COMPETITION AND LEARNING IN REAL OPTIONS

LEANDRO GORNO AND FELIPE S. IACHAN

FGV\EPGE

Abstract. We study a product development race by introducing competition under private information in an otherwise standard real-options environment. Firms observe the stochastic-evolution of their own potential payoffs from entry into a market, while forming conjectures about the state of their opponents. They face a trade-off between immediate entry or delay, in which the costs of waiting are enhanced by the endogenous threat of preemption by a rival. We characterize Markov Perfect Equilibria through a coupled system of differential equations: forward-looking value functions determine optimal exercise rules, while backward-looking beliefs determine the probability of an opponent's exercise. The equilibrium displays an intensity of competition that first builds up, then subsides slowly towards a steady-state. We additionally compute the equilibrium and illustrate some of its qualitative properties, such as the long-lived transient dynamics and the consequences of heterogeneity between competitors.

Keywords: real options, uncertainty, investment, learning, competition.

JEL Codes: C73, D92, G31.

Introduction

Real options, such as the option of bringing a newly developed product into a market, lack the precise contractual terms of their financial counterparts. Most notably, they typically do not expire on a clear deadline but might lose a significant share of their value as soon as competitors enter the market. Knowledge about the conditions of these competitors and the likelihood of their entry is strategically important for firms conducting product development, but is also limited by industrial secrecy.

We study this situation through a continuous-time real option framework. The model features both rivalry between the options held by different players and incomplete information. Each player faces a choice between the current exercise of the option (entry) or its delay. Delay has both a relative benefit and a cost. The benefit originates from the potential evolution of a state which determines the payoff upon exercise, as it is common in the investment under uncertainty literature. For instance, as time passes, the product in question can be further developed to become more profitable, market conditions can improve, or entry costs can decrease. Each player is privately informed about her own state. In additional to discounting, the relative cost of a delayed exercise includes the possibility that an opponent might enter the market first, reducing the original player's payoff. Beliefs about the likelihood of an opponent's entry in the near future are therefore key for determining optimal exercise strategies.

We analyze Markov-Perfect Equilibria of this game. The relevant states for each player are the value in case of current exercise and calendar time. In a recursive formulation, three essential objects are jointly determined. First, a player's value function. This value function takes into account the player's belief about the likelihood of an opponent's entry as time passes. The hazard rate of that entry, our second object, is central for the recursive formulation of the value function. Third, the value function itself induces an optimal exercise boundary which is time dependent. In equilibrium, the likelihood of the state hitting the exercise boundary and the hazard rate that shapes the value function need to be mutually consistent.

With symmetric players, the equilibrium is described by a pair of integral equations. A forward-looking equation obtains for the optimal exercise boundary. A backward-looking equation tracks the hazard rate. Although this hazard rate is sufficient for all strategic interactions, it is also possible to fully characterize the conditional belief distribution over the other players' states.

The equilibrium displays a fundamental time dependence. The passage of time unravels two forces. First, it initially becomes more likely that one of the opponents might exercise her option soon. Both the presence of a positive drift and random increments to payoffs from exercise contribute towards this effect. Second, the fact that her opponents have not exercised their options provides a player with information about the distribution of their states. Intuitively, if the opponents have not yet exercised after a long time, it becomes less likely that they are in a strong enough position to do so in the near future.

In the limit as time goes to infinity, both forces balance out, inducing convergence to a steady-state. Each player behaves as if facing a constant threat of entry, which can be represented by a modification to their discount rate. We provide an explicit formula for that discount modification.

Both the optimal exercise boundary and the induced hazard rate for the first entry display meaningful non-monotonicities. Competition intensity[1] initially increases, reaching a peak, then slowly subsides towards a constant level. The mechanism is the following. If a player attributes a high likelihood to the game ending soon, she becomes more aggressive, exercising the option at lower values. Therefore, she is more likely to bring the game to an end herself. This, in turn, makes opponents more aggressive.

Whenever there is common knowledge that options are unlikely to be exercised early on, competition is initially very weak. In particular, if initial conditions are known, agents are willing to wait for payoff improvements. As time elapses, the effects of randomness in payoff improvements come into play and the information about opponents' conditions becomes less precise. Both a positive drift and volatility contribute for making the opponent more likely to face a high state, creating an increasing perception of the threat of exercise. Players become more aggressive, exercising options at lower values, and the game more likely to end. Over time, however, the absence of previous exercise begins to be interpreted as a strong signal that the opponents are in a weak competitive

[1]We can make the notion of competition intensity precise by defining two related measures with similar qualitative behavior: the arrival rate of a player's defeat and, alternatively, how aggressively players behave (in terms of how close they get to zero profits in their optimal exercise boundaries).

position. In such circumstances, the perceived chances that a rival will exercise in the near future recede towards a stable level. As a consequence, exercise thresholds also approach a constant.

The model we study is a natural extension of standard real option environments and can serve as a useful framework for applied work. We provide an exploratory numerical analysis in Section 4. We choose a calibration that would ensure that half the firms would cross a break-even threshold in their product development within two years and the remaining half would do it over the following four years. Those two moments plus the choice of a discount rate are sufficient for all the parameters of the model and could be matched to empirical moments.

The equilibrium computed exhibits some noticeable dynamic features. At the most aggressive moment, firms observe a competition-modified discount rate which is 70% larger than the baseline one. Competition, as measured by the instantaneous arrival rate of a possible entry by an opponent builds up steeply over the first five years and then starts receding. Although convergence to a steady state occurs, it is remarkably slow, with half-lives for these arrival rates which are multiple-decades long.

While this example suggests the importance of the magnitudes involved, these qualitative features are robust to other calibrations. Over time, there is a variation of approximately 20% in the economic surplus required for strategic exercise. As a practical matter, it follows that reduced-form approaches imposing a single modified discount rate to take into account competition can be deceiving and lead to significant value losses. From a research point of view, the lesson is that focusing on stationary situations might be a bad approximation to actual behavior, as transient dynamics are extremely long lived. The results suggest that, if one wants to adopt a reduced-form approach to model competitive-entry, it is advisable to allow for enough flexibility in the parametric specification of the effective discount rate to accommodate the dynamic behavior of relevant measures of competition.

The numerical framework can also be used for studying asymmetric competition and comparative dynamics across different industries, as characterized by the expected speed of build up in product developments and its volatility. We illustrate how each perturbation, such as endowing a firm with a initial advantage, has both mechanical effects (as it takes that firm closer to the exercise boundary) and strategic ones (as the opponent sees stronger competition initially and responds more aggressively). Strategic effects tend not to be uniform throughout time and usually change directions. The intuition is that if one's opponent becomes more aggressive in the initial years, one should respond more aggressively during that period; however, once those years pass without leading to any entry, one has a stronger signal that the opponent was never in a particularly strong position. As such, competition recedes.

Related literature - This paper is related to a growing literature on dynamic contests and competitive real options. In particular, the game we study belongs to the class of optimal-stopping games, as initially laid out by Dutta and Rustichini (1993), and is related to the class of preemption games, notably studied in Fudenberg and Tirole (1985).

Some work has applied optimal stopping games to incorporate strategic components into a real options framework. An early example is Grenadier (1996), which studies real estate market dynamics. In that set-up, all players share a common state describing market conditions, which is publicly observable. Similar environments in that aspect are present in Grenadier (2002) and Weeds (2002).[2]

Relative to that standard set-up, we take into account two novel features. First, each firm is subject to a particular state describing its payoffs if the option is exercised. This introduces a multidimensional aspect into the problem. Second, each firm is privately informed about the evolution of its this state, while other firms can only draw some noisy inference about that variable. To the best of our knowledge, the only paper to include both these features is Hopenhayn and Squintani (2011).[3]

A key distinction between that paper and ours lies in the stochastic process driving payoffs. Hopenhayn and Squintani (2011) study a non-decreasing process, so that exercise can only become more valuable as time passes. Our paper is a more direct extension of the traditional investment under uncertainty benchmark: payoffs follow a Brownian motion with drift, so a deterioration of current conditions that makes exercise less profitable is possible. First, in applications, a decrease in expected payoffs from exercise can originate from the worsening of market conditions. Second,

[2]A good review of prior work is available in Grenadier (2000).

[3]Multidimensionality without private information is present in Thijssen (2010). Lambrecht and Perraudin (2003) study an environment with a common randomly evolving payoff state and private information regarding a static, non-evolving, exercise cost. Quah and Strulovici (2013) study an individual optimal stopping problem in the presence of non-stationary discounting.

even if product designs can only improve, one might become aware of previous overestimates of their profitability, making decreases in expected profitability arguably as natural as increases for R&D applications.

Importantly, the choice of the stochastic process driving the exercise payoffs is critical for the results. Hopenhayn and Squintani (2011) obtain a degree of competition that increases towards perfect competition as time passes. In the setting we study, the threat of entry of competitor is perceived as time varying and non-monotonic. As we show on the next few sections, an opponent's exercise is initially perceived as unlikely, given that R&D progress requires some time to build up. Then, as time passes, the threat of entry becomes more intense. Eventually, however, a long time without an opponent's entry is understood as a sign of bad payoff conditions for that opponent and of unlikely entry in the near future. These last two effects balance out in the long-run, leading to a well-defined limit level of perceived competition. The nature of the stochastic process, which allows for bad news about profitability, is essential for this non-monotonicity.

Additionally, our approach relies on a coupled system of differential equations: a forward-looking value function (or equivalently an exercise boundary) 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 nascent mean-field games literature.[4] Some important Macroeconomic applications, which rely on General Equilibrium Theory instead of strategic behavior of large firms, are presented in Achdou et al. (2014).

1. Some Preliminary Intuition

In order to develop some preliminary intuition, let us start with the simplest possible continuous-time scenario. Two symmetric players compete in a winner-takes-it-all race to develop a product and be the first entrant into a market.

We look at the problem from the perspective of a given agent, Player 1, who does not observe the level of development of her opponent, Player 2. Player 1 privately observes the evolution of the expected profitability of her product $X(t)$ and discounts the future at a rate r. Suppose that the cost of product introduction into the market is fixed at $K > 0$, so that $X(t) - K$ is the net payoff from exercise at time t. We assume that the initial condition $x_0 = 0$ is known and that $X(t)$

[4] See for instance Lions and Lasry (2007) and Bensoussan et al. (2013).

follows a standard (with zero drift and unit variance) Brownian motion. Since we are considering a symmetric situation, the same holds for the unobserved state of the other player. If the opponent enters, the game ends and the Player 1 obtains a payoff of zero.

The player needs to form conjectures about the instantaneous arrival rate of a defeat or, equivalently in this case, the hazard rate of the single opponent's entry. Suppose that Player 1 conjectures this arrival rate to be constant over time and denote it by $\lambda > 0$. It is a trivial extension of well known results (Dixit and Pindyck, 1994; Mcdonald and Siegel, 1986)[5], that in this case the value function is stationary and satisfies the following Hamilton-Jacobi-Bellman (HJB) equation

$$rV\left(x\right) = \max\left\{-\lambda V\left(x\right) + \frac{1}{2}V''(x), r\left(x-K\right)\right\}.$$

The maximization above is between the continuation value and immediate exercise that leads to a payoff of $x - K$ when the state of Player 1's product is $X(t) = x$. The first term inside the brackets, that refers to the continuation value, is the combination of the possible occurrence of the defeat, which happens with intensity λ and leads to immediate jump of the player's value to zero (that is, a value loss of the difference$0 - V\left(x\right)$), and the effects of volatility in the immediate evolution of the state.

The solution is characterized by a constant boundary β, so that exercise is optimal if, and only if, $X(t) \geq \beta$.[6] Optimality of the boundary implies that

$$\beta = K + \frac{1}{\sqrt{2\left(r+\lambda\right)}}.$$

Static net-present value (NPV) maximization, which does not take into account the option value of an entry delay, leads to investment whenever $X(t) \geq K$. Therefore, the boundary β displays a positive wedge relative to this static investment criterion. That wedge is decreasing in both λ and r. It is also well known[7] that the wedge would increase in the variance of the state $X(t)$, which here we have kept normalized at the value of one. Notice additionally that the arrival rate and

[5]For a recent and formal treatment of one-dimensional stochastic control and stopping problems in Economics, see Strulovici and Szydlowski (2015).

[6]In the region below that boundary, *i.e.* when $X(t) < \beta$, the value function satisfies$V\left(x\right) = C_0 exp\left(\gamma \cdot x\right)$, for some positive constant C_0 and $\gamma = \sqrt{2\left(r+\lambda\right)}$, which is the single positive root associated with the differential equation $\frac{1}{2}V^{\prime\prime}(x) = \left(r+\lambda\right)V$. Both the constant C_0 and the optimal exercise boundary are pinned down by the smooth pasting conditions$V\left(\beta\right) = C_0 exp\left(\gamma \cdot \beta\right) = \beta - K$ and $V^{'}\left(\beta\right) = \gamma V\left(\beta\right) = \frac{\partial(x-K)}{\partial x}|_{x=\beta} = 1$.

[7]See Dixit and Pindyck (1994).

the discount rate play analogous roles, so an increase in the arrival rate of the option expiration induced by the opponent's exercise is equivalent to increase in time discount.

Additionally, by varying λ, one can span situations with different underlying degrees of competition. For instance, for $\lambda = 0$, we recover the optimal decision of a monopolist, who does not face the possibility of an entry by a rival. On the other hand, in the the limit in which λ tends to infinity, competition becomes so fierce that the boundary, β, approaches the exercise cost, K, from above and investment occurs for arbitrarily small NPV opportunities. This is a profit dissipation result that is reminiscent of Bertrand competition. As we show later, equilibrium exercise boundaries are bounded by those two extreme cases.

We can now see that the best-reply to the belief of a constant arrival rate is a stationary boundary. The key difficulty in this environment is that a constant boundary does not induce a constant arrival rate for the end of the game.

To illustrate that, first we use the fact that a standard Brownian Motion eventually hits any constant and finite boundary β with probability 1. The left-hand side panel in Figure 1 below plots the density over time for that absorption. The right-hand side panel plots the hazard rate, which is the relevant variable for the individual problem, as agents are interested in the likelihood of the game ending any instant given that it has not ended previously. This hazard rate already displays a few important features. It starts at zero and is continuous, showing that the game is very unlikely to end right after $t = 0$. Second, as volatility in $X(t)$ accumulates, the hazard rate increases. Last, it eventually asymptotes to a constant. As we show later, all of these features are shared by the equilibrium hazard rate.

Whenever the hazard rate changes over time, the value function of the player inherits that time dependence. Once this value function becomes non-stationary, the optimal threshold $\beta(t)$ also varies with time. For instance, all else held fixed, a belief associated with a larger $\lambda(t)$ is expected to induce a more aggressive behavior, as represented by lower exercise threshold $\beta(t)$. Equilibrium needs to be characterized by an optimal stopping rule $\beta(t)$ and a hazard rate $\lambda(t)$ that are mutually consistent. We extend this simple set-up and introduce the problem more formally in the next section.

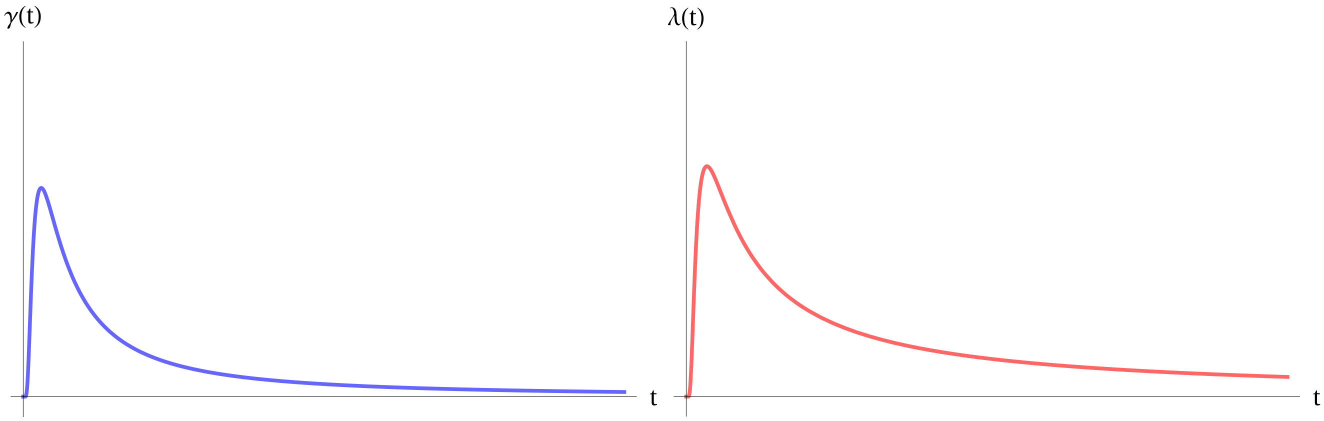

(A) Hitting density, $\gamma(t)$. (B) Hazard rate, $\lambda(t)$.

FIGURE 1. Hitting-time density and hazard rate induced by a constant boundary β.

2. Model

2.1. **Description of the game.** Time is continuous and the horizon is infinite. There is a finite set $\boldsymbol{N}$ of players, indexed by $n \in \{1, 2, ..., N\}$. The discount rate is $r > 0$ for every player. Each player $n \in \boldsymbol{N}$ has a position $X_n(t)$, where $X_n \equiv \{X_n(t)\}_{t \geq 0}$ is stochastic process starting at a condition x_n^0 which is common knowledge[8] and satisfying

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

where Z_n is a Wiener process and $\mu_n \geq 0$ and $\sigma_n > 0$ are constant player-specific drift and volatility. The processes $\{Z_n\}_{n \in \boldsymbol{N}}$ are independent.

The positions can be understood to represent the development state of projects, measured as a gross expected payoff from current exercise. Their evolution is private information, so each player knows her own progress but does not know the progress of her opponents. While we are restricting attention to increments in the states that are independent across agents, the drift term can incorporate common changes in the values of exercise, while innovations represent deviations from that deterministic path.

Each player decides at every instant whether to exercise the option or wait for more information. If player n exercises when $X_n(t) = x$, the game ends and she obtains a payoff of $x - K_n$, while her

[8]The set-up can be easily extended to take into account a non-degenerate distribution of initial conditions, incorporating incomplete information at $t = 0$.

opponents get 0. We assume that the exercise cost is positive, i.e. $K_n > 0$, and that there is no running cost for staying in the game, so that waiting at $t = 0$ is optimal. We call this game the *competitive option game*.

2.2. **Information, strategies, and payoffs.** Let $\mathcal{F}_n \equiv \{\mathcal{F}_n(t)\}_{t\geq 0}$ be the filtration generated by X_n, for $n \in \boldsymbol{N}$. A strategy profile is a n-tuple $\{\hat{\tau}_n\}_{n\in N}$, where $\hat{\tau}_n$ is a $\mathcal{F}_n-$stopping time for each $n \in \boldsymbol{N}$. We allow stopping times to be infinite when a player decides to never submit her project and get payoff of zero.

Let $\mathcal{F}$ be the filtration generated by $\{X_n\}_{n\in N}$. Notice that $\mathcal{F}$ contains more information than observed by each player individually. The game ends at the $\mathcal{F}-$stopping time

$$\hat{\tau} \equiv \min_{n\in \boldsymbol{N}} \hat{\tau}_n,$$

i.e., the game ends whenever the first player exercises her option. If any player exercises her option before player n, we say that player n was defeated.

The expected discounted payoff of player n at time $t \geq 0$ obtained by using strategy $\hat{\tau}_n$ when her opponents uses the strategy profile $\hat{\tau}_{-n}$ is given by

$$J_n(\hat{\tau}_n, \hat{\tau}_{-n}|t) \equiv \mathbb{E}\left\{e^{-r(\hat{\tau}_n - t)} 1_{\hat{\tau}_n = \hat{\tau}} \left(X^n_{\hat{\tau}_n} - K_n\right) \middle| \mathcal{F}_n(t)\right\}.$$

Agents can only observe the passage of time and the evolution of their own position $\{X_n(t)\}$. A strategy $\hat{\tau}_n$ for player n is a *Markov strategy* if the conditional stopping times

$$\hat{\tau}_n(t) \equiv \hat{\tau}_n | \mathcal{F}_n(t)$$

only depend on her current position $X_n(t)$ and the time elapsed since the start of the game.

Let $\mathcal{S}_n$ (resp. $\mathcal{M}_n$) be the set of strategies (resp. Markov strategies) for player n. As usual, denote by $\mathcal{S}_{-n}$ (resp. $\mathcal{M}_{-n}$) the sets of strategy (resp. Markov strategy) profiles for the opponents of player n. We restrict attention to the set of Markov strategy profiles.[9]

[9]Given that the components of (X^1, X^2) are independent strong Markov processes, it is natural to expect that Markov strategies should be enough for each player to best respond to her opponents, even if they play in a non-Markovian way. More formally, we conjecture that

$$\sup_{\hat{\tau}_n \in \mathcal{M}_n} J_n(\hat{\tau}_n, \hat{\tau}_{-n}|t) = \sup_{\hat{\tau}_n \in \mathcal{S}_n} J_n(\hat{\tau}_n, \hat{\tau}_{-n}|t)$$

for all $\hat{\tau}_{-n} \in \mathcal{S}_{-n}$. This would mean that we can restrict consideration to strategies in $\mathcal{M}_n$.

Thus, we define the expected discounted payoff of player n at time $t \geq 0$ and state x_n from following strategy $\hat{\tau}_n$ when her opponents use strategies $\hat{\tau}_{-n}$ by:

$$U_n(\hat{\tau}_n, \hat{\tau}_{-n}|x, t) \equiv \mathbb{E}\left\{e^{-r(\hat{\tau}_n - t)} 1_{\hat{\tau}_n = \hat{\tau}} \left(X_n(\hat{\tau}_n) - K_n\right) \middle| X_n(t) = x\right\}.$$

For any Markov strategy $\hat{\tau}_n \in \mathcal{M}_n$, we have

$$U_n(\hat{\tau}_n, \hat{\tau}_{-n}|X_n(t), t) = J_n(\hat{\tau}_n, \hat{\tau}_{-n}|t)$$

almost surely for all $t \geq 0$. That is, expected discounted payoffs under Markov strategies only depend on the state $(X_n(t), t)$.

2.3. **Equilibrium.** An *equilibrium* for the competitive option game is a strategy profile $\tau = (\tau_1, ..., \tau_N) \in \prod_{n=1}^N \mathcal{S}_n$ such that each τ_n maximizes the expected discounted payoffs of player n after any history holding strategies τ_{-n} fixed for all other players. Formally,

$$J_n(\tau_n, \tau_{-n}|t) \geq J_n(\hat{\tau}_n, \tau_{-n}|t)$$

almost surely for all $\hat{\tau}_n \in \mathcal{S}_n$, $t \geq 0$ and $n \in \boldsymbol{N}$.

A *Markov perfect equilibrium* (MPE) is an equilibrium in Markov strategies. Given the discussion in the previous section, a MPE can be characterized as a strategy profile $(\tau_1, ..., \tau_N) \in \prod_{n=1}^N \mathcal{M}_n$ such that

$$U_n(\tau_n, \tau_{-n}|x_n, t) \geq U_n(\hat{\tau}_n, \tau_{-n}|x_n, t)$$

for all $\tau_n \in \mathcal{M}_n$, $x_n \in \mathbb{R}$, $t \geq 0$ and $n \in \boldsymbol{N}$.

2.4. **A simpler recursive formulation.** Fix a MPE τ, which is not required to be symmetric even if all players share the same parameters. Let $V_n(x, t)$ be the equilibrium payoff of player n starting at time $t \geq 0$ from state x, i.e.,

$$V_n(x, t) \equiv \sup_{\hat{\tau}_n \in \mathcal{M}_n} U_n(\hat{\tau}_n, \tau_{-n}|x, t).$$

For notation simplicity, let us leave implicit the dependence on the state (x, t) and write V_n to represent the value above.

The associated Hamilton-Jacobi-Bellman (HJB) equation for this value is

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

where $\lambda_n(t)$ is the equilibrium arrival rate of the defeat of player n. In other words, $\lambda_n(t)$ is the equilibrium arrival rate of end of the game induced by the exercise from any of the opponents of player n, conditional on the game not having ended previously.

The first term inside the maximization is the value of continuation, while 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_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 time dependence originates from the arrival rate: as discussed in the introductory section, were the arrival rate time-invariant, the value function would be stationary.

The HJB equation is solved as a free-boundary problem of the partial differential equation (PDE)

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

with free-boundary conditions given by

$$V_n(\beta_n(t), t) = \beta_n(t) - K_n \tag{3}$$

and

$$\frac{\partial V_n(x,t)}{\partial x}|_{x=\beta(t)} = 1, \tag{4}$$

where $\beta_n\left(t\right)$ is a free-boundary, which might depend on t. Equation 3 represents the value-matching condition at the boundary, while equation 4 is the smooth-pasting condition.

2.5. **Equivalent MPE definition.** Before moving towards a characterization, we present an alternative equilibrium notion, which relies on a recursive description of mutually consistent continuation payoffs, optimal exercise boundaries, and perceived arrival rates of any opponents' exercise. This alternative notion is equivalent to MPE, as we show.

Definition 1. A competitive exercise equilibrium consists of a collection $\{V_n, \beta_n, \lambda_n\}_{n\in N}$, such that, for each $n \in \boldsymbol{N}$, $V_n : \mathbb{R} \times \mathbb{R}_+ \to \mathbb{R}$, $\beta_n \in \ell^\infty(\mathbb{R}_+)$, and $\lambda_n : \mathbb{R}_+ \to \mathbb{R}_+ \cup \{+\infty\}$ and

(1) V_n solves the HJB equation (1) taking λ_n as given,
(2) β_n is an optimal exercise boundary given V_n and λ_n,
(3) λ_n is the arrival rate of the end of the game induced by the exercise decision of any of the opponents. Formally, λ_n is the hazard rate associated with the cumulative distribution of the random variable $\tilde{\tau} \equiv \inf\{t \geq 0 | \exists w \in N \backslash \{n\},\ X_t^w \geq \beta^w(t)\}$.

The following analogy can be drawn. A competitive exercise equilibrium behaves in a similar way as a recursive competitive equilibrium from General Equilibrium Theory. Value functions and optimal policies solve each agent's dynamic problem. They do so by using at each instant in time the least information necessary about the rest of the economy: here, the arrival rate for the end of the game works as a price, summarizing all relevant information about the opponents' behavior.

To show the equivalence, fix a MPE τ. First, notice that

$$V_n(x_n, t) = U_n(\tau_n, \tau_{-n} | x_n, t).$$

Second, the induced stopping boundary for player n is

$$\beta_n(t) = \sup\{x_n \in \mathbb{R} | V_n(x_n, t) > x_n - K_n\}.$$

Finally, the equilibrium strategy can be recovered as

$$\tau_n = \inf\{t > 0 | X_t^n \geq \beta_n(t)\}.$$

In light of this equivalence, in what follows, we use any of the equivalent descriptions of equilibria interchangeably.

3. Results

3.1. **Bounds on exercise boundaries.** We start by bounding the optimal exercise threshold. As it is intuitive, the behavior of a competitive player lies 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.

Formally, our first result bounds the exercise times in any MPE by eliminating dominated strategies. In order to state the theorem, define individual specific constant boundaries $\underline{\beta}_n \equiv K_n$ 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 boundary and $\overline{\beta}_n$ the stationary boundary that prevails for the optimal exercise of a monopolist. The number ξ_n is the positive root of the characteristic polynomial associated with the ordinary differential equation $rV_n + \mu_n(\partial V_n/\partial x) - (1/2)\sigma_n^2(\partial^2 V_n/\partial x^2) = 0$, which describes the evolution of the value of continuation in the absence of the possibility of the end of the game induced by an opponent's exercise.

Using these constant boundaries, 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 boundary and the (least aggressive) monopolistic boundary. The next result shows that the ranking of the two fixed boundaries is translated to these stopping times and, more importantly, that these stopping times bound the MPE strategies in our environment.

Proposition 1. *Let* $(\tau_1, ..., \tau_N)$ *be a MPE with exercise boundaries* $(\beta_1, ..., \beta_N)$*. Then,* $\underline{\tau}_n \leq \tau_n \leq \overline{\tau}_n$ *and* $\underline{\beta}_n \leq \beta_n \leq \overline{\beta}_n$ *for every player* $n \in \boldsymbol{N}$*.*

Proof. In the appendix. □

Proposition 1 is important for constraining possible equilibrium exercise boundaries and stopping times. It 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 arrival of defeat. Before proceeding towards the characterization of this limit, we study how conditional probability distributions and beliefs evolve in this setting. These are important objects for the individual behavior, as they describe the intensity of competition a player expects to face over time.

3.2. **Equilibrium exercise densities and belief evolution.** To characterize Markov Perfect Equilibria, we first resort to an intermediate result that describes the evolution of a Brownian motion density when subject to a given stopping boundary β_n. This result is directly related to the distribution of an agent's stopping time and will also be important for characterizing equilibrium beliefs about conditions of opponents and the likelihood of their exercise. To simplify its characterization, we make the following uniform differentiability assumption:

Assumption 1. *For each player $n \in \boldsymbol{N}$, the equilibrium boundary β_n is continuously differentiable on $(0, +\infty)$ with uniformly bounded derivative.*

We denote the density of the current state for paths that have not previously hit the stopping 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 Kolmogorov forward equation

$$\frac{\partial f_n}{\partial t} = -\mu_n \frac{\partial f_n}{\partial x} + \frac{1}{2}\sigma_n^2 \frac{\partial^2 f_n}{\partial x^2}, \ \ for\ x < \beta(t). \tag{5}$$

The interpretation of this equation is the following. 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 $f_x > 0$ ($f_x < 0$), a given state x loses (gains) density in proportion to the drift μ_n. The second term originates from the volatility in the 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. Since we assume that the initial condition is deterministic, it is described by a Dirac function[10], i.e. a degenerate unit mass concentration at a point, as in

$$f_n(x_n, 0) = \delta\left(x_n - x_n^0\right). \tag{6}$$

[10]Formally, a generalized function (distribution) on the real number line that is zero everywhere except at zero, with an integral of one over the entire real line.

Additionally, given that β_n works as an absorbing boundary, the density vanishes at that boundary, implying the following boundary condition for the PDE in 5

$$f_n\left(\beta_n\left(t\right),t\right)=0. \tag{7}$$

We can use equations (5)-(7) to characterize the probability distribution of the state x_t^n, conditional on it not having been previously absorbed, as well as 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), which we denote byγ_n. This is the content of the next result.

Proposition 2. *Under Assumption 1, the density f_n admits the following integral representation*

$$f_n(x_n,t)=\frac{\phi\left(\frac{x_n-x_n^0-\mu_n t}{\sigma_n\sqrt{t}}\right)}{\sigma_n\sqrt{t}}-\int_0^t\frac{\phi\left(\frac{x_n-\beta_n(h)-\mu_n(t-h)}{\sigma_n\sqrt{t-h}}\right)}{\sigma_n\sqrt{t-h}}\gamma_n(h)dh. \tag{8}$$

In turn, the exercise density γ_n is characterized by

$$\gamma_n(t)=\frac{\phi\left(A_n(t)\right)A_n(t)}{t}-\int_0^t\frac{\phi\left(B_n(t,h)\right)B_n(t,h)}{t-h}\gamma_n(h)dh, \tag{9}$$

where

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

Proof. In the appendix. □

The interpretation of equation 8 is the following. The first term on the right-hand side is always positive and describes the density of a Brownian motion without taking into account absorption. However, some paths that would have reached $X_n(t,\omega)=x_n$ have crossed the boundary previously at some time $h<t$ and need to be subtracted. At instant $h<t$, a density $\gamma_n(h)$ of paths is absorbed at state $X_n(h)=\beta_n\left(h\right)$. 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 8 integrates over all $0\leq h<t$, effectively subtracting all paths that have been previously absorbed.

Notice, however, that the characterization of the density f_n would be incomplete without a description of the absorption density $\gamma_n(t)$. That absorption rate can be obtained as a function of the mass that is near the boundary β at time t. As this mass is proportional to the slope of the density in the limit as x_n approaches the boundary, the rate $\gamma_n(t)$ can be obtained from the density formula 8. It is also worth noting that equation 9 is quite convenient for computational purposes[11], because it has a recursive backward-looking structure and can easily be 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.

Definition 2. For each player $n \in \boldsymbol{N}$, the unconditional distribution of her exercise time is defined as $\Gamma_n(t) \equiv \int_0^t \gamma_n(h)\, dh$.

Together equations 8 and 9 fully characterize the dynamics of the individual state conditional on any arbitrary boundary. Whenever we restrict attention to the equilibrium boundary β_n, they can be used to describe the equilibrium beliefs of the opponents of player n. As previously discussed, that includes more information than what is strictly necessary in order to compute the optimal policies of those players. For that, it is sufficient to describe the distribution of the game's end time as perceived by them, which is asufficient statistic for the individual problem.

3.3. **Equilibrium conditional exercise rates.** We now consider the decision problem of player n and obtain the conditional arrival rate of the end of the game as perceived by her, $\lambda_n(t)$. In order to do that, note that the game is over the first time a player exercises an option. This means that we need to find the distribution of the random variable given by the first stopping of an opponent of player n, that is,

$$\tau_{[-n]} \equiv \min_{m \neq n} \tau_m.$$

That random variable is characterized by the cumulative distribution function

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

[11]Equation 9 belongs to the class of Volterra integral equations of the 2nd kind.

with associated density function given by $g_{[-n]}(t) \equiv \frac{dG_{[-n]}(t)}{dt}$. The conditional arrival rate of an end of the game induced by the opponents of player n, which is essential for the description of her HJB equation is then simply

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

Proposition 3. *The arrival rate of a defeat for player n is given by the sum of the hazard rates associated to the unconditional distributions of the exercise times of her opponents, i.e.,*

$$\lambda_n(t) = \sum_{m \neq n} \left(\frac{\gamma_m(t)}{1 - \Gamma_m(t)} \right). \tag{10}$$

Proof. Notice that

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

Notice that, as a consequence of the independence between payoff increments across players, the arrival rate of the defeat of any player is simply given by the sum over all of her opponents' boundary-crossing rates. Loosely, keeping strategies fixed, if one doubles the number of players, the arrival rates of defeats 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 result for the hazard rate on the total number of players is still true in the limit. □

3.4. **Optimal Policy.** We now turn our attention back to the value function and the optimal policy. We first describe the payoff from an arbitrary policy, conditional on the equilibrium arrival rate of a defeat.

Definition 3. Take any bounded boundary $\hat{\beta}$ for player n and a current state (x, t) such that $x < \hat{\beta}(t)$. We define a conditional exercise density recursively as

$$g_n^{\hat{\beta}}(s|x,t) \equiv \frac{\phi(A_{s,t}) A_{s,t}}{s-t} - \int_t^s \frac{\phi(B_{s,h}) B_{s,h}}{s-h} g_n^{\hat{\beta}}(h|x,t)\, dh, \tag{11}$$

where

$$A_{s,t} \equiv \frac{\hat{\beta}(s) - x - \mu_n (s-t)}{\sigma_n \sqrt{s-t}}, \qquad B_{s,h} \equiv \frac{\hat{\beta}(s) - \hat{\beta}(h) - \mu_n (s-h)}{\sigma_n \sqrt{s-h}}.$$

Given that the boundary $\hat{\beta}$ is bounded, the process X_n reaches it in finite time with probability 1 for non-negative drift. As a consequence, $g_n^{\hat{\beta}}(s|x,t)$ forms a probability density over $(t, +\infty)$.[12] Notice that this density is defined over time and not the position of player n. In particular, this density is used for describing how likely player n is to reach this exercise boundary for the first time at any future interval of time given her current state. Importantly, it does not take into account the possibility that the game might end before that interval is reached.

Therefore, to fully describe the expected payoff from following this arbitrary boundary, one also needs to incorporate the equilibrium arrival rate of the defeat of player n. From inspection of the HJB equation in the interior of the region where player n does not exercise, as in equation 2, one can notice that this arrival rate plays a role that is analogous to an increase in the discount rate. Motivated by this observation, we define the effective discount factor below.

Definition 4. The effective instantaneous discount rate for player n is defined as $r + \lambda_n(t)$ and her effective discount factor between times t and $s > t$ is given by $e^{-\rho_n(s,t)}$, in which

$$\rho_n(h,t) \equiv \int_t^h [r + \lambda_n(s)]\, ds. \tag{12}$$

Once the effective discount and the exercise probability distributions have been defined, we can write the continuation payoff from following $\hat{\beta}$.

Definition 5. The continuation payoff from following a bounded exercise boundary $\hat{\beta}$ given a current state (x,t) is given by

$$W_n^{\hat{\beta}}(x,t) \equiv \int_t^\infty e^{-\rho_n(h,t)} \left(\hat{\beta}(h) - K_n\right) g_n^{\hat{\beta}}(h|x,t)\, dh.$$

[12]For negative drift, X_n fails to cross $\hat{\beta}$ with positive probability. As a result, $g_n^{\hat{\beta}}(s|x,t)$ integrates to less than 1 over $(t, +\infty)$. To define a proper density over the extended reals, one needs to assign a positive probability mass at $+\infty$.

Conditional on a commitment to this arbitrary boundary for $s > t$, instantaneously optimal exercise at time t requires

$$\hat{\beta}(t) - K_n \equiv W_n^{\hat{\beta}}\left(\hat{\beta}(t), t\right) = \int_t^{\infty} e^{-\rho_n(h,t)} \left(\hat{\beta}(h) - K_n\right) g_n^{\hat{\beta}}\left(h|\hat{\beta}(t), t\right) dh. \tag{13}$$

Equation 13 is a necessary requirement for optimality. Given any continuation payoff induced by a policy, the instantaneously optimal policy is given by a trigger. Immediate exercise should occur whenever the surplus from it, *i.e.* $X_n(t) - K_n$, exceeds this critical value. The trigger itself is given by the continuation value described in the right-hand side of that equation.

This reasoning provides us with an interpretation of the optimal instantaneous exercise trigger $\hat{\beta}(t)$. This trigger is such that the surplus from current exercise exactly matches the expected discounted surplus from future exercise. Discounting here is given by the effective discount factor, which takes into account the probability of defeat. Expectations are taken with respect to the density defined over the future exercise times of player n.

Notice, however, that equation 13 is not sufficient for optimality. For example, the perfectly competitive extreme $\hat{\beta}(t) = K_n$, $\forall t$, always satisfies it despite being inconsistent with an equilibrium. Indeed, as typically the case in the optimal stopping literature, optimality also requires smooth-pasting, so that

$$1 = \left.\frac{\partial W_n^{\hat{\beta}}(x,t)}{\partial x}\right|_{x=\hat{\beta}(t)}, \forall t. \tag{14}$$

The main technical difficulty in dealing directly with equations 13 and 14 lies in the differentiation of the conditional exercise density with respect to the current state. We circumvent that problem by using truncated Laplace transforms in the derivation of the proposition below.

Proposition 4. *Under Assumption 1, the equilibrium exercise boundary of each player $n \in \boldsymbol{N}$ satisfies the following integro-differential equation*

$$\beta_n(t) - K_n = \int_t^{\infty} e^{-\rho_n(h,t)} \frac{\phi\left(\frac{\beta_n(h)-\beta_n(t)-\mu_n(h-t)}{\sigma_n\sqrt{h-t}}\right)}{\sigma_n\sqrt{h-t}} \tag{15}$$

$$\left\{\sigma_n^2 + \left[\left(\frac{\beta_n(h) - \beta_n(t)}{h-t}\right) - 2\frac{d\beta_n(h)}{dh} + \mu_n\right](\beta_n(h) - K_n)\right\} dh,$$

while the value function in the interior of the no-exercise region is described by

$$V_n(x,t) = \frac{1}{2}\int_t^{\infty} e^{-\rho_n(h,t)} \frac{\phi\left(\frac{\beta_n(h)-x-\mu_n(h-t)}{\sigma_n\sqrt{h-t}}\right)}{\sigma_n\sqrt{h-t}} \left\{\sigma_n^2 + \left[\left(\frac{\beta_n(h)-x}{h-t}\right) - 2\frac{d\beta_n(h)}{dh} + \mu_n\right](\beta_n(h)-K_n)\right\} dh. \tag{16}$$

Proposition 4 characterizes the equilibrium exercise boundary as a fixed point of an operator on the right-hand side. Notice that equation (15) does not require the separate computation of the evolution of the density over future exercise times: it is embedded in the operator. This is a feature which is common to some analytic representations of the value of American call-options, as derived by McKean (1965), Kim (1990), and Jamshidian (1992).[13] Moreover, the value function is fully determined by the behavior of the boundary.

3.5. **Steady-state.** In this section, we characterize the long-run equilibrium dynamics of the equilibrium. We show that optimal exercise thresholds and arrival rates converge towards a steady state. Moreover, we can provide an explicit description of the limit in terms of the exogenous parameters of the model:

Proposition 5. *Let $(\tau_1, ..., \tau_N)$ be a MPE with exercise boundaries $(\beta_1, ..., \beta_N)$. Then, for every player $n \in \boldsymbol{N}$, we have*

$$\lim_{t\to+\infty} \beta_n(t) = \beta_n^*,$$

where β_n^* is the optimal exercise boundary of a monopolist with (modified) discount rate

$$r_n^* \equiv r + \frac{1}{2}\sum_{m\neq n} \frac{\max\{\mu_m, 0\}^2}{\sigma_m^2}. \tag{17}$$

Proof. In the appendix. □

Proposition 5 reveals the long-run determinants of equilibrium strategies. To provide some intuition, let's start with the simplest case of two players, $N = 2$, and no drift, $\mu_n = 0$, $\forall n \in N$.

Consider an opponent that is supposed, at time t, to be at either at either a low x_l state or a high x_h state, with $x_l < x_h \leq \beta(t)$. As time passes from t to $t + \Delta$, random increments affect

[13] See Chiarella et al. (2004) for a survey of the integral representations for American financial options.

this agent's payoffs and any density at $x \in \{x_l, x_h\}$ is diffused across other states. Neighboring states exchange more mass and distant ones, less so. Naturally, an agent at an x_t that is close to the boundary is more likely to cross it over the next time interval. As the information that the opponent did not cross the boundary between t and $t + \Delta$ arrives, belief updates indicate that she was relatively more likely to have been at the low $x_t = x_l$ than at the high $x_t = x_h$ value, which contributes relatively more to the mass absorbed by the boundary. Therefore information that no crossing occurred is interpreted as news of a higher likelihood of lower states both in the recent past (t) as in the present ($t + \Delta$). Under the assumption of no drift, this means that the conditional belief distribution over the opponent's states loses mass close to the boundary and becomes flatter. As a consequence, it becomes increasingly less likely that an opponent that has never previously exercised her option will do it over a future time interval. As both players engage in this reasoning, exercise thresholds converge towards the policies chosen by monopolists.

Once a positive drift is present, the passage of time pushes probability masses towards higher states, that is, in the direction of the exercise boundary. So, while a signal that an opponent was at a lower state in the past moves the belief density away from the boundary, a positive drift works as a countervailing force. The resultant of these two opposing forces is a well defined limit with an arrival rate which is strictly above zero and is also increasing in the opponent's innovation drift. Additionally, it is also decreasing in the volatility of the opponent's innovations, as a higher volatility forces the belief update in a direction that favors lower states.

Equation 17 has also consequences for industry-wide limit dynamics. Suppose an industry defined by fast innovation processes, represented by high μ_n for some of the players. This industry becomes more competitive in the limit, effective discount rates increase, and products are brought to market under lower profit expectations than they would be absent concerns about competition. As the value functions are forward looking, that increased competition is also propagated towards the transition phase. A similar conclusion follows from an increase in the number of opponents.

4. Simulations

In this section, we present results from simulations and comparative dynamics. First, we compute the equilibrium for a simple symmetric 2-player set-up. We normalize the payoff units to set the exercise cost to unity, i.e., $K = 1$, and the initial condition to $x_0 = 0$. To provide a clear meaning to time, we set the reference time unit to a year and the interest rate $r = 10\%$. We then calibrate 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 ($x_t = K$) within the first two years. The second condition is that the out of the remaining histories, half should cross it within the next four years. We obtain $\mu = 0.04$ and $\sigma = 0.96$.

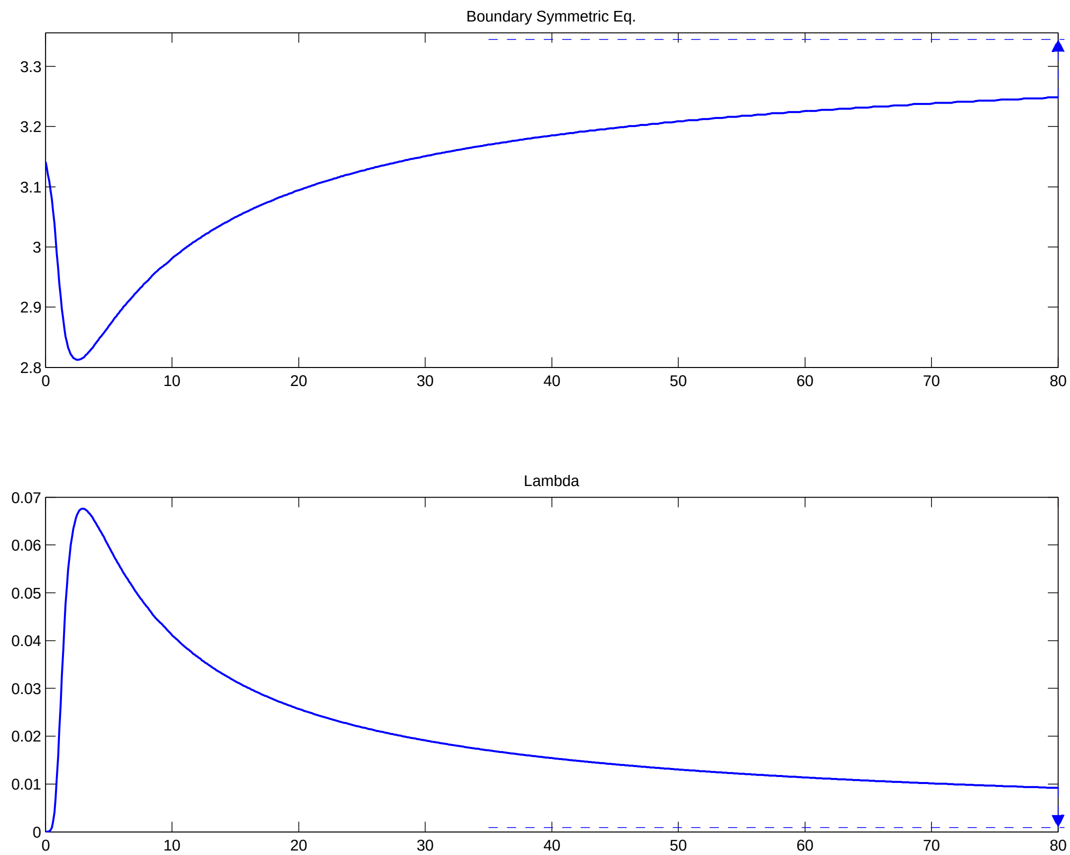

FIGURE 2. Baseline Equilibrium Characterization.
Symmetric parameters set to $K = 1$, $x_0 = 0$, $\mu = 0.04$, and $\sigma = 0.96$. The arrows and dotted lines mark asymptotic limits.

Figure 2 plots the symmetric equilibrium exercise boundaries and the hazard rates. The dotted lines mark 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 hazard rate of almost 7 percentage points, which means that the effective discount rate can be increased by up to 70% relative to the original discount rate of $r = 10\%$. Notice that this magnitude would get significantly larger in the presence of more opponents.

Second, as the value function is forward-looking, the exercise boundary anticipates changes in the hazard rate, hitting its most aggressive point of approximately $\beta(t) = 2.8$ before the hazard rate reaches its peak. It then recedes towards the steady state value of $\lim_{t\to+\infty}\beta(t) = 3.35$. For these baseline parameter values, the zero net-present value boundary is given by $\underline{\beta} = 1$, while the monopoly boundary is $\overline{\beta} = 3.36$. We can see then that the variation in the equilibrium exercise boundaries over time covers more than a fifth of that range. Therefore, while it is well-known that uncertainty can create a large distance between 0-NPV rules and optimal exercise, this simulation exercise shows that gap can be greatly reduced in the presence of short-term competition, while still converging very close to its maximum in the long-run.

Last, another striking feature of the simulation is that convergence towards the steady state is very slow. In the later phase, hazard rates display half-lives that are more than decades-long. Nonetheless, most of the quantitatively meaningful effects are restricted to the first twenty years.

We next 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.

4.1. **An initial lead.** We now study the case in which Player i has a technological lead. She starts at $x_0^i = 0.5$, half the original distance from zero net-present value. The opponent, Player j, still starts at $x_0^j = 0$. The initial lead of player i is common knowledge to both players and all other parameters are kept the same as in the previous section. Results are plotted in Figure 3.

A lead for Player i would, all else held constant, increase the hazard rate of the defeat imposed on Player j. If Player j did not change her exercise boundary, Player i would still be subject to the same arrival rates of defeat and would not have any incentives to change her exercise boundary, which does not depend on the initial condition. Nevertheless, as a consequence of the improved initial condition, she would still be more likely to hit that same boundary earlier. In the presence of a more likely early defeat, Player j has incentives to become more aggressive, increasing the

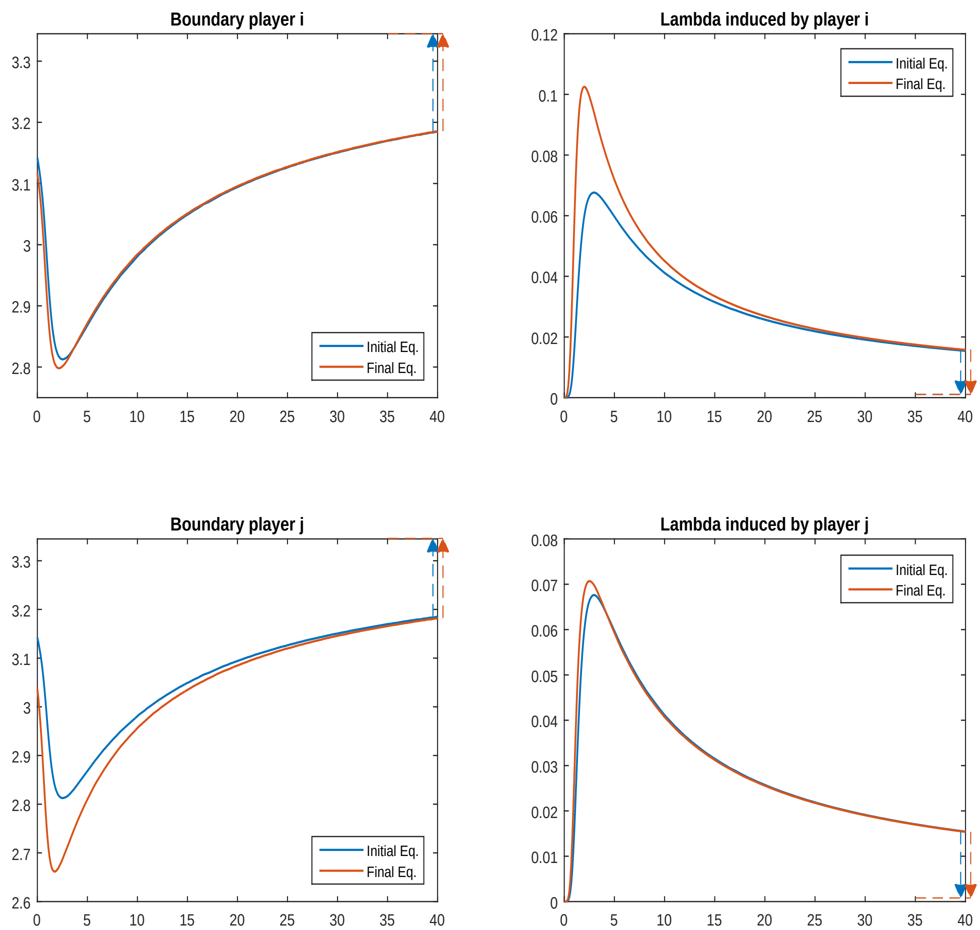

FIGURE 3. Equilibrium comparison with an initial lead for Player i. The arrows and dotted lines indicate asymptotic limits.

likelihood of an early exercise. To this, Player i has incentives to reply with a more aggressive (lower) exercise boundary.

The final equilibrium consequences can be seen in Figure 3. In the equilibrium with a initial lead for Player i, both agents behave more aggressively. Hazard rates increase uniformly, making early entry by any player more likely. Interestingly, most of the quantitative response of the equilibrium boundaries is concentrated on Player j, since the arrival rates of her defeat (which are induced by Player i) respond much more strongly. Asymptotically, however, the effects of the initial lead vanish.

4.2. **A technological lead and faster product development.** We now suppose that one player, Player i, has faster payoff improvements than Player j. In particular, $\mu_i = 0.08$ is twice the benchmark rate to which Player j is still subjected, $\mu_j = 0.04$.

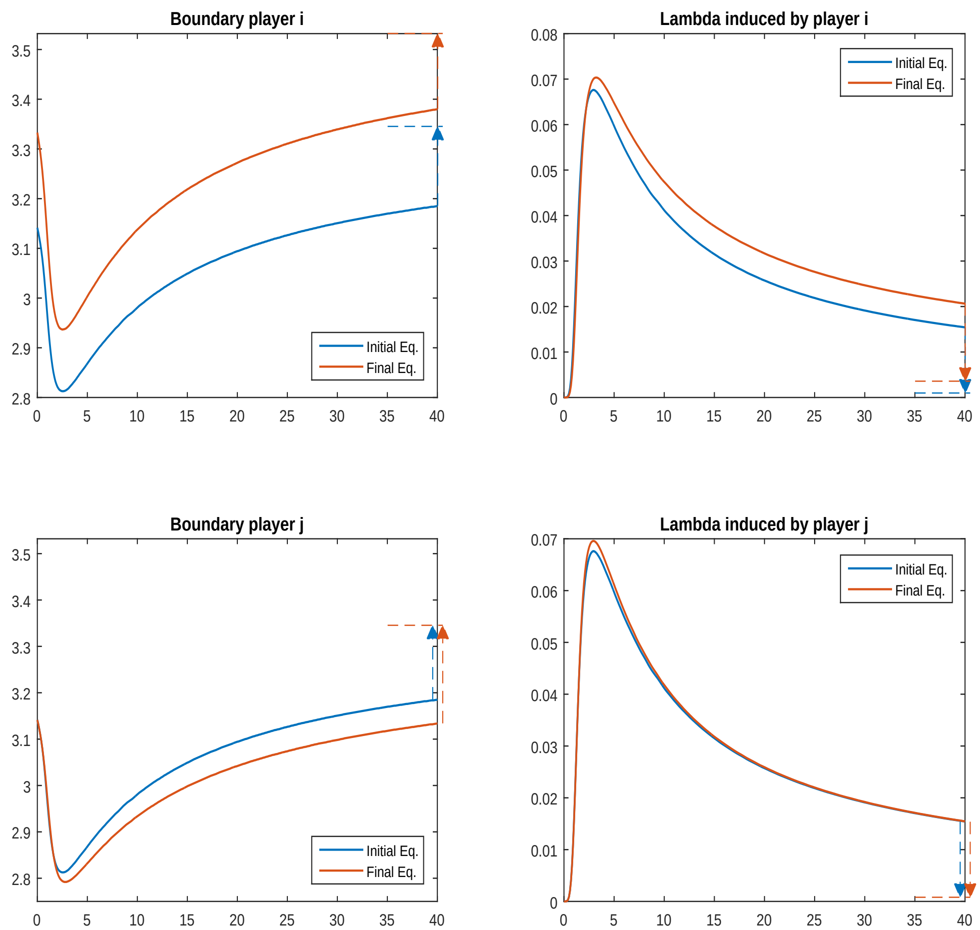

FIGURE 4. Equilibrium comparison when Player i is subject to larger expected payoff increments.
The arrows and dotted lines indicate asymptotic limits.

Given that Player i is subjected to faster payoff improvements, she has always weakly higher incentives to wait instead of exercising earlier. As a consequence, we can see in the top left panel of Figure 4 that her optimal exercise threshold becomes uniformly less aggressive (higher). Nonetheless, given the original increase in the payoff drift, she still arrives faster at that boundary and imposes a higher defeat rate on Player j, as seen in the top-right panel.

Given this, Player j has incentives to behave more aggressively, and the change of her boundary, through an equilibrium feedback effect that turns out to be quantitatively weak, helps partially offset the dovish incentives that a higher drift creates for Player i.

In this case, unlike in the case of a simple initial lead, there are asymptotic effects. The higher drift means that Player i is more intensely pushed against her own boundary asymptotically. Although Player j replies with a boundary that converges asymptotically to a higher value as a response, that has no consequences on the arrival rate of defeat that she imposes on Player i in the limit, which only depends on Player j's own drift and volatility, not on the level of the asymptotic threshold, as indicated by Equation 17.

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. It is illustrated in Figure 5. 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 i, it keeps λ_i (the defeat rate as imposed by player j) fixed. Notice that an increased drift would make this firm less aggressive, as illustrated by the upward displacement of the boundary relative to the baseline (lower drift) situation. In equilibrium, however, as in the previous exercise, despite this less aggressive boundary, the higher rate of innovation increases the perceived intensity of competition. This effect, therefore, dampens the tendency for less aggressive behavior. The line labeled final equilibrium illustrates then that industries with higher rates of innovation face higher entry cut-offs, but not high enough as to revert a direct effect that a higher drift has of increasing hazard rates (which can be seen on the panel on the right).

4.3. **Increased randomness in payoff evolution.** As a final exercise we study comparisons between a symmetric industry with a lower exposure to randomness in the payoff process (labeled initial equilibrium in Figure 6) and one with a higher level (labeled final equilibrium). This enhanced effect of uncertainty can originate either from more volatile market conditions or from more uncertainty in the product development stage.

We can notice that the increased volatility raises the option value from delayed entry, leading to less aggressive exercise strategies. The traditional intuition from non-competitive environments

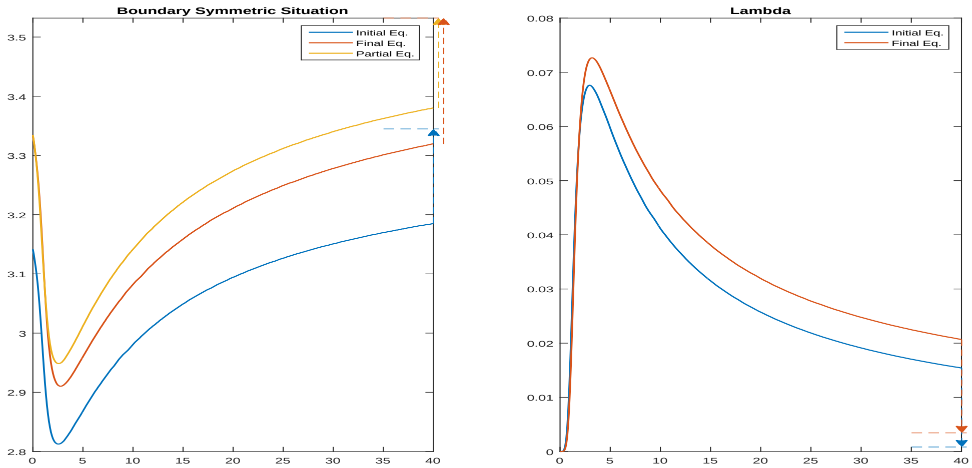

FIGURE 5. Consequences of symmetric doubling of drift in the payoff process, from $\mu = 0.04$ (Initial Equilibrium) to $\mu = 0.08$ (Final Equilibrium).
Partial Equilibrium refers to a situation in which beliefs about opponents exercise rates are kept fixed, but own drift is believed to be at the new level. The arrows and dotted lines indicate asymptotic limits.

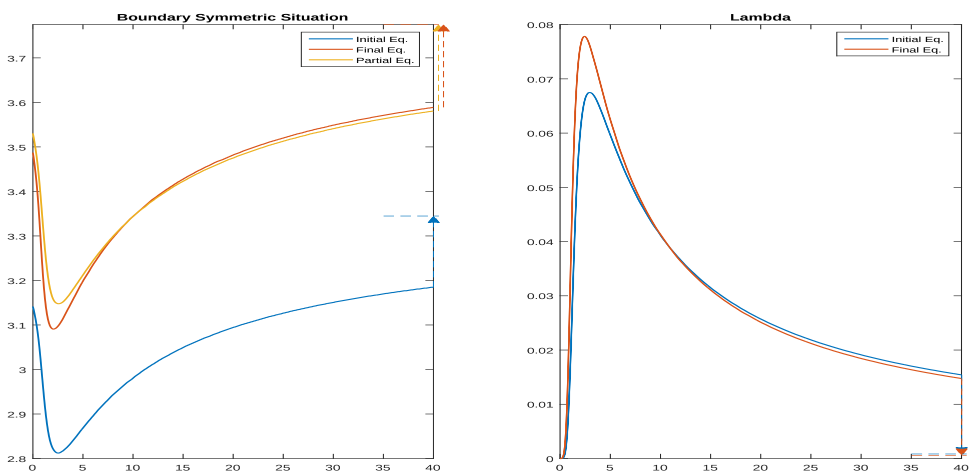

FIGURE 6. A cross industry comparison between a high volatility industry (final equilibrium) and a lower volatility one (initial equilibrium).
Partial Equilibrium refers to a situation in which beliefs about the opponent's exercise rates are kept fixed at the initial equilibrium, but own volatility is believed to be at the high level. The arrows and dotted lines indicate asymptotic limits.

is exhibited in the line marked as partial equilibrium, in the left panel, as it ignores the change in competition but takes into account the consequences of an increased volatility in a firm's own product development process.

As the right-hand side panel indicates, in the short-run, the direct effect of a higher volatility pushing agents more strongly against any exercise boundary dominates, increasing short-run exercise hazard rates despite the less aggressive exercise strategies. In the long-run, however, more volatility means that an opponent that has not previously exercised is very unlikely to be close to exercising in the near future. As a consequence, long-run competition becomes less intense in more volatile industries. We can, therefore, conclude that the equilibrium effects from more volatile product development conditions are not symmetric over time, as higher uncertainty tends to intensify entry and competition in the short-run, while having the opposite effect in case entry is not observed in the initial years.

5. Conclusion

We study a competitive framework in which agents are privately informed about the evolution of the profitability of an investment option. Our model naturally extends the canonical investment under uncertainty model by incorporating the possibility that this opportunity might expire due to strategic preemption by an opponent.

In an equilibrium, each player's optimal exercise boundary depends on the belief this player has about the arrival rate of a possible defeat. This belief changes over time, as the passage of time without any option exercise is itself informative about the conditions of opponents. We develop methods for characterizing the dynamics of player beliefs and equilibrium exercise strategies. These methods are likely to be useful for understanding dynamic strategic behavior in other classes of games.

We show that competition has economically important dynamics, first picking up and then receding towards a stationary long-run situation. In future work, the framework we study can be used to address questions related to optimal technological development policies and the value of information in technological competition. We illustrate some of this potential with equilibrium computation and comparative dynamics exercises, deriving some initial lessons for applied work.

First, competition can have large consequences on the effective rate of time discount. Also, the intensity of competition changes dramatically over time and the transition dynamics towards a steady-state are very long-lived. As a consequence, any rule-of-thumb analysis that focuses on either constant effective discount rates in an attempt to incorporate the consequences of competition or even restricts attention to analytical stationary limits can lead to large errors.

References

Achdou, Yves, Francisco J Buera, Jean-Michel Lasry, Pierre-Louis Lions, and Benjamin Moll, "Partial differential equation models in macroeconomics," *Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences*, 2014, *372* (2028), 20130397.

Bensoussan, Alain, Jens Frehse, and Phillip Yam, *Mean field games and mean field type control theory*, Springer, 2013.

Chiarella, Carl, Andrew Ziogas, and Adam Kucera, *A survey of the integral representation of American option prices*, University of Technology Sydney, 2004.

Dixit, Avinash K and Robert S Pindyck, "Investment under uncertainty, 1994," *Princeton UP, Princeton*, 1994.

Dutta, Prajit K and Aldo Rustichini, "A theory of stopping time games with applications to product innovations and asset sales," *Economic Theory*, 1993, *3* (4), 743–763.

Fudenberg, Drew and Jean Tirole, "Preemption and rent equalization in the adoption of new technology," *The Review of Economic Studies*, 1985, *52* (3), 383–401.

Grenadier, Steven R, "The strategic exercise of options: development cascades and overbuilding in real estate markets," *The Journal of Finance*, 1996, *51* (5), 1653–1679.

———, "Option exercise games: the intersection of real options and game theory," *Journal of Applied Corporate Finance*, 2000, *13* (2), 99–107.

———, "Option exercise games: An application to the equilibrium investment strategies of firms," *Review of financial studies*, 2002, *15* (3), 691–721.

Hopenhayn, Hugo A and Francesco Squintani, "Preemption games with private information," *The Review of Economic Studies*, 2011, *78* (2), 667–692.

Jamshidian, Farshid, "An analysis of American options," *Review of Futures Markets*, 1992, *11* (1), 72–80.

Kim, In Joon, "The analytic valuation of American options," *Review of financial studies*, 1990, *3* (4), 547–572.

Lambrecht, Bart and William Perraudin, "Real options and preemption under incomplete information," *Journal of Economic Dynamics and Control*, 2003, *27* (4), 619–643.

Lions, Pierre-Louis and Jean-Michel Lasry, "Mean field games," 2007.

Mcdonald, Robert and Daniel Siegel, "The value of waiting to invest," *The Quarterly Journal Economics*, 1986, *101* (4), 707–728.

McKean, Henry P, "Appendix: A free boundary problem for the heat equation arising from a problem in mathematical economics," *Sloan Management Review*, 1965, *6* (2), 32.

Quah, John K-H and Bruno Strulovici, "Discounting, values, and decisions," *Journal of Political Economy*, 2013, *121* (5), 896–939.

Strulovici, Bruno and Martin Szydlowski, "On the smoothness of value functions and the existence of optimal strategies in diffusion models," *Journal of Economic Theory*, 2015.

Thijssen, Jacco JJ, "Preemption in a real option game with a first mover advantage and player-specific uncertainty," *Journal of Economic Theory*, 2010, *145* (6), 2448–2462.

Weeds, Helen, "Strategic delay in a real options model of R&D competition," *The Review of Economic Studies*, 2002, *69* (3), 729–747.

Appendix A

Proofs omitted from the main text.

Proof of Proposition 1. It is easy to show that a value matching and payoff monotonicity conditions hold, so that V_n and β_n satisfy $V_n(\overline{\beta}_n, t) = \overline{\beta}_n - k_n$ and

$$\beta_n(t) = \inf \{x \in \mathbb{R} | V_n(x, t) \leq x - k_n\}.$$

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. Since $k_n = \underline{\beta}_n > \beta_n$, we have $V(\beta_n(t_0), t_0) < 0$. This cannot happen in equilibrium as never exercising (i.e., $\tau_n = +\infty$) 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 2. $F_n(x_n, t)$ counts all the Brownian paths which lie in $(-\infty, x_n]$ at time t and are strictly below β_n for all times in $[0, t)$. Note that Γ_n is the first-passage distribution associated to boundary β_n and $\Phi\left(\frac{x_n - \beta_n(h) - \mu_n(t-h)}{\sigma_n\sqrt{t-h}}\right)$ counts all the Brownian paths which start at $\beta_n(h)$ at time h and lie in $(-\infty, x_n]$ at time t. Therefore, we obtain

$$F_n(x_n, t) = \Phi\left(\frac{x_n - x_n^0 - \mu_n t}{\sigma_n \sqrt{t}}\right)$$

$$-\int_0^t \Phi\left(\frac{x_n - \beta_n(h) - \mu_n(t-h)}{\sigma_n\sqrt{t-h}}\right) d\Gamma_n(h).$$

Under Assumption 1, Γ_n is continuously differentiable with respect to time and we can therefore compute the space derivative

$$f_n(x_n,t) = \frac{1}{\sigma_n\sqrt{t}}\phi\left(\frac{x_n - x_n^0 - \mu_n t}{\sigma_n\sqrt{t}}\right)$$
$$-\int_0^t \frac{1}{\sigma_n\sqrt{t-h}}\phi\left(\frac{x_n - \beta_n(h) - \mu_n(t-h)}{\sigma_n\sqrt{t-h}}\right)\gamma_h(h)\,dh.$$

Another differentiation w.r.t. x_n yields:

$$\frac{\partial f_n(x_n,t)}{\partial x_n} = \frac{1}{\sigma_n^2 t}\phi'\left(\frac{x_n - x_n^0 - \mu_n t}{\sigma_n\sqrt{t}}\right)$$
$$-\int_0^t \frac{1}{\sigma_n^2(t-h)}\phi'\left(\frac{x_n - \beta_n(h) - \mu_n(t-h)}{\sigma_n\sqrt{t-h}}\right)\gamma_h(h)\,dh.$$

It is easy to verify that $\phi'(z) = -z\phi(z)$ for all $z \in \mathbb{R}$. Moreover, it is well-known from the theory of Kolmogorov Forward Equations that

$$\gamma_n(t) = -\sigma_n^2 \left.\frac{\partial f_n(x_n,t)}{\partial x_n}\right|_{x_n=\beta_n(t)}.$$

Combining these two facts, the results follow immediately. □

Proof of Proposition 4. The value function satisfies

$$(r+\lambda_n(t))V_n = \mu_n\frac{\partial V_n}{\partial x} + \frac{\sigma_n^2}{2}\frac{\partial^2 V_n}{\partial x^2} + \frac{\partial V_n}{\partial t}$$

for all (x,t) in the interior of the no-exercise region. Define an auxiliary function $\tilde{V}_n$ by setting $\tilde{V}_n(x,t) := V_n(x,t)\mathbb{1}\{x < \beta_n(t)\}$ for each $(x,t) \in \mathbb{R}\times\mathbb{R}_+$. Since $\beta_n(t) \in \left[k, \overline{\beta}_n\right]$ by Proposition 1, the auxiliary function $\tilde{V}_n(\cdot,t)$, unlike the value function itself, is absolutely integrable for every $t \geq 0$:

$$\int_{-\infty}^{+\infty}\left|\tilde{V}_n(x,t)\right|dx \leq \int_{-\infty}^{\beta_n(t)} V_n(x,t)dx \leq \int_{-\infty}^{\overline{\beta}_n}\overline{V}_n(x)dx = \xi_n^{-2} < +\infty,$$

where $\overline{V}_n$ is the value that player n would obtain as a monopolist and ξ_n is the only positive root of equation $(1/2)\sigma_n^2\xi_n^2 + \mu_n\xi_n - r = 0$. Let L_n denote the Fourier transform of $\tilde{V}_n$:

$$L_n \equiv L_n(\omega,t) := \int_{-\infty}^{+\infty} e^{-i\omega x}\tilde{V}_n(x,t)dx = \int_{-\infty}^{\beta_n(t)} e^{-i\omega x}V_n(x,t)dx.$$

Using the HJB equation, value-matching and smooth pasting, it is easy to verify that L_n satisfies the ODE

$$\frac{\partial L_n}{\partial t} = \delta_n L_n - \psi_n,$$

where

$$\delta_n \equiv \delta_n(\omega,t) := r + \lambda_n(t) + \frac{1}{2}\sigma_n^2\omega^2 - \mu_n i\omega.$$

and

$$\psi_n \equiv \psi_n(\omega,t) := e^{-i\omega\beta_n(t)}\left[\left(\mu_n + \frac{1}{2}\sigma_n^2 i\omega - \beta_n'(t)\right)(\beta_n(t) - K_n) + \frac{1}{2}\sigma_n^2\right].$$

Note that Proposition 5 implies the V_n converges. Thus, the ODE above has a unique forward solution:

$$L_n(\omega,t) = \int_t^\infty e^{-\int_t^h \delta_n(\omega,s)ds}\psi(\omega,h)dh.$$

We now proceed to invert this transform using standard inversion formulas. Exchanging the order of integrals, we can write:

$$\begin{aligned}
&\frac{1}{2\pi}\int_{-\infty}^{+\infty} e^{i\omega x} L_n(\omega,t)d\omega \\
&= \int_t^{\infty}\left(\frac{1}{2\pi}\int_{-\infty}^{+\infty} e^{i\omega x}\left[e^{-\int_t^h \delta_n(\omega,s)ds}\psi(\omega,h)\right]d\omega\right)dh, \\
&= \int_t^{\infty} e^{-\rho_n(h,t)}\left(\frac{1}{2\pi}\int_{-\infty}^{+\infty} e^{i\omega x}\left[e^{\mu_n(h-t)i\omega-\frac{1}{2}\sigma_n^2(h-t)\omega^2}\psi(\omega,h)\right]d\omega\right)dh.
\end{aligned}$$

Thus, we only need to compute the Fourier inverse

$$\frac{1}{2\pi}\int_{-\infty}^{+\infty} e^{i\omega x}\left[e^{\mu_n(h-t)i\omega-\frac{1}{2}\sigma_n^2(h-t)\omega^2}\psi(\omega,h)\right]d\omega = \int_{-\infty}^{+\infty} e^{-A\omega^2+B\omega}\left(C\omega+D\right)d\omega,$$

where

$$\begin{aligned}
A &:= \frac{1}{2}\sigma_n^2(h-t), \\
B &:= \left[x-\beta_n(h)+\mu_n(h-t)\right]i, \\
C &:= \frac{1}{4\pi}\sigma_n^2\left(\beta_n(h)-K_n\right)i, \\
D &:= \frac{1}{2\pi}\left[\left(\mu_n-\beta_n'(h)\right)\left(\beta_n(h)-K_n\right)+\frac{1}{2}\sigma_n^2\right].
\end{aligned}$$

It is easy to show that

$$\int_{-\infty}^{+\infty} e^{-A\omega^2+B\omega}\left(C\omega+D\right)d\omega = e^{\frac{B^2}{4A}}\sqrt{\frac{\pi}{A}}\left(\frac{BC}{2A}+D\right).$$

Note that

$$e^{\frac{B^2}{4A}} = e^{-\frac{1}{2}\left(\frac{\beta_n(h)-x-\mu_n(h-t)}{\sigma_n\sqrt{h-t}}\right)^2} = \sqrt{2\pi}\phi\left(\frac{\beta_n(h)-x-\mu_n(h-t)}{\sigma_n\sqrt{h-t}}\right),$$

and

$$\frac{BC}{2A}+D = \frac{1}{4\pi}\left\{\sigma_n^2+\left[\left(\frac{\beta_n(t)-x}{h-t}\right)-2\beta_n'(t)+\mu_n\right]\left(\beta_n(t)-K_n\right)\right\}.$$

Thus,

$$\begin{aligned}
&e^{\frac{B^2}{4A}}\sqrt{\frac{\pi}{A}}\left(\frac{BC}{2A}+D\right) = \\
&\frac{\phi\left(\frac{\beta_n(h)-x-\mu_n(h-t)}{\sigma_n\sqrt{h-t}}\right)}{2\sigma_n\sqrt{h-t}}\left\{\sigma_n^2+\left[\left(\frac{\beta_n(h)-x}{h-t}\right)-2\beta_n'(h)+\mu_n\right]\left(\beta_n(h)-K_n\right)\right\}.
\end{aligned}$$

From the theory of Fourier transforms, we know that

$$\frac{\tilde{V}_n(x-,t)+\tilde{V}_n(x+,t)}{2} = \int_t^{\infty} e^{-\rho_n(h,t)}e^{\frac{B^2}{4A}}\sqrt{\frac{\pi}{A}}\left(\frac{BC}{2A}+D\right)dh$$

for all $x\in\mathbb{R}$. On the one hand, in the no-exercise region, we have $\tilde{V}_n(\cdot,t)=V_n(\cdot,t)$. Thus, Equation 16 follows immediately from continuity of the value function. On the other hand, when $x=\beta_n(t)$, we have

$$\frac{\tilde{V}_n(x-,t)+\tilde{V}_n(x+,t)}{2} = \frac{V_n(\beta_n(t)-,t)+0}{2} = \frac{1}{2}V_n(\beta_n(t),t)$$

by definition of $\tilde{V}_n$ and value-matching. As a result,

$$V_n(\beta_n(t),t) = 2\,e^{\frac{B^2}{4A}}\sqrt{\frac{\pi}{A}}\left(\frac{BC}{2A}+D\right)\Big|_{x=\beta_n(t)}.$$

Evaluating the formula for the RHS above, we obtain Equation 15. □

In order to prove Proposition 5, we need the following two lemmas.

Lemma 1. *Let β_F and β_G are the optimal exercise boundaries of a player facing exogenous stopping distributions F and G. Fix $t \in \mathbb{R}_+$ and suppose that*

$$\frac{1-F(t+h)}{1-F(t)} \geq \frac{1-G(t+h)}{1-G(t)}$$

for all $h \in \mathbb{R}_+$. Then, $\beta_F(t) \geq \beta_G(t)$.

Proof. There is no loss in restricting attention to stopping times τ such that $X^x(\tau) \geq k$ a.s. For every such stopping time, we have

$$\mathbb{E}\left\{e^{-r\tau}\left(\frac{1-F(t+\tau)}{1-F(t)} - \frac{1-G(t+\tau)}{1-G(t)}\right)(X_n^x(\tau)-k)\right\} \geq 0.$$

This implies that the recursive values satisfy $V_F(x,t) \geq V_G(x,t)$ for all $x \in \mathbb{R}$. Hence, $V_G(x,t) > x-k$ implies $V_F(x,t) > x-k$. Hence, optimality of continuation under G implies that continuation is also optimal under F. It follows that $\beta_F(t) \geq \beta_G(t)$. □

Lemma 2. *We define, for each $t, h \in \mathbb{R}_+$ and player $n \in \boldsymbol{N}$,*

$$\Lambda_n(t,h) := -\ln\left(\frac{1-\Gamma_n(t+h)}{1-\Gamma_n(t)}\right).$$

Then, for every player $n \in \boldsymbol{N}$ and $h \in \mathbb{R}_+$, we have

$$\lim_{t\to+\infty}\left(\frac{\Lambda_n(t,h)}{h}\right) = \frac{1}{2}\left(\frac{\max\{\mu_n, 0\}}{\sigma_n}\right)^2.$$

Proof. Optimal exercise boundaries satisfy $k_n \leq \beta_n \leq \beta_n^M$ for every $n \in \boldsymbol{N}$. Let Γ_n^K and Γ_n^M be the absorption probabilities associated with constant exercise boundaries k_n and β^M. Clearly, $\Gamma_n^M(t) \leq \Gamma_n(t) \leq \Gamma_n^K(t)$ 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

$$\limsup_{t\to+\infty} \Lambda_n(t,h) \leq Ah. \tag{18}$$

Clearly, $\Gamma_n^M(t) < \Gamma_n^k(t)$ for all $t > 0$. Hence, for every $t > 0$ and $h \in \mathbb{R}_+$, we have

$$\frac{1-\Gamma_n(t+h)}{1-\Gamma_n(t)} > \frac{1-\Gamma_n^k(t+h)}{1-\Gamma_n^M(t)}$$

Thus,

$$\Lambda_n(t,h) < -\ln\left(\frac{1-\Gamma_n^k(t+h)}{1-\Gamma_n^M(t)}\right).$$

Using L'Hôpital's rule, we can explicitly compute:

$$\lim_{t\to+\infty}\left(\frac{1-\Gamma_n^k(t+h)}{1-\Gamma_n^M(t)}\right) = \begin{cases} \exp\left(\frac{\mu_n}{\sigma_n^2}\left(\beta_n^M - k_n + \frac{1}{2}\mu_n h\right)\right)\left(\frac{k_n - x_n^0}{\beta_n^M - x_n^0}\right) & \mu_n \geq 0 \\ \frac{1-\exp\left(\frac{2\mu_n(k_n - x_n^0)}{\sigma_n^2}\right)}{1-\exp\left(\frac{2\mu_n(\beta_n^M - x_n^0)}{\sigma_n^2}\right)} & \mu_n < 0. \end{cases}$$

It follows that

$$\limsup_{t\to+\infty} \Lambda_n(t,h) \leq \lim_{t\to+\infty}\left[-\ln\left(\frac{1-\Gamma_n^k(t+h)}{1-\Gamma_n^M(t)}\right)\right],$$

$$= -\ln\left[\lim_{t\to+\infty}\left(\frac{1-\Gamma_n^k(t+h)}{1-\Gamma_n^M(t)}\right)\right],$$
$$= Ah + B,$$

where we define

$$A := \begin{cases} \frac{1}{2}\left(\frac{\mu_n}{\sigma_n}\right)^2 & \mu_n > 0 \\ 0 & \mu_n \le 0 \end{cases}$$

and

$$B := \begin{cases} \frac{\mu_n}{\sigma_n^2}\left(\beta_n^M - k_n\right) + \ln\left(\frac{\beta_n^M - x_n^0}{k_n - x_n^0}\right) & \mu_n > 0 \\ 0 & \mu_n = 0 \\ \ln\left(\frac{1-\exp\left(\frac{2\mu_n(k_n - x_n^0)}{\sigma_n^2}\right)}{1-\exp\left(\frac{2\mu_n(\beta_n^M - x_n^0)}{\sigma_n^2}\right)}\right) & \mu_n < 0. \end{cases}$$

Running a symmetric argument, we can obtain the following lower bound for the limit inferior:

$$\liminf_{t\to+\infty}\Lambda_n(t,h) \ge Ah - B,$$

where A and B are defined as before.

Next, we will show that, for all $h \in \mathbb{R}_+$, we have

$$\lim_{t\to+\infty}\Lambda_n(t,h) = Ah.$$

Suppose first that the limit on the left-hand side above exists for some $h \in \mathbb{R}_+$. Then, for every $m \in \mathbb{N}$, we have

$$\Lambda_n(t,mh) = -\ln\left(\frac{1-\Gamma_n(t+mh)}{1-\Gamma_n(t)}\right),$$
$$= -\ln\left[\prod_{i=1}^m\left(\frac{1-\Gamma_n(t+ih)}{1-\Gamma_n(t+(i-1)h)}\right)\right],$$
$$= \sum_{i=1}^m\left[-\ln\left(\frac{1-\Gamma_n(t+ih)}{1-\Gamma_n(t+(i-1)h)}\right)\right],$$
$$= \sum_{i=1}^m\Lambda_n(t+ih,h).$$

This formally implies that

$$\lim_{t\to+\infty}\Lambda_n(t,mh) = \lim_{t\to+\infty}\sum_{i=1}^m\Lambda_n(t+ih,h),$$
$$= \sum_{i=1}^m\lim_{t\to+\infty}\Lambda_n(t+ih,h),$$
$$= m\lim_{t\to+\infty}\Lambda_n(t,h).$$

Reversing the derivation proves that the limit in the LHS must also exist. It follows that

$$\lim_{t\to+\infty}\Lambda_n(t,mh) = \liminf_{t\to+\infty}\Lambda_n(t,mh) = \limsup_{t\to+\infty}\Lambda_n(t,mh).$$

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

$$Amh - B \leq \lim_{t\to+\infty} \Lambda_n(t, mh) \leq Amh + B.$$

Combined with the linearity derived above, this implies that

$$Ah - \frac{B}{m} \leq \lim_{t\to+\infty} \Lambda_n(t, h) \leq Ah + \frac{B}{m}.$$

Since this inequality holds for every $m \in \mathbb{N}$, we must have

$$Ah \leq \lim_{t\to+\infty} \Lambda_n(t, h) \leq Ah.$$

This establishes that, if the limit exists, it must be in fact equal to Ah.

Finally, we consider the existence of the limit. It is enough to show that $\lim_{i\to+\infty} \Lambda_n(t_i, h) = Ah$ holds for every increasing unbounded sequence of times $\{t_i\}$. Note that a real sequence converges to the number $Ah \in \mathbb{R}$ if and only if every sub-sequence has a convergent sub-sub-sequence converging to Ah. So fix any increasing sequence of times $\{t_i\}$ such that $\lim_{i\to+\infty} t_i = +\infty$ and define $\{\alpha_i\}$ as

$$\alpha_i := \Lambda_n(t_i, h).$$

Since α_i eventually belongs to the compact interval $[0, Ah + B + 1]$, there is no loss in assuming that this holds for all $i \in \mathbb{N}$. Now pick an arbitrary sub-sequence $\{\alpha_{i_j}\}$. Since $\{\alpha_{i_j}\}$ is contained in a compact interval, $\{\alpha_{i_j}\}$ has a convergent sub-sub-sequence $\{\alpha_{i_{j_l}}\}$. Since $\lim_{l\to+\infty} \alpha_{i_{j_l}}$ exists by construction, we can run the argument above to show that $\lim_{l\to+\infty} \alpha_{i_{j_l}} \to Ah$. Since the sub-sequence $\{\alpha_{i_j}\}$ was arbitrary, we conclude that $\lim_{i\to+\infty} \alpha_i = Ah$, completing the proof. $\square$

Proof of Proposition 5. Consider player $n \in \{1, ..., N\}$. Proposition 1 implies that $\hat{\beta}_n$ must satisfy $k_n \leq \hat{\beta}_n(t) \leq \beta_n^M$ for all $t \in \mathbb{R}_+$. By Lemma 2, for every $h \in \mathbb{R}_+$, we have

$$\lim_{t\to+\infty} \left(\frac{\Lambda_n(t, h)}{h}\right) = \frac{1}{2}\left(\frac{\max\{\mu_n, 0\}}{\sigma_n}\right)^2.$$

Now, for every $t \in \mathbb{R}_+$, define

$$\overline{\lambda}_n(t) := \sup_{\substack{s\geq t\\ h\geq 0}} \left(\frac{\Lambda_n(s, h)}{h}\right)$$

and

$$\underline{\lambda}_n(t) := \inf_{\substack{s\geq t\\ h\geq 0}} \left(\frac{\Lambda_n(s, h)}{h}\right).$$

Notice that, for all $t, h \in \mathbb{R}_+$, we have

$$\underline{\lambda}_n(t) \leq \frac{\Lambda_n(t, h)}{h} \leq \overline{\lambda}_n(t).$$

Moreover,

$$\lim_{t\to+\infty} \underline{\lambda}_n(t) = \lim_{t\to+\infty} \overline{\lambda}_n(t) = \frac{1}{2}\left(\frac{\max\{\mu_n, 0\}}{\sigma_n}\right)^2.$$

Define $\overline{\beta}_n(t)$ and $\underline{\beta}_n(t)$ as the constant boundaries of a monopolist facing constant hazard rates $\sum_{m\neq n} \underline{\lambda}_m(t)$ and $\sum_{m\neq n} \overline{\lambda}_m(t)$, respectively, for all $s \geq t$. Clearly,

$$\lim_{t\to+\infty} \underline{\beta}_n(t) = \lim_{t\to+\infty} \overline{\beta}_n(t) = \beta_n^*.$$

Moreover, since $\underline{\lambda}_n(t)h \leq \Lambda_n(t,h) \leq \overline{\lambda}_n(t)h$ for all $t, h \in \mathbb{R}_+$, Lemma 1 implies that, for all $t \in \mathbb{R}_+$, we have

$$\underline{\beta}_n(t) \leq \hat{\beta}_n(t) \leq \overline{\beta}_n(t).$$

Therefore, we conclude that

$$\lim_{t \to +\infty} \hat{\beta}_n(t) = \beta_n^*,$$

as claimed. □