financ. Stud. 15, 691–721.

Guo, Y., Roesler, A.-K., 2018. Private Learning and Exit Decisions in Collaboration. Working paper.

Hopenhayn, H.A., Squintani, F., 2011. Preemption games with private information. Rev. Econ. Stud. 78, 667–692.

Hopenhayn, H.A., Squintani, F., 2015. Patent rights and innovation disclosure. Rev. Econ. Stud. 83, 199–230.

Jackson, K., Kreinin, A., Zhang, W., 2009. Randomization in the first hitting time problem. Stat. Probab. Lett. 79, 2422–2428.

Jamshidian, F., 1992. An analysis of American options. Rev. Futures Mark. 11, 72–80.

Kim, I.J., 1990. The analytic valuation of American options. Rev. Financ. Stud. 3, 547–572.

Kolodner, I.I., 1956. Free-boundary problem for the heat equation with applications to problems of change of phase. I. General method of solution. Commun. Pure Appl. Math. 9, 1–31.

Lambrecht, B., Perraudin, W., 2003. Real options and preemption under incomplete information. J. Econ. Dyn. Control 27, 619–643.

Laraki, R., Solan, E., Vieille, N., 2005. Continuous-time games of timing. J. Econ. Theory 120, 206–238.

Lasry, J.-M., Lions, P.-L., 2007. Mean field games. Jpn. J. Math. 2, 229–260.

Lehmann, A., 2002. Smoothness of first passage time distributions and a new integral equation for the first passage time density of continuous Markov processes. Adv. Appl. Probab. 34, 869–887.

Martinez, S., Picco, P., San Martin, J., 1998. Domain of attraction of quasi-stationary distributions for the Brownian motion with drift. Adv. Appl. Probab. 30, 385–408.

Martinez, S., San Martin, J., 1994. Quasi-stationary distributions for a Brownian motion with drift and associated limit laws. J. Appl. Probab. 31, 911–920.

Mcdonald, R., Siegel, D., 1986. The value of waiting to invest. Q. J. Econ. 101, 707–728.

McKean, H.P., 1965. Appendix: a free boundary problem for the heat equation arising from a problem in mathematical economics. Sloan Manag. Rev. 6, 32.

Peskir, G., Shiryaev, A., 2006. Optimal Stopping and Free-Boundary Problems. Springer.

Quah, J.K.-H., Strulovici, B., 2013. Discounting, values, and decisions. J. Polit. Econ. 121, 896–939.

Seel, C., Strack, P., 2013. Gambling in contests. J. Econ. Theory 148, 2033–2048.

Strulovici, B., Szydlowski, M., 2015. On the smoothness of value functions and the existence of optimal strategies in diffusion models. J. Econ. Theory.

Thijssen, J.J., 2010. Preemption in a real option game with a first mover advantage and player-specific uncertainty. J. Econ. Theory 145, 2448–2462.

Trigeorgis, L., 1995. Real Options in Capital Investment: Models, Strategies, and Applications. Greenwood Publishing Group.

Weeds, H., 2002. Strategic delay in a real options model of R&D competition. Rev. Econ. Stud. 69, 729–747.

Zygmund, A., 1951. A remark on characteristic functions. In: Proceedings of the Second Berkeley Symposium on Mathematical Statistics and Probability, pp. 369–372.