

The Dynamics of Bertrand Price Competition with Cost-Reducing Investments[†]

Fedor Iskhakov*

CEPAR, University of New South Wales

John Rust[‡]

Georgetown University

Bertel Schjerning[§]

University of Copenhagen

March, 2013

Abstract: We present a dynamic extension of the classic static model of Bertrand price competition that allows competing duopolists to undertake cost-reducing investments in an attempt to “leapfrog” their rival to attain low-cost leadership — at least temporarily. We show that leapfrogging occurs in equilibrium, resolving the *Bertrand investment paradox*, i.e. leapfrogging explains why firms have an *ex ante* incentive to undertake cost-reducing investments even though they realize that simultaneous investments to acquire the state of the art production technology would result in Bertrand price competition in the product market that drives their *ex post* profits to zero. Our analysis provides a new interpretation of “price wars”. Instead of constituting a punishment for a breakdown of tacit collusion, price wars are fully competitive outcomes that occur when one firm leapfrogs its rival to become the new low cost leader. We show that the equilibrium involves *investment preemption* only when the firms invest in a deterministically alternating fashion and technological progress is deterministic. We prove that when technological progress is deterministic and firms move in an alternating fashion, the game has a unique Markov perfect equilibrium. When technological progress is stochastic or if firms move simultaneously, equilibria are generally not unique. Unlike the static Bertrand model, the equilibria of the dynamic Bertrand model are generally inefficient. Instead of having too little investment in equilibrium, we show that duopoly investments generally *exceed* the socially optimum level. Yet, we show that when investment decisions are simultaneous there is a “monopoly” equilibrium when one firm makes all the investments, and this equilibrium is efficient. However, efficient non-monopoly equilibria also exist, demonstrating that it is possible for firms to achieve efficient dynamic coordination in their investments while their customers also benefit from technological progress in the form of lower prices.

Keywords: duopoly, Bertrand-Nash price competition, Bertrand paradox, Bertrand investment paradox, leapfrogging, cost-reducing investments, technological improvement, dynamic models of competition, Markov-perfect equilibrium, tacit collusion, price wars, coordination and anti-coordination games, strategic preemption

JEL classification: D92, L11, L13

[†] This paper is part of the Frisch Centre project “1307 Strategic research programme on pensions” financed by the Ministry of Labour, Norway. We acknowledge helpful comments from Michael Baye, Jeffrey Campbell, Dan Cao, Ulrich Doraszelski, Joseph E. Harrington, Jr., Dan Kovenock, Rober Lagunoff, Stephen Morris, Michael Riordan, Jean-Marc Robin, David Salant, Karl Schmedders, Che-Lin Su, and other participants at seminars at Columbia, University of Copenhagen, CREST, Georgetown, Princeton, the Chicago Booth School of Business, the 3rd CAPCP conference at Pennsylvania State University, the 2012 conferences of the Society of Economic Dynamics, the Society for Computational Economics, the North American Summer Meetings of the Econometric Society, the 4th World Congress of the Game Theory Society, the Initiative for Computational Economics at Chicago and Zurich (ICE 2012 and ZICE 2013), and the NSF/NBER CEME Conference on “The Econometrics of Dynamic Games.” The unabridged version of this paper with proofs of main results is [available for download](#).

* Email: f.iskhakov@unsw.edu.au

[‡] **Correspondence address:** Department of Economics, Georgetown University, Washington, DC, phone: (301) 801-0081, email: jrl393@georgetown.edu

[§] Email: bertel.schjerning@econ.ku.dk

1 Introduction

Given the large theoretical literature on price competition since the original work of Bertrand (1883), it is surprising that we still know relatively little about price competition in the presence of production cost uncertainty. For example Routledge (2010) states that “However, there is a notable gap in the research. There are no equilibrium existence results for the classical Bertrand model when there is discrete cost uncertainty.” (p. 357). Note that Routledge’s analysis and most other recent theoretical models of Bertrand price competition with production cost uncertainty are *static*. Even less is known about Bertrand price competition in dynamic models where firms also compete by undertaking cost-reducing investments. In these environments firms face uncertainty about their rivals’ investment decisions as well as uncertainty about the timing of technological innovations that can affect future prices and costs of production.

This paper analyzes a dynamic version of the static “textbook” Bertrand-Nash duopoly game where firms can make investment decisions as well as pricing decisions. At any point in time, a firm can replace its current production facility with a new state of the art production facility that enables it to produce at a lower marginal cost. The state of the art production cost evolves stochastically and exogenously, but the technology adoption decisions of the firms are fully endogenous. The term *leapfrogging* is used to describe the longer run investment competition between the two duopolists where the higher cost firm purchases a state of the art production technology that reduces its marginal cost relative to its rival and allows it to attain, at least temporarily, a position of low cost leadership. Except for Giovannetti (2001), leapfrogging behavior has been viewed as *incompatible* with Bertrand price competition.

When firms produce goods that are perfect substitutes using constant returns to scale production technologies with no capacity constraints, it is well known that Bertrand equilibrium results in zero profits for the high cost firm. The motivation for the high cost firm to undertake a cost-reducing investment is to obtain a production cost advantage over its rival. The low cost firm does earn positive profits, charging a price equal to the marginal cost of production of its higher cost rival. However, whenever both firms have the same marginal cost of production, the Bertrand equilibrium price is their common marginal cost, and both firms earn zero profits. Baye and Kovenock (2008) describe this as the *Bertrand paradox*.

A new paradox arises when we try to extend the static Bertrand price competition to a dynamic model where cost-reducing investments occur due to exogenous technological improvements. If both firms have equal opportunity to acquire the state of the art technology at the same investment cost, there is no way for firms to ensure anything more than a temporary advantage over their rivals unless they are successful in undertaking some sort of *strategic investment deterrence* or preemptive investments. These markets can be said to be *contestable due to ease of investment* similar to the notion of markets that are *contestable due to ease of entry* of Baumol, Panzar and Willig (1982). Either firm can become a low cost leader by simply acquiring the state of the art technology. However if both firms do this simultaneously, the resulting Bertrand price competition ensures that *ex post* profits are zero. If the firms expect this, the *ex ante* return on their investments will be negative, but then neither firm has any incentive to undertake cost-reducing investments. We refer to this as the *Bertrand investment paradox*.

We solve the Bertrand investment paradox by showing that at least one of the firms will undertake cost-reducing investments in any equilibrium of the model: no investment by both firms can only be an equilibrium outcome when the cost of investing is prohibitively high. Indeed we will show that paradoxically, far from investing too infrequently, investment generally occurs *too frequently* in the equilibria of this model. This equilibrium *over-investment* (which includes *duplicative investments*) is a reflection of coordination failures that is the source of inefficiency in our model.

Thus, our results lead to a new paradox: even though the static Bertrand equilibrium is well known to be efficient (i.e. production is done by the low cost firm), once we consider the simplest dynamic extension of the static Bertrand model, we find that most equilibria are inefficient. We do show that there are efficient equilibria in our model, but we show that these are *asymmetric pure strategy equilibria*. These include *monopoly equilibria* where only one firm never invests and the other does all of the investing and sets a price equal to the marginal cost of production of the high cost, non-investing firm. Although investment competition in the non-monopoly equilibria of the model does benefit consumers by lowering costs and prices in the long run, it does generally come at the cost of some inefficiency due to coordination failures. However, perhaps surprisingly, we show that there also exist fully efficient, non-monopoly, asymmetric, pure strategy equilibria as well.

Our model provides a new interpretation for the concept of a *price war*. Price paths in the equilibria of our model are piece-wise flat, with periods of significant price declines just after one of the firms invests and displaces its rival to become the low cost leader. We call the large drop in prices when this happens a “price war”. However in our model these periodic price wars are part of a fully competitive outcome where the firms are behaving as Bertrand price competitors in every period. Thus, our notion of a price war is very different from the standard interpretation of a price war in the literature, where price wars are a punishment device to deter tacitly colluding firms from cheating. The key difference in the prediction of our model and the standard model of tacit collusion is that price wars in our model are very brief, lasting only a single period, and lead to *permanent* price reductions, whereas in the model of tacit collusion, price wars can extend over multiple periods but the low prices are ultimately only *temporary* since prices are predicted to *rise* at the end of a price war.

We review the previous literature on leapfrogging and Bertrand price competition in section 2, including a conjecture by Riordan and Salant (1994) that there is a unique equilibrium involving *investment preemption* and no leapfrogging equilibria when firms are Bertrand price competitors. We present our model in section 3. The model has a natural “absorbing state” when the improvement in the state of the art cost of production asymptotically achieves its lowest possible value (e.g. a zero marginal cost of production). We show how the solution to the dynamic game can be decomposed starting from the solution to what we refer to as the “end game” when the state of the art marginal cost of production has reached this lowest cost absorbing state. We characterize socially optimal investment strategies in section 4. These results are key to establishing the efficiency of equilibria in the subsequent analysis. We present our main results in section 5 and our conclusions in section 6, where we relate our results to the larger literature on the impact of market structure on innovation and prices, including Schumpeter’s conjecture that there should be greater innovation under monopoly than duopoly or other market structures where firms have less market power.

2 Previous work on leapfrogging: the Riordan and Salant conjecture

The earliest work on leapfrogging that we are aware of include Fudenberg et. al. (1983) and Reinganum (1985). Reinganum introduced a subgame perfect Nash equilibrium model of Schum-

peterian competition with *drastic innovations* in which an incumbent enjoys a temporary period of monopoly power, but is eventually displaced by a challenger whose R&D expenditures result in a new innovation that enables it to leapfrog the incumbent to become the new temporary monopolist. This early work in the literature on *patent races* and models of *research and development* (R&D) focused on the strategies of competing firms where there is a continuous choice of R&D expenditures with the goal of producing a patent or a drastic innovation that could not be easily duplicated by rivals. In Reinganum's model the incumbent invests less than challengers to produce another drastic innovation, and the incumbents overinvest. However Gilbert and Newbery (1982) and Vickers (1986) argued that an incumbent monopolist has a strong incentive to engage in preemptive patenting. In their models there is no leapfrogging, since the incumbent monopolist repeatedly outspends/outbids its rivals to maintain its incumbency.

Our study is less related to this literature than to a separate literature on adoption of cost-reducing innovations, which is binary adopt/don't adopt decision rather than a continuous investment decision. Prior to the work of Giovannetti (2001), the main result in the literature on cost reducing investment under duopoly in the presence of downstream Bertrand price competition was that *leapfrogging investments cannot occur in equilibrium*. Instead, the main result from this literature is that all equilibria involve *strategic preemption*, i.e. a situation where one of the duopolists undertakes all cost-reducing investments at times determined to deter any leapfrogging investments by the firm's rival.

Riordan and Salant (1994) analyzed a dynamic Bertrand duopoly game of pricing and investment very similar to the one we analyze here, except that they assumed that firms move in an alternating fashion and that technological progress is deterministic. Under the continuous time, non-stochastic framework that Riordan and Salant used, equilibrium strategies consist of an *ex ante* fixed sequence of dates at which firms upgrade their production facilities. They proved that "If firms choose adoption dates in a game of timing and if the downstream market structure is a Bertrand duopoly, the equilibrium adoption pattern displays rent-dissipating increasing dominance; i.e. all adoptions are by the same firm and the discounted value of profits is zero." (p. 247). The rent dissipation result can be viewed as a dynamic generalization of the zero profit result in a static symmetric cost Bertrand duopoly. The threat of investment by the high cost firm forces the low cost leader to invest at a sequence of times that drives its discounted profits to zero. However

unlike the static Bertrand equilibrium, this preemption equilibrium is *fully inefficient* — all rent *and* all social surplus is dissipated by the excessively frequent investments of the low cost leader.

Riordan and Salant conjectured that their results did not depend on their alternating move assumption about investments: “These heuristic ideas do not rely on the alternating move structure that underlies our definition of an equilibrium adoption pattern. We believe the same limit results hold if firms move simultaneously at each stage of the discrete games in the definition. The alternating move sequence obviates examining mixed strategy equilibria for some subgames of the sequence of discrete games.” (p. 255). We show that Riordan and Salant’s conjecture is incorrect, though we do characterize the conditions under which there exists a unique Markov Perfect equilibrium involving strategic preemption.

Giovannetti analyzed a discrete-time, simultaneous move game of investment and Bertrand price competition similar to the one we analyze here. He showed that leapfrogging could indeed emerge in his model, and he also found equilibria that do not result in complete preemption but exhibit what he called *increasing asymmetry* where the low cost leader is the only firm that invests over long intervals, increasing its advantage over its high cost rival.

Giovannetti’s analysis was conducted in an environment of deterministic technological progress, similar to Riordan and Salant. We extend both of their analyses by considering a model where a) firms can invest either simultaneously, or in alternating fashion (including a model where the right of move evolves stochastically according to a two state Markov chain), and b) technological improvement occurs stochastically according to an exogenous Markov process (which includes deterministic technological progress as a special case).

We confirm Giovannetti’s finding that there are multiple equilibria when the firms make their decisions simultaneously, but go farther by characterizing (and computing) *all Markov perfect equilibria*. We show that when the state space is continuous (i.e. when the set of possible marginal costs of production is a subinterval of the positive real line) the model has a continuum of equilibria. However when we restrict the model so that the state space is finite (i.e. where there is a finite set of possible marginal costs of production), the number of equilibria depends on the set of allowable *equilibrium selection rules*. If we exclude the possibility of stochastic equilibrium selection rules, then we show there will be a finite (even) number of possible equilibria, and if we allow stochastic equilibrium selection rules, then there are a continuum of possible equilibria of the model even

when the state space is finite.

We also analyze equilibria of a version of the game where firms move alternately in discrete time, enabling us to match our results more closely to Riordan and Salant (1994) since Giovannetti’s analysis was restricted to a simultaneous move model only. We prove that there is a unique equilibrium involving strategic preemption when technological progress is deterministic, but not when technology evolves stochastically.¹ We show that preemption arises when the right to move alternates *deterministically* whereas leapfrogging occurs when it alternates stochastically. Further, we show that full rent dissipation only holds in the limit in *continuous time*. The preempting firm earns positive profits when time is discrete, but these profits tend to zero as the interval between time periods tends to zero, $\Delta t \rightarrow 0$, so that Riordan and Salant’s preemption equilibrium with full rent dissipation emerges as limiting special case of our framework.

In general, there are far fewer equilibria when firms invest alternately than simultaneously. In particular, the two “monopoly equilibria” and the zero expected profit symmetric mixed strategy equilibrium cannot be supported in the alternating move framework.² However when technological progress is stochastic, there are multiple equilibria even in the alternating move version of the game, and most of them involve leapfrogging rather than preemption. Further, since the monopoly equilibria are no longer sustainable in the alternating move setting, it follows that *all* of the equilibria in the alternating move game are inefficient.³

We present simulations of non-monopoly equilibria that involve both simultaneous investments by the firms (at nodes of the game tree where both firms play mixed strategies), as well and alternating investments that can occur under equilibrium selection rules where the firms play pure strategies at nodes, capturing the intuitive notion of leapfrogging investments by the two firms as they vie for temporary low cost leadership. These patterns can emerge regardless of whether the firms are assumed to invest simultaneously or alternately as long as technological progress is stochastic. We find equilibria where one firm exhibits persistent low cost leadership over its oppo-

¹Strictly speaking, the equilibrium is not unique, but rather there are *two equilibria* that differ as to which of the two firms is designated to play the role of preemptive investor.

²We discuss certain exceptions where the monopoly equilibria are sustainable in the alternating move game, including the obvious special case where only one of the firms has the right to invest in every period.

³We have been able to construct examples of equilibria in alternating move games that are fully efficient, but these are in examples with a small number of possible cost states that are not robust to changes in the model parameters, particularly to increasing the number of possible cost states in the model.

ment, and equilibria involving “sniping” where a high cost opponent displaces the low cost leader to become the new (permanent) low cost leader, even though it has spent a long period of time as the high cost follower.

Besides the literature cited above there is some connection of our work to the large literature on patent races and research and development (R&D) (e.g. Reinganum, 1981, 1985). However we believe that there are fundamental differences between models of R&D and patents and our model that involve investments in a cost-reducing technology that improves exogenously (stochastically or deterministically) over time. The R&D decision is inherently more of a continuous investment with an uncertain payoff (with higher investments increasing the probability of success, and improving the distribution of payoffs), whereas our analysis involves a simple binary decision by the firms about whether or not to acquire a new machine or production facility that can enable it to produce at the state of the art marginal cost of production c . Furthermore, patents generate temporary periods of monopoly, and even when R&D capital stocks are not protected by patents they are nevertheless still at least semi-proprietary and thus not easy for their rivals to replicate. These features make patent and R&D competition far less *contestable* and more likely to be subject to various types of inefficiencies than the competition over cost-reducing investments that we analyze in this paper.

Goettler and Gordon (2011) is a recent seminal structural empirical study of Bertrand price setting and R&D competition by Intel and AMD.⁴ They find strong spillover effects and that competition between these firms *reduces* investment spending on R&D and rates of innovation relative to a monopoly situation or the social optimum. This is the opposite of our finding that duopoly competition generally *increases* the rate of adoption of new lower cost technologies compared to a monopolist. Besides spillover effects (which are internalized by the monopolist and social planner), the difference in conclusions may also be driven by the fact that computer chips are *durable goods* and the consumers in Goettler and Gordon’s model face dynamic decisions about when to replace an existing computer and buy a new one, whereas the consumers in our model make static

⁴Another excellent recent analysis of the dynamics of duopoly investment competition with Bertrand price competition is Dubé, Hitsch and Chingagupta (2010) which analyzes investment by producers of *platforms* (producers of video games). Consumers in their model also face much more difficult dynamic decisions about which video console to buy, and these decisions depend on expectations of decisions of current and future software developers that produce video games for the consoles, as well as on the future pricing of the consoles. In this framework interesting and complicated issues of *network externalities* and *market tipping* arise that do not arise in our simpler framework where consumer decisions are assumed to be essentially static.

purchase decisions in every period. When consumers make periodic durable purchases, firms may have an incentive to innovate faster because this can create *technological obsolescence* that can cause consumers to buy durable goods more frequently. This may lead even a monopolist to choose a socially inefficient investment strategy.⁵

3 The Model

We consider a market consisting of two firms producing an identical good. We assume that the two firms are price setters, have no fixed costs and can produce the good at a constant marginal cost of c_1 and c_2 , respectively. We assume constant return to scale production technologies so that neither firm ever faces a binding capacity constraint. While it is more realistic to include capacity constraints as an additional motive for investment, we believe that it is of interest to start by considering the simplest possible extension of the classic Bertrand price competition model to a multi-period setting, ignoring capacity constraints. Binding capacity constraints provide a separate motivation for leapfrogging investments than the simpler situation that we consider here. It is considerably more difficult to solve a model where capacity constraints are both choices and state variables, and we anticipate the equilibria of such a model will be considerably more complex than the ones we find in the simpler setting studied here, and we already find a very complex set of equilibria and equilibrium outcomes.⁶

Our model does allow for switching costs and idiosyncratic factors that affect consumer demand, resulting in less than perfectly elastic demand. When these are present, if one of the firms slightly undercuts its rival's price, it will not capture all of its rival's market share. However we believe it is of interest to consider whether leapfrogging is possible even in the limiting "pure Bertrand" case where consumer demand is perfectly elastic. This represents the most challenging case for leapfrogging, since the severe price cutting incentives unleashed by the purest version of

⁵Rust (1986) showed that monopoly producers of durable goods have an incentive to engage in planned obsolescence, though the monopolist's choice of durability could be either higher or lower than the social optimum. In an R&D context, the analog of planned obsolescence can be achieved by innovating at a faster rate than the social optimum.

⁶Kreps and Scheinkman (1983) showed that in a two period game, if duopolists set prices in period two given capacity investment decisions made in period one, then the equilibrium of this two period Bertrand model is identical to the equilibrium of the static model of Cournot quantity competition. We are interested in whether this logic will persist in a multiple period extension of our model. Kovenock and Deneckere (1996) showed that even in a two period model, if firms have different unit costs of production the Kreps and Scheinkman result can fail to hold.

Bertrand price competition creates the sharpest possible case for the “Bertrand investment paradox” that we noted in the Introduction.

We rule out the possibility of entry and exit and assume that the market is forever a duopoly. Ruling out entry and exit can be viewed as a worst case scenario for the viability of leapfrogging equilibrium, since the entry of a new competitor provides another mechanism by which high cost firms can be leapfrogged by lower cost ones (i.e. the new entrants). We also assume that the firms do not engage in explicit collusion. The equilibrium concept does not rule out the possibility of tacit collusion, though the use of the Markov-perfect solution concept effectively rules out many possible tacitly collusive equilibria that rely on history-dependent strategies and incredible threats to engage in price wars as a means of deterring cheating and enabling the two firms to coordinate on a high collusive price.

On the other hand, we will show that the set of Markov-perfect equilibria is very large, and include equilibria that involve coordinated investments that are reminiscent of tacit collusion. For example, we show there are equilibria where there are long alternating intervals during which one of the firms attains persistent low cost leadership and the opponent rarely or never invests. This enables the low cost leader to charge a price equal to the marginal cost of production of the high cost follower that generates considerable profits. Then after a brief price war in which the high cost follower leapfrogs the low cost leader, the new low cost leader enjoys a long epoch of low cost leadership and high profits.

These alternating periods of muted competition with infrequent price wars resemble tacit collusion, but are not sustained by complex threats of punishment for defecting from a tacitly collusive equilibrium. Instead, these are just examples of the large number of Markov perfect equilibria that can emerge in our model that display a high degree of coordination, even though it is not enforced by any sort of “trigger strategy” or punishment scheme such as are analyzed in the literature on supergames. We also find much more “competitive” equilibria where the firms undertake alternating investments that are accompanied by a series of price wars that successively drive down prices to the consumer while giving each firm temporary intervals of time where it is the low cost leader and thereby the ability to earn positive profits.

3.1 Consumers

Under our assumption of perfectly elastic consumer demand, as is well known, the Bertrand equilibrium for two firms with constant returns to scale production technologies and no capacity constraints is for the lower cost firm to serve the entire market at a price $p(c_1, c_2)$ equal to the marginal cost of production of the higher cost rival

$$p(c_1, c_2) = \max[c_1, c_2]. \quad (1)$$

In the case where both firms have the same marginal cost of production we obtain the classic result that Bertrand price competition leads to zero profits for both firms at a price equal to their common marginal cost of production.

3.2 Production Technology and Technological Progress

We assume that the two firms have the ability to make an investment to acquire a new production facility (plant) to replace their existing plant. Exogenous stochastic technological progress drives down the marginal cost of production of the state of the art production plant over time. Suppose that the technology for producing the good improves according to an exogenous first order Markov process specified below. If the current state of the art marginal cost of production is c , let $K(c)$ be the cost of investing in the plant that embodies this state of the art production technology.

We assume that the state of the art production technology entails constant marginal costs of production (equal to c) and no capacity constraints. Assume there are no costs of disposal of an existing production plant, or equivalently, the disposal costs do not depend on the vintage of the existing plant and are embedded as part of the new investment cost $K(c)$. If either one of the firms purchases the state of the art plant, then after a one period lag (constituting the “time to build” the new production facility), the firm can produce at the new marginal cost c .

We allow the fixed investment cost $K(c)$ to depend on c . This can capture different technological possibilities, such as the possibility that it is more expensive to invest in a plant that is capable of producing at a lower marginal ($K'(c) > 0$), or situations where technological improvements lower both the marginal cost of production c and the cost of building a new plant ($K'(c) < 0$). Clearly, even in the monopoly case, if investment costs are too high, then there may be a point at which

the potential gains from lower costs of production using the state of the art production plant are insufficient to justify incurring the investment cost $K(c)$. This situation is even more complicated in a duopoly, since if the competition between the firms leads to leapfrogging behavior, then neither firm will be able to capture the entire benefit of investments to lower its cost of production: some of these benefits will be passed on to consumers in the form of lower prices. If *all* of the benefits are passed on to consumers, the duopolists may not have an incentive to invest for *any* positive value of $K(c)$. This is the Bertrand investment paradox that we discussed in the introduction.

Let c_t be the marginal cost of production under the state of the art production technology at time t . Each period the firms face a simple binary investment decision: firm j can decide not to invest and continue to produce using its existing production facility at the marginal cost $c_{j,t}$. If firm j pays the investment cost $K(c)$ and acquires the state of the art production plant with marginal cost c_t , then when this new plant comes on line, firm j will be able to produce at the marginal cost $c_t < c_{j,t}$.

Given the one period lag to build a new production facility, if a firm does invest at the start of period t , it will not be able to produce using the new facility until period $t + 1$. If there has been no improvement in the technology since the time firm 1 upgraded its plant, then $c_{2,t+1} \geq c_{1,t+1} = c_t = c_{t+1}$. If a technological innovation occurs at time $t + 1$, then $c_{2,t+1} \geq c_{1,t+1} = c_t > c_{t+1}$ and firm 1's new plant is already slightly behind the frontier the moment it comes online. If c_t is a continuous stochastic process the state space S for this model is the pyramid $S = \{(c_1, c_2, c) | c_1 \geq c \text{ and } c_2 \geq c \text{ and } c \leq c_0\}$ in R^3 , where $c_0 > 0$ is the initial state of technology at $t = 0$, and thus (c_0, c_0, c_0) is the apex of the pyramid. In cases where we restrict c_t to a finite set of possible values in $[0, c_0]$ the state space is a finite subset of S .

Suppose that both firms believe that the state of the art technology for producing the good evolves stochastically according to a Markov process with transition probability $\pi(c_{t+1}|c_t)$. Specifically, suppose that with probability $\pi(c_t|c_t)$ we have $c_{t+1} = c_t$ (i.e. there is no improvement in the state of the art technology at $t + 1$), and with probability $1 - \pi(c_t|c_t)$ the technology improves, so that $c_{t+1} < c_t$ and c_{t+1} is a draw from some distribution over the interval $[0, c_t]$. An example of a convenient functional form for such a distribution is the Beta distribution. However the presentation of the model and most of our results do not depend on specific functional form assumptions about π except that we prove the uniqueness of equilibrium when firms invest in an alternating

fashion and $\pi(c_t|c_t) = 0$ when $c_t > 0$.⁷ We will refer to this special case as *strictly monotonic* technological improvement, i.e. a situation where the state of the art improves in every period until it reaches the absorbing state, $c_t = 0$. Throughout the paper we will use the condition $\pi(c_t|c_t) = 0$ to denote this case, bearing in mind that it only applies when $c_t > 0$.

Note that completely deterministic technological progress is characterized by the condition $\pi(c_{t+1}|c_t) \in \{0, 1\}$ for any c_t, c_{t+1} , and therefore should be distinguished from strictly monotonic improvement. Apart from the case where $\pi(c_t|c_t) = 1$ and technological progress stops at some c_t , deterministic technological improvement is strictly monotonic, but not vice versa.

3.3 Timing of Moves and Unobserved Investment Cost Shocks

Let $m_t \in \{0, 1, 2\}$ be a state variable that governs which of the two firms are allowed to undertake an investment at time t . We will assume that $\{m_t\}$ evolves as an exogenous two state Markov chain with transition probability $f(m_{t+1}|m_t)$ independent of the other state variables $(c_{1,t}, c_{2,t}, c_t)$. As noted above, we assume that in every period, the firms simultaneously set their prices after having made their investment choices. However their investment choices may or may not be made simultaneously. The value $m_t = 0$ denotes a situation where the firms make their investment choices simultaneously, $m_t = 1$ indicates a state where only firm 1 is allowed to invest, and $m_t = 2$ is the state where only firm 2 can invest.

In this paper we analyze two variants of the game: 1) a *simultaneous move* investment game where $m_1 = 0$ and $f(0|m_t) = 1$ (so $m_t = 0$ with probability 1 for all t), and 2) *alternating move* investment games, with either deterministic or random alternation of moves, but where there is no chance that the firms could ever undertake simultaneous investments (i.e. where $m_1 \in \{1, 2\}$ and $f(0|m_t) = 0$ for all t). Under either the alternating or simultaneous move specifications, each firm always observes the investment decision of its opponent after the investment decision is made. However, in the simultaneous move game, the firms must make their investment decisions based on their assessment of the probability their opponent will invest. In the alternating move game, since only one of the firms can invest at each time t , the mover can condition its decision on the

⁷The state $c_t = 0$ is a natural absorbing state for the marginal cost, so $\pi(0|0) = 1$. However we can also allow for an absorbing state to occur at a positive level of costs that represent the minimum physically possible cost of production under any technology.

investment decision of its opponent if it was the opponent's turn to move in the previous period. The alternating move specification can potentially reduce some of the strategic uncertainty that arises in a fully simultaneous move specification of the game.

We interpret random alternating moves as a way of reflecting *asynchronicity* of timing of decisions in a discrete time model that can occur in continuous time models that have the property that the probability of two firms making investment decisions at the exact same instant of time is zero, an idea that Doraszelski and Judd (2011) have shown can dramatically simplify the calculation of equilibria of the game. In discrete time alternating move games, there are cases where equilibrium has been shown to be unique (see, e.g. Lagunoff and Matsui, 1997 or Bowlus and Seitz, 2006). We are interested in conditions under which unique equilibria will emerge in asynchronous move versions of our model.

The timing of events in the model is as follows. At the start of period t each firm knows the costs of production $(c_{1,t}, c_{2,t})$, and both learn the current values of c_t and m_t . If $m_t = 0$, then the firms simultaneously decide whether or not to invest. We assume that both firms know each others' marginal cost of production, i.e. there is common knowledge of state $(c_{1,t}, c_{2,t}, c_t, m_t)$. Further, both firms have equal access to the new technology by paying the investment cost $K(c_t)$ to acquire the current state of the art technology with marginal cost of production c_t . We also consider versions of the model where each firm j , $j \in \{1, 2\}$ incurs idiosyncratic "disruption costs" $\epsilon_{t,j} = (\epsilon_{0,t,j}, \epsilon_{1,t,j})$ associated with each of the choices of not to invest $(\epsilon_{0,t,j})$ and investing $(\epsilon_{1,t,j})$, respectively. These shocks are private information to each firm j , though there is common knowledge that these shocks are distributed independently of each other with common bivariate density $q(\epsilon)$. In our analysis below q will actually be a Type I extreme value distribution, though our analysis also applies to other distributions as well.

These costs, if negative, can be interpreted as benefits to investing. Benefits may include things such as temporary price cuts in the investment cost $K(c)$, tax benefits, or government subsidies that are unique to each firm. Let $\eta\epsilon_{t,1}$ be the idiosyncratic disruption costs involved in acquiring the state of the art production technology for firm 1, and let $\eta\epsilon_{t,2}$ be the corresponding costs for firm 2, where $\eta \geq 0$ is a scaling parameter.

For tractability, we assume that it is common knowledge among the two firms that $\{\epsilon_{t,1}\}$ and $\{\epsilon_{t,2}\}$ are independent *IID* Type I bivariate extreme value processes with common scale parameter

$\eta \geq 0$. Firm j observes its current and past idiosyncratic investment shocks $\{\varepsilon_{t,j}\}$, but does not observe its future shocks or its opponent's past, present, or future idiosyncratic investment cost shocks. After each firm decides whether or not to invest in the latest technology, the firms then independently and *simultaneously* set the *prices* for their products, where production is done in period t with their existing plant. The Bertrand equilibrium price function $p(c_1, c_2) = \max[c_1, c_2]$ is the unique Nash equilibrium outcome of the simultaneous move pricing stage game.

The one period time-to-build assumption implies that even if both firms invest in new plants at time t , their marginal costs of production in period t are $c_{1,t}$ and $c_{2,t}$, respectively, since they have to wait until period $t + 1$ for the new plant to be installed, and must produce in period t using their existing plants. However in period $t + 1$ we have $c_{1,t+1} = c_t$ and $c_{2,t+1} = c_t$, since the new plants the firms purchased in period t have now become operational.

Our analysis will also consider limits of equilibria for sequences of games in which the time period between successive moves tends to zero. This enables us to use our discrete time framework to approximate the equilibria of continuous time versions of the game. In particular we are interested in drawing a correspondence between our results and those of Riordan and Salant (1994) who proved their key result about strategic preemption as a continuous time limit of a sequence of discrete time games where the time between moves tends to zero.

Let the time interval between moves be $\Delta t > 0$. To ensure comparability of results, it is important that we take continuous limits in the same way as Riordan and Salant did, and this requires that payoffs and the firms' discount factors depend on Δt in the manner that they assumed. Thus, similar to Riordan and Salant (1994) we assume that the discount factor is $\beta = \exp\{-r\Delta t\}$ and we multiply per-period flow payoffs by Δt also to reflect that profits earned per period tend to zero as the length of each period Δt tends to zero. Further, we need to specify how the discrete time transition probability for technological progress, $\pi(c_{t+1}|c_t)$, depends on the time interval Δt . To match the case of deterministic technological progress that Riordan and Salant analyzed, where technological improvement of the state of the art production process in continuous time was assumed to be governed by a deterministic non-increasing function $c(t)$, we consider corresponding discrete time deterministic sequences for the state of the art costs given by

$$(c_0, c_1, c_2, c_3, \dots) = (c(0), c(\Delta t), c(2\Delta t), c(3\Delta t), \dots). \quad (2)$$

3.4 Solution concept

Assume that the two firms are expected discounted profit maximizers and have a common discount factor $\beta \in (0, 1)$. The relevant solution concept that we adopt for this dynamic game between the two firms is the standard concept of *Markov-perfect equilibrium* (MPE).

In a MPE, the firms' investment and pricing decision rules are restricted to be functions of the current state, $(c_{1,t}, c_{2,t}, c_t, m_t)$. When there are multiple equilibria in this game, the Markovian assumption restricts the "equilibrium selection rule" to depend only on the current value of the state variable. We will discuss this issue further below.

The firms' pricing decisions only depend on their current production costs $(c_{1,t}, c_{2,t})$ in accordance with the static Bertrand equilibrium outcome. However the firms' investment decisions also depend on the value of the state of the art marginal cost of production c_t and the designated mover m_t . Further, if $\eta > 0$, each firm $j \in \{1, 2\}$ has private information about cost shocks $(\epsilon_{0,t,j}, \epsilon_{1,t,j})$ affecting its decision to not invest, or invest, respectively.

Definition 1. *A Stationary Markov Perfect Equilibrium of the duopoly investment and pricing game consists of a pair of strategies $(P_j(c_1, c_2, c), p_j(c_1, c_2))$, $j \in \{1, 2\}$ where $P_j(c_1, c_2, c, m) \in [0, 1]$ is firm j 's probability of investing and $p_j(c_1, c_2) = \max[c_1, c_2]$ is firm j 's pricing decision. The investment rules $P_j(c_1, c_2, c, m)$ must maximize the expected discounted value of firm j 's future profit stream taking into account then investment and pricing strategies of its opponent.*

Note that under our assumptions, consumer purchases of the good is a purely static decision, and consequently there are no dynamic effects of pricing for the firms, unlike in the cases of durable goods where consumer expectations of future prices affects their timing of new durable purchases which cause pricing decisions to be a fully dynamic decision as in Goettler and Gordon (2011). Thus in our case, the pricing decision is given by the simple static Bertrand equilibrium in every period, $p_t = \max[c_{1,t}, c_{2,t}]$. The only dynamic decision in our model is firms' investment decisions. We allow the investment strategies of the firms to be probabilistic to allow for the possibility of mixed strategy equilibria.

When $\eta > 0$ the firms' investment strategies also depend on their privately observed cost shocks. Thus, we can write these strategies as $P_j(c_1, c_2, c, m, \epsilon_{0,j}, \epsilon_{1,j})$, for $j \in \{1, 2\}$. Note that since $(\epsilon_{0,j}, \epsilon_{1,j})$ are cost shocks affecting *new investment* rather than *production* using the firm's

existing production technology, the firms' pricing strategies remain functions only of (c_1, c_2) just as in the case where $\eta = 0$. Since the private shocks to investment $(\epsilon_{0,j}, \epsilon_{1,j})$ are continuously distributed, then when $\eta > 0$ with probability 1 the firms will be using *pure strategies*.⁸ When we integrate over the idiosyncratic shocks $(\epsilon_{0,j}, \epsilon_{1,j})$ affecting firms' investment choices, we obtain *conditional investment probabilities* $P_j(c_1, c_2, c, m)$ for firms $j = \{1, 2\}$. The choice probabilities in the case $\eta > 0$ resemble mixed strategies that can exist in the case $\eta = 0$. It is not hard to show that as $\eta \rightarrow 0$, the set of conditional investment probabilities that arise in various stage game equilibria for the incomplete information game when $\eta > 0$ will converge to the set of pure and mixed strategies in the limiting complete information game where $\eta = 0$.

To derive the functional equations characterizing a stationary Markov-perfect equilibrium, we focus on the situation where each firm faces the same cost $K(c)$ of investment, though it is straightforward to allow one of the firms to have an *investment cost advantage*.⁹ Suppose the current (mutually observed) state is (c_1, c_2, c, m) , i.e. firm 1 has a marginal cost of production c_1 , firm 2 has a marginal cost of production c_2 , and the marginal cost of production using the current best technology is c and m denotes which of the firms (or both if $m = 0$) has the right to make a move and invest. Since we have assumed that the two firms can both invest in the current best technology at the same cost $K(c)$, it is tempting to conjecture that there should be a "symmetric equilibrium" where by "symmetric" we mean an equilibrium where the decision rule and value function for firm 1 depends on the state (c_1, c_2, c) , and similarly for firm 2, and these value functions and decision rules are *anonymous* (also called *exchangeable*) in the sense that

$$V_1(c_1, c_2, c, m, \epsilon_0, \epsilon_1) = V_2(c_2, c_1, c, m, \epsilon_0, \epsilon_1) \quad (3)$$

where $V_1(c_1, c_2, c, m, \epsilon_0, \epsilon_1)$ is the value function for firm 1 when the mutually observed state is (c_1, c_2, c, m) , and the privately observed costs/benefits for firm 1 for investing and not investing in the current state of the art technology are ϵ_0 and ϵ_1 , respectively. V_2 is the corresponding value

⁸Doraszelski and Escobar (2010) have also used additively-separable *iid* unobservable cost shocks as a computational device to *purify* dynamic games where players have discrete decisions and where there is a potential for mixed strategy equilibria.

⁹In this case there would be two investment cost functions, K_1 and K_2 , and firm 1 would have an investment cost advantage if $K_1(c) \leq K_2(c)$ for all $c \geq 0$. It is also possible to extend the model to allow for higher investment costs for the firm that falls behind: in such case we can use investment cost functions $K_1(c, c_1, c_2)$ and $K_2(c, c_1, c_2)$ such that $K_1(c, c_1, c_2) > K_2(c, c_1, c_2)$ if $c_1 > c_2$, and vice versa if $c_1 < c_2$.

function for firm 2, but where $m^\sim$ denotes the complementary state to m .¹⁰

Thus, the intuitive content of the symmetry condition (3) is that the value function for the firms only depend on the values of the state variables, not on their identities or the arbitrary labels “firm 1” and “firm 2”. However it turns out that most of the “interesting” equilibria of the game are *asymmetric*, i.e. where the symmetry condition condition (3) does not hold. For this reason our analysis will not restrict attention only to symmetric equilibria of the game, but rather our main conclusions and results about the existence of leapfrogging and the efficiency of equilibria will be based on an analysis of *all* stationary equilibria of the game.

When $\eta > 0$ and the idiosyncratic cost/benefits from investing or not investing are $(\epsilon_{0,j}, \epsilon_{1,j})$ for $j = 1, 2$, it is not hard to show that the value functions V_j , $j = 1, 2$ take the form

$$V_j(c_1, c_2, c, m, \epsilon_{0,j}, \epsilon_{1,j}) = \max[v_{I,j}(c_1, c_2, c, m) + \eta \epsilon_{0,j}, v_{N,j}(c_1, c_2, c, m) + \eta \epsilon_{1,j}] \quad (4)$$

where, when $m = 0$, $v_{N,j}(c_1, c_2, c, m)$ denotes the expected value to firm j if it does not invest in the latest technology, and $v_{I,j}(c_1, c_2, c, m)$ is the expected value to firm j if it invests. However when $m \in \{1, 2\}$, the functions $v_{N,j}$ and $v_{I,j}$ have slightly different interpretations. When $m = 1$ (so it is firm 1’s turn to invest), then $v_{I,1}(c_1, c_2, c, 1)$ and $v_{N,1}(c_1, c_2, c, 1)$ denote the expected values to firm 1 if it does and does not invest, respectively, and this is identical to the interpretation of form the values to firm 1 in the simultaneous move case. However when $m = 1$ it is not firm 2’s turn to invest, so $v_{I,2}(c_1, c_2, c, 1)$ and $v_{N,2}(c_1, c_2, c, 1)$ are interpreted as the expected future value to firm 2 *conditional on firm 1’s investment choice*. Thus, $v_{N,2}(c_1, c_2, c, 1)$ represents firm 2’s expected discounted profits given that it is firm 1’s turn to invest and that firm 2 learns that firm 1 did not invest at its turn. Similarly $v_{I,2}(c_1, c_2, c, 1)$ represents firm 2’s expected discounted profits given that firm 1 invested on its turn.

Let $r_1(c_1, c_2)$ be the expected profits that firm 1 earns in a single period equilibrium play of the Bertrand-Nash pricing game when the two firms have costs of production c_1 and c_2 , respectively.

$$r_1(c_1, c_2) = \begin{cases} 0 & \text{if } c_1 \geq c_2 \\ \max[c_1, c_2] - c_1 & \text{otherwise.} \end{cases} \quad (5)$$

¹⁰For the simultaneous move game, $m = 0$ and $m^\sim = 0$, but in an alternating move game, if $m = 1$ then $m^\sim = 2$ and if $m = 2$ then $m^\sim = 1$.

and the profits for firm 2, $r_2(c_1, c_2)$ are defined symmetrically, so we have $r_2(c_1, c_2) = r_1(c_2, c_1)$.¹¹

The formula for the expected profits associated with *not* investing (after taking expectations over player j 's privately observed idiosyncratic shocks (ϵ_0, ϵ_1) but conditional on the publicly observed state variables (c_1, c_2, c)) is given by:

$$v_{N,j}(c_1, c_2, c, m) = r_j(c_1, c_2) + \beta EV_j(c_1, c_2, c, m, 0), \quad (6)$$

where $EV_j(c_1, c_2, m, c)$ denotes the conditional expectation of firm j 's value function (next period) $V_j(c_1, c_2, c, m, \epsilon_{0,j}, \epsilon_{1,j})$ given that it does not invest this period (represented by the last 0 argument in EV_j), conditional on the current state (c_1, c_2, c, m) .

The formula for the expected profits associated with investing (after taking conditional expectations over firm j 's privately observed idiosyncratic shocks $(\epsilon_{0,j}, \epsilon_{1,j})$ but conditional on the publicly observed state variables (c_1, c_2, c, m)) is given by

$$v_{I,j}(c_1, c_2, c, m) = r_j(c_1, c_2) - K(c) + \beta EV_j(c_1, c_2, c, m, 1), \quad (7)$$

where $EV_j(c_1, c_2, c, m, 1)$ is firm j 's conditional expectation of its next period value function given in equation (4) and given that it invests (the last argument, now equal to 1), conditional on (c_1, c_2, c, m) .

To compute the conditional expectations $EV_j(c_1, c_2, c, m, 0)$ and $EV_j(c_1, c_2, c, m, 1)$ we invoke a well known property of the extreme value family of random variables — “max stability” (i.e. a family of random variables closed under the max operator). The max-stability property implies that the expectation over the idiosyncratic IID cost shocks $(\epsilon_{0,j}, \epsilon_{1,j})$ is given by the standard “log-sum” formula when these shocks have the Type I extreme value distribution. Thus, after taking expectations over $(\epsilon_{0,j}, \epsilon_{1,j})$ in the equation for V_j in (4) above, we have

$$\begin{aligned} & \int_{\epsilon_0^j} \int_{\epsilon_1^j} V_j(c_1, c_2, c, m, \epsilon_{0,j}, \epsilon_{1,j}) q(\epsilon_{0,j}) q(\epsilon_{1,j}) d\epsilon_{1,j} d\epsilon_{0,j} = \\ & \eta \log \left[\exp\{v_{N,j}(c_1, c_2, c, m)/\eta\} + \exp\{v_{I,j}(c_1, c_2, c, m)/\eta\} \right]. \end{aligned} \quad (8)$$

The log-sum formula provides a closed-form expression for the conditional expectation of the value functions $V_j(c_1, c_2, c, m, \epsilon_{0,j}, \epsilon_{1,j})$ for each firm j , where V_j is the maximum of the value of

¹¹When we consider limiting equilibria of discrete time games when the time interval $\Delta t \downarrow 0$, we multiply the single period profit flows r_j by Δt , so that $r_j(c_1, c_2)$ should be interpreted as the profit flow earned over a unit time interval $\Delta t = 1$, and thus $r_j(c_1, c_2, c)\Delta t$ is the profit flow earned by firm j over a time interval $\Delta t < 1$.

not investing or investing as we can see from equation (4) above. This means that we do not need to resort to numerical integration to compute the double integral in the left hand side of equation (8) with respect to the next-period values of $(\epsilon_{0,j}, \epsilon_{1,j})$. However we do need to compute the two functions $v_{N,j}(c_1, c_2, c, m)$ and $v_{I,j}(c_1, c_2, c, m)$ for both firms $j = 1, 2$. We will describe one algorithm for doing this below.

To simplify notation, we let $\phi(v_{N,j}(c_1, c_2, c, m), v_{I,j}(c_1, c_2, c, m))$ be the log-sum formula given above in equation (8), that is define ϕ as¹²

$$\phi(v_{N,j}(c_1, c_2, c, m), v_{I,j}(c_1, c_2, c, m)) \equiv \eta \log [\exp\{v_{N,j}(c_1, c_2, c, m)/\eta\} + \exp\{v_{I,j}(c_1, c_2, c, m)/\eta\}]. \quad (9)$$

Let $P_1(c_1, c_2, c, m)$ be firm 2's belief about the probability that firm 1 will invest if the mutually observed state is (c_1, c_2, c, m) . Consider first the case where $m = 0$, so the two firms move simultaneously in this case. Firm 1's investment decision is probabilistic from the standpoint of firm 2 because firm 1's decision depends on the cost benefits/shocks $(\epsilon_{0,1}, \epsilon_{1,1})$ that only firm 1 observes. But since firm 2 knows the probability distribution of these shocks, it can calculate P_1 as the following binary logit formula

$$P_1(c_1, c_2, c, m) = \frac{\exp\{v_{I,1}(c_1, c_2, c, m)/\eta\}}{\exp\{v_{N,1}(c_1, c_2, c, m)/\eta\} + \exp\{v_{I,1}(c_1, c_2, c, m)/\eta\}} \quad (10)$$

Firm 2's belief of firm 1's probability of not investing is of course simply $1 - P_1(c_1, c_2, c, m)$. Firm 1's belief of the probability that firm 2 will invest, $P_2(c_1, c_2, c, m)$ is defined similarly.

$$P_2(c_1, c_2, c, m) = \frac{\exp\{v_{I,2}(c_1, c_2, c, m)/\eta\}}{\exp\{v_{N,2}(c_1, c_2, c, m)/\eta\} + \exp\{v_{I,2}(c_1, c_2, c, m)/\eta\}} \quad (11)$$

It is not hard to show that

$$\lim_{\eta \rightarrow 0} P_1(c_1, c_2, c, m) = I\{v_{I,1}(c_1, c_2, c, m) \geq v_{N,1}(c_1, c_2, c, m)\}, \quad (12)$$

and a similar limit holds for $P_2(c_1, c_2, c, m)$. Most of our results in section 5 will be for the limiting complete information case where $\eta = 0$, though many of our results continue to hold for the incomplete information case where $\eta > 0$.

We now present the Bellman equations for the firms in the case where both firms move simultaneously in every period. We can simplify notation by dropping the m argument in the value

¹²The ϕ function is also sometimes called the “smoothed max” function since we have $\lim_{\eta \rightarrow 0} \phi(a, b) = \max[a, b]$.

functions and investment probabilities and write the Bellman equations for the simultaneous move version of the problem as follows

$$\begin{aligned}
v_{N,1}(c_1, c_2, c) &= r_1(c_1, c_2) + \beta \int_0^c [P_2(c_1, c_2, c) \phi(v_{N,1}(c_1, c, c'), v_{I,1}(c_1, c, c')) + \\
&\quad (1 - P_2(c_1, c_2, c)) \phi(v_{N,1}(c_1, c_2, c'), v_{I,1}(c_1, c_2, c'))] \pi(dc'|c). \\
v_{I,1}(c_1, c_2, c) &= r_1(c_1, c_2) - K(c) + \beta \int_0^c [P_2(c_1, c_2, c) \phi(v_{N,1}(c, c, c'), v_{I,1}(c, c, c')) + \\
&\quad (1 - P_2(c_1, c_2, c)) \phi(v_{N,1}(c, c_2, c'), v_{I,1}(c, c_2, c'))] \pi(dc'|c). \tag{13}
\end{aligned}$$

$$\begin{aligned}
v_{N,2}(c_1, c_2, c) &= r_2(c_1, c_2) + \beta \int_0^c [P_1(c_1, c_2, c) \phi(v_{N,2}(c, c_2, c'), v_{I,2}(c, c_2, c')) + \\
&\quad (1 - P_1(c_1, c_2, c)) \phi(v_{N,2}(c_1, c_2, c'), v_{I,2}(c_1, c_2, c'))] \pi(dc'|c). \\
v_{I,2}(c_1, c_2, c) &= r_2(c_1, c_2) - K(c) + \beta \int_0^c [P_1(c_1, c_2, c) \phi(v_{N,2}(c, c, c'), v_{I,2}(c, c, c')) + \\
&\quad (1 - P_1(c_1, c_2, c)) \phi(v_{N,2}(c_1, c, c'), v_{I,2}(c_1, c, c'))] \pi(dc'|c). \tag{14}
\end{aligned}$$

In the alternating move case, the Bellman equations for the two firms lead to a system of eight functional equations for $(v_{N,j}(c_1, c_2, c, m), v_{I,j}(c_1, c_2, c, m))$ for $j, m \in \{1, 2\}$. Below we write out the four Bellman equations for firm 1, but we omit the value functions for firm 2 to save space since they are defined similarly.

$$\begin{aligned}
v_{N,1}(c_1, c_2, c, 1) &= r_1(c_1, c_2) + \beta f(1|1) \int_0^c \phi(v_{N,1}(c_1, c_2, c', 1), v_{I,1}(c_1, c_2, c', 1)) \pi(dc'|c) + \\
&\quad \beta f(2|1) \int_0^c \rho(c_1, c_2, c') \pi(dc'|c) \\
v_{I,1}(c_1, c_2, c, 1) &= r_1(c_1, c_2) - K(c) + \beta f(1|1) \int_0^c \phi(v_{N,1}(c, c_2, c', 1), v_{I,1}(c, c_2, c', 1)) \pi(dc'|c) + \\
&\quad \beta f(2|1) \int_0^c \rho(c, c_2, c') \pi(dc'|c) \\
v_{N,1}(c_1, c_2, c, 2) &= r_1(c_1, c_2) + \beta f(1|2) \int_0^c \phi(v_{N,1}(c_1, c_2, c', 1), v_{I,1}(c_1, c_2, c', 1)) \pi(dc'|c) + \\
&\quad \beta f(2|2) \int_0^c \rho(c_1, c_2, c') \pi(dc'|c) \\
v_{I,1}(c_1, c_2, c, 2) &= r_1(c_1, c_2) + \beta f(1|2) \int_0^c \phi(v_{N,1}(c_1, c, c', 1), v_{I,1}(c_1, c, c', 1)) \pi(dc'|c) + \\
&\quad \beta f(2|2) \int_0^c \rho(c_1, c, c') \pi(dc'|c). \tag{15}
\end{aligned}$$

where

$$\rho(c_1, c_2, c) = P_2(c_1, c_2, c, 2) v_{I,1}(c_1, c_2, c, 2) + [1 - P_2(c_1, c_2, c, 2)] v_{N,1}(c_1, c_2, c, 2). \tag{16}$$

Note that $P_2(c_1, c_2, c, 1) = 0$, since firm 2 is not allowed to invest when it is firm 1's turn to invest, $m = 1$. A similar restriction holds for $P_1(c_1, c_2, c, c, 2)$.

4 Socially optimal production and investment

We will compare the investment outcomes under duopoly to those that would emerge under a social planning solution that maximizes total expected discounted consumer and producer surplus. The static model of Bertrand price competition is also efficient in a static sense: the firm with lower marginal cost produces the good. However the static model begs the question of potential redundancy in production costs when there are two firms. The static model treats the investment costs necessary to produce the production plant of the two firms as a sunk cost that is ignored in the social planning calculation. In a dynamic model, the social planner does account for these investment costs. Clearly, under our assumptions about production technology (any plant has unlimited production capacity at a constant marginal cost of production) it only makes sense for the social planner to operate a single plant. Thus, the duopoly equilibrium can be inefficient due to duplicative investments that a social planner would not undertake. However we will show that inefficiency in the duopoly equilibrium manifests itself in other ways as well.

4.1 Optimal investment rule in discrete time

Our model of consumer demand is based on the implicit assumption that consumers have quasi-linear preferences; the surplus they receive from consuming the good at a price of p is some initial level of willingness to pay net of p . The social planning solution entails selling the good at the marginal cost of production, and adopting an efficient investment strategy that minimizes the expected discounted costs of production. Let c_1 be the marginal cost of production of the current production plant, and let c be the marginal cost of production of the current state-of-the-art production process, which we continue to assume evolves as an exogenous first order Markov process with transition probability $\pi(c'|c)$ and its evolution is beyond the purview of the social planner. All the social planner needs to do is to determine an *optimal investment strategy* for the production of the good. Since consumers are in effect risk-neutral with regard to the price of the good (due to the quasi-linearity assumption), there is no benefit to “price stabilization” on the

part of the social planner so it is optimal to simply provide the good to consumers at the current marginal cost of production.

Let $C(c_1, c)$ be the *smallest* present discounted value of costs of investment and production when the plant operated by the social planner has marginal cost c_1 and the state-of-the-art technology has a marginal cost of $c \leq c_1$. The minimization occurs over all feasible investment and production strategies, but subject to the constraint that the planner must produce enough in every period to satisfy the unit mass of consumers in the market. We have

$$C(c_1, c) = \min \left\{ c_1 + \beta \int_0^c C(c_1, c') \pi(dc'|c), c_1 + K(c) + \beta \int_0^c C(c, c') \pi(dc'|c) \right\}, \quad (17)$$

where the first component corresponds to the case when investment is not made, and cost c_1 is carried in the future, and the second component corresponds to the case when new state of the art cost c is acquired for additional expense of $K(c)$. The optimal investment strategy takes the form of a *cutoff rule* where it is optimal to invest in the state-of-the-art technology if the current cost c_1 is above a cutoff threshold $\bar{c}_1(c)$. Otherwise it is optimal to produce the good using the existing plant with marginal cost c_1 . The cutoff rule $\bar{c}_1(c)$ is the indifference point in (17), and thus it is the solution to the equation

$$K(c) = \beta \int_0^c [C(\bar{c}_1(c), c') - C(c, c')] \pi(dc'|c), \quad (18)$$

if it exists, and $\bar{c}_1(c) = c_0$ otherwise.¹³

At the optimal cutoff $\bar{c}_1(c)$ the social planner is indifferent between continuing to produce using its current plant with marginal cost c_1 or investing in the state-of-the-art plant with marginal cost of production c . When c_1 is below the threshold, the drop in expected future operating costs is insufficiently large to justify undertaking the investment, and when c_1 is above the threshold, there is a strictly positive net benefit from investing.

Theorem 1 (Socially optimal solution). *The socially optimal (cost minimizing) investment rule $\mathfrak{u}(c_1, c)$ is given by*

$$\mathfrak{u}(c_1, c) = \begin{cases} 1 \text{ (invest)} & \text{if } c_1 > \bar{c}_1(c) \\ 0 \text{ (don't invest)} & \text{if } c_1 \leq \bar{c}_1(c) \end{cases} \quad (19)$$

¹³In problems where the support of $\{c_t\}$ is a finite set, the cutoff $\bar{c}_1(c)$ is defined as the smallest value of c_1 in the support of $\{c_t\}$ such that $K(c) > \beta \int_0^c [C(c_1, c') - C(c, c')] \pi(dc'|c)$.

where $\bar{c}_1(c)$ is the solution to equation (18) and the value function C in equation (18) is the unique solution to the Bellman equation (17).

Proof. The proof follows immediately from the preceding discussion. $\square$

Theorem 2 (Social optimality of monopoly solution). *The socially optimal investment rule $\mathfrak{v}(c_1, c)$ is identical to the profit maximizing investment decision rule of a monopolist who faces the same discount factor β and the same technological process $\{c_t\}$ with transition probability π as the social planner, assuming that in every period the monopolist can charge a price of c_0 equal to the initial value of the state-of-the-art production technology.*

Proof. Since the monopolist is constrained to charge a price no higher than c_0 every period, it follows that the monopolist maximizes expected discounted value of profits by charging a price of $p = c_0$ every period and adopting a cost-minimizing production and investment strategy. $\square$

Note that since the state space for the social planning problem is a triangle $S = \{(c_1, c) | c_1 \in [0, c_0], c \in [0, c_1]\}$ (or finite subset of the triangle S if the support of $\{c_t\}$ is a finite subset of $[0, c_0]$), it follows that investment will never occur along the 45 degree line ((c, c) points in S). For problems with continuous transition densities π where the support of $\{c_t\}$ equals $[0, c_0]$, there will generally be a non-empty *inaction region* below the 45 degree line and above the optimal investment threshold $\bar{c}_1(c)$ where it will not be optimal to invest in the state of the art technology. The following lemma formalizes this statement.

Lemma 1. *If $K(c) > 0$ for all $c \geq 0$, then $\mathfrak{v}(c, c) = 0$ and $\bar{c}_1(c) > c$.*

Proof. The proof follows from (18) and (19). $\square$

Figure 1 provides an illustrative calculation of the optimal investment threshold and optimal investment decision rule in a case where the state of the art production technology improves in a deterministic fashion. The optimal investment threshold $c_1(c)$ is the boundary between the red “continuation region” and the blue “stopping region” in Figure 1, treated as a function of c , though the figure is rotated so that c is the vertical or y axis.

We have implicitly assumed that the cost of investment $K(c)$ is not so high that the social planner would never want to invest in a new technology. Theorem 3 below provides a bound on




Figure 1: Socially Optimal Investment Policy

Notes: State of the art production technology improves in a deterministic fashion with the state of the art marginal costs of production decreasing linearly from $c_0 = 5$ to 0 in equally spaced steps of 0.001. Fixed investment cost of $K(c) = 10$, continuous time interest rate of $r = 0.05$, time interval between successive periods of $\Delta t = 0.016$, so that $\beta = \exp\{-r\Delta t\} = 0.9992$. Technological progress results in the state of the art marginal cost of production reaching its absorbing state of 0 after 5000 discrete time steps, which corresponds to a duration of $T = 80$ in continuous time. The optimal decision rule $\mathfrak{t}(c_1, c)$ is for investment to occur whenever the state (c_1, c) is in the blue region and no investment occurs when the state is in the red region.

the costs of investments that must be satisfied for investment to occur under the socially optimum solution. The proof of Theorem 3, and all subsequent proofs (other than short proofs that can be given in a few lines) are provided in Appendix A.

Theorem 3 (Necessary and sufficient condition for investment by the social planner). *It is optimal for the social planner to invest at some state $(c'_1, c') \in S$ for $c'_1 \leq c_1$ and $c' \leq c'_1$ if and only if there exists $c \in [0, c_1] \cap \text{supp}(\{c_t\})$ (where $\text{supp}(\{c_t\})$ is the support of the Markov process for the state of the state of the art technology $\{c_t\}$) for which*

$$\frac{\beta(c_1 - c)}{(1 - \beta)} > K(c). \quad (20)$$

The conditions under which the social planner will invest in some future state plays a central role when we analyze the duopoly investment dynamics in section 5. To make this condition applicable to the duopoly case we need to consider a social planner that controls two production plants with constant returns to scale production technologies and marginal costs of production of $c_1 > 0$ and $c_2 > 0$. Assuming that the state of the art production technology evolves according to

the same Markov process π and investment cost $K(c)$ as in our analysis of a social planner that operates only a single production facility, we say that investment cost $K(c)$ are *not prohibitively high* at the state (c_1, c_2, c) where $c \leq \min[c_1, c_2]$ if there exists a $c' \in [0, c]$ satisfying

$$\frac{\beta(\min[c_1, c_2] - c')}{1 - \beta} > K(c') \quad (21)$$

It is clear from Theorem 3 that investment cost is not prohibitively high if and only if the investment is socially optimal, so in the rest of the paper we use both terms interchangeably.

Clearly, a social planner who has two production facilities that have constant returns to scale and no capacity constraints will only want to operate the plant that has the lower marginal cost of production. If investment costs are not prohibitively high, the social planner will find it optimal to close the high cost plant and replace the more efficient of these two plants by an even more efficient state of the art plant at some point in the future.

4.2 Social surplus and measure for efficiency

We conclude this section by defining social surplus in the duopoly case, and defining our measure of the efficiency of the duopoly equilibrium. From Theorem 2 above, the socially optimal investment strategy coincides with the optimal investment strategy of a monopolist producer. Let $S(c_1, c)$ be the expected discounted surplus for the social planner under an optimal investment and production plan when the state of the economy is (c_1, c) . We have

$$S(c_1, c) = \frac{c_0}{(1 - \beta)} - C(c_1, c) \quad (22)$$

where C is the solution to the Bellman equation for costs (17). This social surplus is the difference between the present value of the maximal willingness of consumers to pay for the good, $c_0/(1 - \beta)$, less the minimized cost of producing the good and installing upgraded production plants as technology improves under the planner's optimal investment strategy. This measure of surplus gives 100% of the surplus to consumers and zero profits to producers, since the social planner can be interpreted as selling the good to consumers at a lump sum price of $C(c_1, c)$, which equals the expected discounted cost of producing the good and investing in the plant and plant upgrades necessary to produce the good as cost-effectively as possible. However the surplus also equals the expected discounted profit of a monopolist who is constrained to charge a price no higher than c_0

every period (consumers' maximal willingness to pay). Under this latter interpretation of surplus, 100% of the surplus goes to the monopolist and none to consumers.

In the duopoly equilibrium, total surplus is the sum of consumer surplus and firm profits. Let $R_d(c_1, c_2, c)$ be the expected discounted revenue of the duopolists and $C_d(c_1, c_2, c)$ be the total expected discounted costs of production and investment by the duopolists. Thus, $\Pi_d(c_1, c_2, c) = R_d(c_1, c_2, c) - C_d(c_1, c_2, c)$ is the total producer surplus, and total surplus in the duopoly equilibrium, $S_d(c_1, c_2, c)$, is given by the sum of consumer and producer surplus

$$\begin{aligned} S_d(c_1, c_2, c) &= \frac{c_0}{(1-\beta)} - R_d(c_1, c_2, c) + \Pi_d(c_1, c_2, c) \\ &= \frac{c_0}{(1-\beta)} - R_d(c_1, c_2, c) + R_d(c_1, c_2, c) - C_d(c_1, c_2, c) \\ &= \frac{c_0}{(1-\beta)} - C_d(c_1, c_2, c). \end{aligned} \tag{23}$$

Clearly a social planner would never have any reason to operate a second, higher cost, production facility given that both production plants have constant returns to scale production functions. So it is evident that there will be inefficiency in the duopoly equilibrium due to redundancy reasons alone. However we will show in the next section that there are other sources of inefficiency in the duopoly equilibrium. Since the social planner will always shut down the higher cost plant, we make the correspondence that the surplus in the duopoly equilibrium in state (c_1, c_2, c) should be compared with the surplus in the social planning state $(\min(c_1, c_2), c)$. Therefore we define the efficiency index $\xi(c_1, c_2, c)$ as the ratio of social surplus under the duopoly equilibrium and the corresponding social planning equilibrium solutions

$$\begin{aligned} \xi(c_1, c_2, c) &= \frac{S_d(c_1, c_2, c)}{S(\min(c_1, c_2), c)} \\ &= \frac{c_0/(1-\beta) - C_d(c_1, c_2, c)}{c_0/(1-\beta) - C(\min(c_1, c_2), c)}. \end{aligned} \tag{24}$$

In the next section we will show how $C_d(c_1, c_2, c)$ can be calculated for any equilibrium of the duopoly game, and that $C_d(c_1, c_2, c) \geq C(\min(c_1, c_2), c)$ for each state (c_1, c_2, c) . This implies that $\xi(c_1, c_2, c) \in [0, 1]$ for all $(c_1, c_2, c) \in S$. While it will generally be the case that $\xi(c_1, c_2, c) < 1$ for most points $(c_1, c_2, c) \in S$, we will characterize conditions under which there also exist efficient duopoly equilibria, i.e. equilibria for which $\xi(c_1, c_2, c) = 1$ for all (c_1, c_2, c) along the equilibrium path.

5 Duopoly Investment Dynamics

We are now in position to solve the model of duopoly investment and pricing described in section 3 and characterize the stationary Markov Perfect equilibria of this model. Under our assumptions the Markov process governing exogenous improvements in production technology has an absorbing state, which without loss of generality we assume to correspond to a marginal production cost $c = 0$. If either of the two firms reach this absorbing state, there is no need for any further cost-reducing investments. If investments do occur, they would only be motivated by transitory shocks (e.g. one time investment tax credits, or subsidies, etc.) but there is no longer any strategic motivation for undertaking further investment. We refer to this zero cost absorbing state as the *end game* of the overall investment game.

The main complication of solving dynamic games compared to static or one shot games is that in the former, the entries of the payoff matrix are generally not specified *a priori* but rather depend on the solution to the game, including the choice of the equilibrium of the game. Thus, we start our analysis with the easiest cases first, showing how we derive the payoffs implied by the Bellman equations (13) and (14) given in section 3, to simultaneously determine equilibrium payoffs and decision rules in the end game.

With the end game solution in hand, we then discuss the solution of the full game where $c > 0$. To solve the full game, we employ a *state space recursion algorithm*, that involves a form of backward induction in the state space of the game, starting in the end-game (absorbing state) $c = 0$ and working backward — but not in time, but in terms of the start of the art production cost c . Iskhakov, Rust and Schjerning (2013) discuss the state recursion algorithm in more detail and show how it can be applied to compute MPE of both the simultaneous move and the alternating move version of the game given a particular equilibrium selection rule. Further, they introduce a new algorithm called the *recursive lexicographical search* (RLS) algorithm that is capable of finding *all* MPE of both the simultaneous and alternating move versions of the game. We used these algorithms to compute illustrative examples below, but our main results are based on proofs of the general properties of the equilibria of this game.

Table 1: End Game Payoff Matrix

		Firm 2	
		Invest	Don't invest
Firm 1	Invest	$-K(0), c_1 - c_2 - K(0)$	$\frac{\beta c_2}{1-\beta} - K(0), c_1 - c_2$
	Don't invest	$0, c_1 - c_2 + \frac{\beta c_1}{1-\beta} - K(0)$	$0, (c_1 - c_2)/(1 - \beta)$

Notes: The payoffs are shown for state $(c_1, c_2, 0)$ where $c_1 \geq c_2$, $\eta = 0$. Only special case when firms' investment probabilities are either 0 or 1 is displayed. In the general case the payoff of firm 2 in no investment case (lower right cell) is $c_1 - c_2 + \beta V_2(c_1, c_2, 0)$, see the proof of Lemma 3 in Appendix A of the unabridged version of the paper for more detail.

5.1 End Game Solutions

Consider first a simultaneous move game, denoted by $m = 0$ in the notation of section 3. The end game at the state $(c_1, c_2, 0)$ corresponds to a “stage game” of the overall investment game where the state of the art cost c has reached the zero cost absorbing state but the marginal costs of the two firms may not necessarily equal 0 since the firms may not yet have invested the required amount $K(0)$ to acquire the state of the art technology that would enable them to produce at zero marginal cost. Table 1 presents the payoff matrix for this game for the case where $\eta = 0$ and $c_1 \geq c_2$.

Note that entries of the payoff matrix in Table 1 are only valid in the special case where the firms' investment probabilities are either 0 or 1, and in such cases there will be no possibility of future investment after the current period t . However if we use the Bellman equation (14) to solve for the payoffs to the firms in the case where the firms' investment probabilities are strictly in the interior of the $[0, 1]$ interval, as is the case in a mixed strategy equilibrium, then the expected discounted payoff to firm 2 in the lower right hand corner of Table 1 is not $(c_1 - c_2)/(1 - \beta)$ but rather $(c_1 - c_2) \cdot (1 + \beta P_1(c_1, c_2, 0))/(1 - \beta + \beta P_1(c_1, c_2, 0))$ which reflects firm 2's expectation that its future profit stream will end at some point in the future when a realization of firm 1's mixed strategy leads it to invest and leapfrog firm 2. Thus unlike static games where the payoff matrix is defined in advance, the payoffs to the firms at each stage game in the overall dynamic game are *endogenously determined* and depend on the equilibrium that is selected in the current and future periods. This makes the problem of finding equilibria of dynamic games a much more complicated undertaking than finding the equilibria of static games.

The duopoly investment game is isomorphic to an *anti-coordination game*. The firms avoid

investing at the same time, since doing so results in *ex post* Bertrand price competition that drives both of their profits to zero. This is the Bertrand investment paradox described in the introduction. The anti-coordination end game generally has three equilibria. The two pure strategy equilibria correspond to outcomes where firm 1 invests and firm 2 doesn't and firm 2 invests and firm 1 doesn't. There is also a third mixed strategy equilibrium where firm 1 invests with probability $P_1(c_1, c_2, 0) \in (0, 1)$ and firm 2 invests with probability $P_2(c_1, c_2, 0) \in (0, 1)$. It is not hard to see that when $c_1 = c_2 = c$ the game is fully symmetric and we have $P_1(c, c, 0) = P_2(c, c, 0)$. However when $c_1 \neq c_2$, then the game is asymmetric and $P_1(c_1, c_2, 0) \neq P_2(c_1, c_2, 0)$. Lemma 3 below shows that $c_1 > c_2$ implies $P_1(c_1, c_2, 0) > P_2(c_1, c_2, 0)$, i.e. *the high-cost follower has a greater probability of investing and leapfrogging the low-cost leader*.

From both the standpoint of the firms and the social planner, the mixed strategy equilibrium is the “bad” equilibrium. In the symmetric case, $c_1 = c_2$, the mixed strategy equilibrium results in zero expected profits for both firms, whereas each of the pure strategy equilibria result in positive profits for the investing firm. In the asymmetric case, the low cost leader reaps a positive profit until one or the other of the firms invests in the state-of-the-art production technology, and earns zero profits thereafter. Intuitively there is a possibility of redundant investment under the mixed strategy equilibrium, which causes it to be inefficient from a social planning perspective.

Let $p_1 \in [0, 1]$ represent firm 2's beliefs about the probability that firm 1 will invest. Adapting the notation from section 3, the *best response probability* for firm 2 is a function of the form

$$p_2 = P_2(p_1), \quad (25)$$

where we have omitted the current state $(c_1, c_2, 0)$ as an additional argument of P_2 to simplify notation. This precise functional form for P_2 is given in equation (11) in section 3, where the values for firm 2 from investing and not investing, respectively, $v_{I,2}(c_1, c_2, 0, p_1)$ and $v_{N,2}(c_1, c_2, 0, p_1)$ are expressed with the additional argument p_1 representing firm 2's beliefs about the probability firm 1 will invest in state $(c_1, c_2, 0)$. The best response probability for firm 1 is defined similarly

$$p_1 = P_1(p_2), \quad (26)$$

Substituting equation (25) into equation (26), we obtain the *second order best response probability*

$$p_1 = P_2(P_1(p_1)). \quad (27)$$




Figure 2: End game equilibria

Notes: The second order best response function of firm 1 is drawn in the point $c_0 = 5$, $c_1 = 1.25$, $c_2 = 3.75$, $c = 0$, investment cost is $K(0) = 10$, $\beta = 0.9513$. The discontinuity is at $p = 0.59$, and the corresponding investment probability of firm 2 in the mixed strategy equilibrium is 0.86.

The fixed points of this function constitute the Nash equilibria of the $(c_1, c_2, 0)$ end game (or Bayesian Nash equilibria in the case $\eta > 0$ where firms have privately observed investment cost shocks). Figure 2 plots the equilibria computed by plotting the second order best response function against the 45 degree line in the case $\eta = 0$. The best response probabilities P_1 and P_2 become discontinuous functions of p_2 and p_1 , respectively, as the variance of the unobserved investment cost shocks $\eta \downarrow 0$. The discontinuities occur at values of these probabilities that constitute mixed strategies of the investment stage game. We can directly compute the set of all equilibria, and in both cases (i.e. $\eta = 0$ and $\eta > 0$) topological index theorems guarantee that there will be an odd number of equilibria for almost all parameter values defining the game (e.g. the discount factor β , the variance parameter η , firm costs, transition probabilities π , etc). Further, a homotopy argument can be used to show that the equilibrium correspondence is continuous so the set of equilibria to the incomplete information investment game when $\eta > 0$ converge to the set of equilibria of the complete information end game where $\eta = 0$ as $\eta \downarrow 0$.

Lemma 2 (Efficiency of equilibria in the simultaneous move end game). *Suppose $\eta = 0$ and $m = 0$ (i.e. simultaneous move game with no investment cost shocks), and $c = 0$. In states where investment is not socially optimal, i.e. $\beta \min(c_1, c_2)/(1 - \beta) < K(0)$, the investment game has a unique pure strategy equilibrium where neither firm invests. When investment is socially optimal, the in-*

vestment game has three subgame perfect Nash equilibria: two efficient pure strategy equilibria and an inefficient mixed strategy equilibrium.

Lemma 3 (Leapfrogging in the mixed strategy equilibrium). *Suppose $\eta = 0$ and $m = 0$ (i.e. simultaneous move game with no investment cost shocks), $c = 0$, and the investment is socially optimal, i.e. $\beta \min(c_1, c_2)/(1 - \beta) > K(0)$. Then if $c_1 > c_2 > 0$, in the mixed strategy equilibrium the probability that firm 1 invests exceeds the probability that firm 2 invests, $P_1(c_1, c_2, 0) > P_2(c_1, c_2, 0)$.*

Lemma 4 (Efficiency of the alternating move end game). *Suppose $\eta = 0$ and $m \neq 0$ (i.e. alternating move game with no investment cost shocks), and $c = 0$. In every end game state $(c_1, c_2, 0)$ there is a unique efficient equilibrium, i.e. both firms invest when it is their turn to invest if and only if investment would be optimal from the point of view of the social planner.*

In the interest of space, the proofs of the lemmas above and all subsequent results (except those that are short and intuitive and are provided in the text) appear in the appendix of an unabridged version of this paper that is available from the authors on request. The end game already provides us with some insight into the nature of inefficiency of the duopoly equilibrium solution. From Lemma 2 we see that in the simultaneous move end game, inefficiency only results from the lack of coordination of the two firms in the *mixed strategy equilibrium*. However the two pure strategy equilibria represent efficient solutions to the anti-coordination game where there is in effect “agreement” that only one of the firms will invest. The efficiency of equilibrium in the alternating move game is less surprising, since the structure of alternating moves provides an *exogenous coordination mechanism* that prevents the simultaneous investment by the firms.

Lemma 3 provides a first indication of the possibility of leapfrogging since the high cost leader has a higher probability of investing to become the (permanent) low cost leader who “wins the game” by acquiring the state-of-the-art plant with zero marginal costs of production.

We conclude this section by pointing out that continuous state versions of the full dynamic investment game has a continuum of equilibria.

Theorem 4 (Number of equilibria in simultaneous move game). *If investment is socially optimal, and the support of the Markov process $\{c_t\}$ for the state of the art marginal costs is the full interval $[0, c_0]$ (i.e. continuous state version), the simultaneous move Bertrand investment and pricing game has a continuum of MPE.*

In the subsequent analysis we will focus on a subclass of games where the support of the Markov process $\{c_t\}$ representing the evolution of the state of the art production technology is a *finite subset of R^1* . This will imply that the state space of the investment game is a finite subset of S where all of the coordinates c_1 , c_2 and c lie in the support of the Markov process $\{c_t\}$. If we further restrict the set of possible equilibrium selection rules to be *deterministic* functions of the current state (c_1, c_2, c) , then we can show that there will only be a finite number of possible equilibria in both the simultaneous and alternating move formulations of the game, though Iskhakov, Rust and Schjerning (2013) show that the number of possible equilibria grows exponentially fast in the total number of points in the state space of the game $|S|$.

5.2 Investment dynamics in the full game

With the end game solutions in hand, we are now ready to discuss the solution of the full game. The end game equilibria give us some insight into what can happen in the full game, but the possibilities in the full game are much richer, since unlike in the end game, if one firm leapfrogs its opponent, the game does not end, but rather the firms must anticipate additional leapfrogging and cost reducing investments in the future. In particular, forms of dynamic coordination may be possible that are not present in the end game, which is closer to a two stage game than to a fully dynamic infinite horizon game.

As we indicated earlier, the full game can be solved using a *state recursion algorithm* that differs from a standard backward induction or *time recursion* by our use of the state of the art technology c as the “time variable” for solving the value functions and equilibria. Consider the simultaneous move game at generic point (c_1, c_2, c) in the state space S . The two firms still face a simultaneous move game about whether to invest or not, but unlike the end game states $(c_1, c_2, 0)$ the firms must also anticipate that the state of the art may improve to some value $c' < c$ in the next period. Thus a wider variety of transitions in the states are possible. If there is no technological innovation and tomorrow’s value of the state of the art c' equals today’s value c , investment by one or both of the firms can change the state to either (c, c_2, c) (if only firm 1 invests), (c_1, c, c) (if only firm 2 invests) or (c, c, c) if both firms 1 and 2 invest. We refer to any of these three states to be *edge states* whereas states of the form (c_1, c_2, c) where $c_1 > c$ and $c_2 > c$ are *interior states*.

Lemma 5 (No investment equilibrium at edge states). *In both the simultaneous and alternating move games with no investment cost shocks (i.e. $\eta = 0$) there is a unique stage equilibrium at all edge states in which neither firm invests.*

When investment cost is not prohibitively high at the interior states (c_1, c_2, c) where $c_1 > c$ and $c_2 > c$ there can be multiple stage equilibria similar to the end game (see Lemmas 2 and 4). However the calculation of the payoff matrices for these states is more complicated because the value functions $v_{I,j}(c_1, c_2, c)$ and $v_{N,j}(c_1, c_2, c)$, $j \in \{1, 2\}$ corresponding to the choices of investing (13) and not investing (14) depend on the firms' expectations of both future technological improvements and whether their opponent will invest. The solutions to the Bellman equations (13) and (14) also depend on the equilibrium selection rule since the choice of equilibrium at lower levels of the game¹⁴ affects the firms' expected payoffs from choosing whether to invest or not at the current state of the game (c_1, c_2, c) . Similar reasoning applies to the Bellman equations in the alternating move game, since we will show below that there are typically multiple equilibria in these games as well. When there are multiple possible equilibria at the interior states of the game, we are free to choose any one of these equilibria, and the latitude to choose different equilibria at the interior nodes generates a vast multiplicity of MPE equilibria in the overall game.

Note that while we have great freedom to select various equilibria at the interior states of the game, we do not have unlimited freedom. This is because various choices of equilibria at the end game and lower levels of the game affect the set of possible equilibria at higher levels of the game. Some equilibrium selections can result in a unique equilibrium at certain higher level interior states, whereas in other interior states there may be 3 or 5 possible equilibria depending on the parameters of the game and the choices of equilibria at lower level nodes. Recursive lexicographic search algorithm (RLS) developed in (Iskhakov, Rust and Schjerning, 2013) can systematically enumerate and compute *all possible equilibria* of a class of *directional dynamic games* that includes the dynamic duopoly investment game as a special case. We have used this algorithm to systematically compute and characterize the set of all possible equilibria in our analysis below, though our main results are based on analytical proofs and we are not simply reporting results from specific computed examples.

¹⁴At states (c_1, c_2, c') for $c' < c$.

However there is a subclass of games for which we can prove that the equilibrium exists and is unique¹⁵.

Theorem 5 (Sufficient conditions for uniqueness). *In the dynamic Bertrand investment and pricing game a sufficient condition for the MPE to be unique is that (i) firms move in alternating fashion (i.e. $m \neq 0$), and (ii) for each c in the support of π we have $\pi(c|c) = 0$.*

Theorem 5 implies that under strictly monotonic technological improvement the alternating move investment game has a unique Markov perfect equilibrium. This is closely related to, but not identical with an assumption of the *deterministic technological progress* as discussed in section 3. There are specific types of non-deterministic technological progress for which Theorem 5 will still hold, resulting in a unique equilibrium to the alternating move game. In subsection 5.5 we will return to this case, by considering further properties of the unique equilibrium that results when $\pi(c|c) = 0$ including the conditions in which it constitutes a discrete time equivalent to Riordan and Salant's (1994) continuous time preemption equilibrium.

It is also helpful to understand why multiple equilibria can arise in the alternating move game when $\pi(c|c) > 0$. The reason is that when there is a positive probability of remaining in any given given state (assuming firms choose not to invest when it is their turn to invest), it follows that each firm's value of *not investing* depends on their belief about the probability their opponent will invest. Thus, by examining the Bellman equations (15) it not hard to see that for firm 1 the value of not investing when it is its turn to invest, $v_{N,1}(c_1, c_2, c, 1)$, depends on $P_2(c_1, c_2, c, 2)$ when $\pi(c|c) > 0$. This implies that $P_1(c_1, c_2, c, 1)$ will depend on $P_2(c_1, c_2, c, 2)$, and similarly, $P_2(c_1, c_2, c, 2)$ will depend on $P_1(c_1, c_2, c, 1)$. This mutual dependency creates the possibility for multiple solutions to the Bellman equations and the firms' investment probabilities and multiple equilibria at various stage games of the alternating move game. We will have more to say about the nature of these equilibria in subsection 5.3, where we show that when $\pi(c|c) > 0$, the stage games of the alternating move investment game are isomorphic to *coordination games*, unlike the stage games in the simultaneous move investment game are isomorphic to *anti-coordination games* as we have already shown in section 5.1.

¹⁵Modulo a relabeling of the firms, so by "unique" we actually mean a situation where there are at two *asymmetric* equilibria of the investment game.

Lemmas 2 and 4 enable us to provide a resolution to the Bertrand investment paradox in the end game. These Lemmas state that as long as investment costs are not prohibitively high (so that a social planner would find it optimal to undertake investment) at least one firm will be willing to invest under either the simultaneous or alternating move formulations of the game. However the result that no investment by both firms cannot be a duopoly equilibrium outcome extends to the full game provided investment costs are not prohibitively high.

Theorem 6 (Solution to Bertrand investment paradox). *If investment is socially optimal at a state point $(c_1, c_2, c) \in S$, then no investment by both firms cannot be an MPE outcome in the subgame starting from (c_1, c_2, c) in either the simultaneous or alternating move versions of the dynamic game. In particular, if investment is socially optimal at the apex (c_0, c_0, c_0) , no investment by both firms cannot be an MPE outcome in the full Bertrand investment and price game.*

5.3 Characterization of the set of equilibrium payoffs

Despite the prevalence of leapfrogging in equilibrium, we now show that “monopoly” equilibria will also exist in the simultaneous move game when there are no investment cost shocks, $\eta = 0$ provided the cost of investment is not prohibitively high. However it is generally not possible to support the monopoly outcome in the alternating move version of the game except for isolated, atypical counterexamples. Note that the monopoly equilibrium we characterize below is *not* the preemption equilibrium analyzed by Riordan and Salant (1994). In contrast to their rent dissipation result, under the monopoly equilibrium below profits are positive and are equal to the maximum possible profits subject to the limit on price, and efficiency is 100%.

Theorem 7 (Necessary and sufficient condition for monopoly equilibrium in the simultaneous move game). *If investment is socially optimal at the apex (c_0, c_0, c_0) , the simultaneous move game has two fully efficient “monopoly” MPE equilibria in which either one or the other firm make all the investments and earns maximum feasible profit.*

We note that the proof of Theorem 7 does not rely on any assumptions about the value functions of the firms or their investment decisions at points (c_1, c_2, c) where $c_2 < c_1$ that are off the

equilibrium path.¹⁶ For this reason, there are potentially multiple different monopoly equilibrium outcomes, with each of these equilibria having the same (monopoly) equilibrium path, but with different stage equilibria “chosen” on the subset of points in S that are off the monopoly equilibrium path. Indeed, our computations reveal that there are generally multiple equilibria that have the same monopoly payoffs. Because they differ only off the equilibrium path, these equilibria are *observationally equivalent*.

Theorem 8 (Existence of a zero profit equilibrium in the simultaneous move game). *Under the same conditions as in Theorem 7, there exist a symmetric equilibrium in the simultaneous move game that results in zero expected payoffs to both firms at all states $(c, c, c') \in S$ with $c' \in [0, c]$, and zero expected payoffs to the high cost firm and positive expected payoffs to the low cost firm in states (c_1, c_2, c) where $c_1 \neq c_2$.*

We proved Theorems 7 and 8 by mathematical induction, and this is the reason we assumed that the support of $\{c_t\}$ is a finite set. However we believe both Theorems still hold when the state space is continuous. We believe we can prove this result by a limiting argument, by considering a sequence of finite state investment games defined on a sequence of finite subsets of S that become dense in the overall continuous state pyramid $S \subseteq R^3$. However in the interest of space we do not attempt to prove this result here and merely state it as a conjecture that we believe to be true.


Conjecture 1 (Extension to continuous state games). *Theorems 7 and 8 also hold when technology evolves according to a continuous state Markov process with full support on the $[0, c_0]$ interval.*

Theorems 7 and 8 are sufficient for us to prove the main results of this subsection, which we have formulated in the following theorem.

Theorem 9 (Triangular payoffs in the simultaneous move game). *Suppose that the $\{c_t\}$ process has finite support and $\eta = 0$. The convex hull of the set of the expected discounted equilibrium payoffs to the two firms in all MPE equilibria of the game at the apex state $(c_0, c_0, c_0) \in S$ in the simultaneous move game is a triangle with vertices at the points $(0, 0)$, $(0, V_M)$ and $(V_M, 0)$ where*

¹⁶Technically speaking any point $(c_1, c_2, c) \in S$ where $c_2 \neq c_0$, and c_0 is the initial starting point of the game is off the equilibrium path for the monopoly equilibrium where firm 1 is the monopolist and firm 2 never invests. However the notion of subgame perfect equilibrium requires us to verify that this is also an equilibrium from any other starting node in the game, including off the equilibrium path points (c_1, c_2, c) where $c_2 < c_1$.

Figure 3: Initial node equilibrium payoffs in the simultaneous move game



Notes: The panels plot payoff maps of the Bertrand investment game with deterministic (a) and random (b) technologies. Parameters are $\eta = 0$, $\beta = 0.9512$, $k_1 = 8.3$, $k_2 = 1$, $n_c = 5$. Parameters of beta distribution for random technology are $a = 1.8$ and $b = 0.4$. Panel (a) displays the initial state payoffs to the two firms in the 192,736,405 equilibria of the game, though there are 63,676 distinct payoff pairs among all of these equilibria. Panel (b) displays the 1,679,461 distinct payoff pairs for the 164,295,079 equilibria that arise under stochastic technology. The color and size of the dots reflect the number of repetitions of a particular payoff combination.

$V_M = v_{N,i}(c_0, c_0, c_0)$ is the expected discounted payoff to firm i in the monopoly equilibrium where firm i is the monopolist investor.

Figure 3 illustrates Theorem 9 by plotting all apex payoffs to the two firms under all possible deterministic equilibrium selection rules in the simultaneous move game where the support of $\{c_t\}$ is the 5 point set $\{0, 1.25, 2.5, 3.75, 5\}$. Panel (a) plots the set of payoffs that occur when technological progress is deterministic, whereas panel (b) shows the much denser set of payoffs that occur when technological progress is stochastic. Though there are actually a greater total number of equilibria (192,736,405) under deterministic technological progress, many of these equilibria are observationally equivalent and we refer to them as *repetitions* of the same payoff point and indicate the number of such repetitions by increasing the size of the payoff point plotted to be proportional to the number of repetitions. Figure 3 shows that when technology is stochastic there are fewer repetitions and so even though there are actually 28 million fewer equilibria, there are actually a substantially greater number (1,679,461 versus 63,676) of distinct payoff points.

It is perhaps not surprising that when firms move in an alternating fashion neither one of them will be able to attain monopoly payoffs in any equilibrium of the alternating move game (except

for some isolated counterexamples we discuss below). When firms make simultaneous investment decisions, the high cost firm has no incentive to deviate from the equilibrium path in which its opponent always invests. However when the firms move in an alternating fashion, the high cost firm will have an incentive to deviate because it knows that its opponent will not be able to invest at the same time (thereby avoiding the Bertrand investment paradox), and once the opponent sees that the firm has invested, it will not have an incentive to immediately invest to leapfrog for a number of periods until it is once again its turn to invest and there has been a sufficient improvement in the state of the art. This creates a temptation for each firm to invest and leapfrog their rival that is not present in the simultaneous move game, and the alternating move structure prevents the firms from undertaking inefficient simultaneous investments, though it also generally prevents either firm from being able to time their investments in a socially optimal way.

Theorem 10 (The set of equilibrium payoffs in the alternating move game). *The (convex hull of the) set of expected discounted equilibrium payoffs to the two firms in all possible equilibria at the apex of the alternating game is a strict subset of the triangle with the same vertices as in Theorem 9.*

The proof of Theorem 10 demonstrates that the zero expected profit mixed strategy equilibrium is no longer sustainable in the alternating move game either. Though it may seem tempting to conclude from Lemma 4 that mixed strategies can never arise in the alternating move game (since Lemma 4 proves that in the *end game* the two firms have only pure, efficient investment strategies i.e. they invest with probability 1 if and only if a social planner would invest), this does not continue to be the case at higher nodes in the game. We find that both pure and mixed strategy stage game equilibria are possible at higher levels nodes of the alternating move game. The intuition as to why this should occur is that even though only one firm invests at any given time, when $\pi(c|c) > 0$ the firms know that there is a positive probability that they will remain in the same state (c_1, c_2, c) for multiple periods until the technology improves and the duration in this state is geometrically distributed. The possibility of remaining in the same state implies that the payoff to each firm from *not investing* depends on their belief about the probability their opponent will invest in this state at its turn.

Further, it is not hard to see that the value of not investing is a *decreasing function* of the

firm's beliefs about the probability that its opponent will invest, however the value to either firm from *investing* does not depend on its beliefs about the probability its opponent will invest in state (c_1, c_2, c) . As a result, there will exist states for which it is optimal for firm 1 not to invest when it is firm 1's turn to invest provided firm 1 believes that firm 2 will also not invest if the state remains at (c_1, c_2, c) when it is firm 2's turn to invest. But because the expected value of not investing decreases in P_2 (firm 2's probability of investing) firm 1 will also want to invest if it believes that firm 2 will invest in state (c_1, c_2, c) when it is firm 2's turn to invest.

In these situations there will be three equilibria of the stage game at (c_1, c_2, c) when $\eta = 0$. But unlike in the case of simultaneous move game, the stage game in the alternating move investment game is a *coordination game* rather than an *anti-coordination game*. That is, there are three equilibria in these stage games: two pure strategy equilibria where *both* firms invest or don't invest, and a third mixed strategy equilibrium where both firms invest with a positive probability. Furthermore, we find that unlike the symmetric, zero profit mixed strategy stage game equilibrium of the simultaneous move game, in the mixed strategy stage game equilibria of the alternating move game the firm with the *lower* cost has a *higher* probability of investing. This is probably due to the fact that the alternating move stage games are coordination rather than anti-coordination games, and is also the reason why a zero profit mixed strategy equilibrium is no longer possible.

Theorem 11 (Alternating move stage games are anti-coordination games when $\pi(c|c) > 0$). *In the alternating move investment game, in states where $\pi(c|c) > 0$ at the higher level nodes $(c_1, c_2, c) \in S$ where $c > 0$ there can be multiple stage game equilibria. When $\eta = 0$ the stage game has the structure of a coordination game with three equilibria: two of the equilibria are pure strategy equilibria where the firms either both invest or both don't invest with probability 1, as well as a third mixed strategy equilibrium.*

Though we consistently observe that inequality (28) holds in all of the numerical solutions of alternating move MPE where $\pi(c|c) > 0$ that we examined, we have not been able to find a general proof of this inequality but we conjecture

Conjecture 2 (Low cost firm has higher probability of investing in the alternating move game). *In any mixed strategy stage game equilibrium of the alternating move game, the low cost leader*

has a higher probability of investing when it is its turn to invest:

$$c_1 > c_2 \implies P_1(c_1, c_2, c, 1) < P_2(c_1, c_2, c, 2). \quad (28)$$

Additionally, in all of our numerical solutions of *simultaneous move* investment game, we found that in the symmetric zero profit mixed strategy equilibrium the high cost firm always has a *higher* probability of investing than the low cost firm. We have been able to prove that this result holds, but only in the symmetric, zero expected profit mixed strategy stage game equilibria, as shown in Theorem 12 below, under a slight strengthening of the condition that cost of investing in new technology are not too high.

Theorem 12 (Leapfrogging in the mixed strategy equilibrium of the simultaneous move game). *Suppose a slightly stronger version of the condition (21) holds in simultaneous move investment game. Namely, if*

$$\frac{\beta(\min[c_1, c_2] - c)}{1 - \beta} > K(c) \quad (29)$$

for some some $(c_1, c_2, c) \in S$ then inequality (29) also holds at all points $(c_1, c_2, c') \in S$ where $c' \leq c$. Then, in the symmetric, zero-payoff mixed strategy equilibrium of Theorem 8 the set of states where the symmetric, zero expected profit mixed strategy stage game equilibria hold are non-decreasing, in the sense that if such an equilibrium holds at the point $(c_1, c_2, c) \in S$, then there is also a zero expected payoff, symmetric mixed strategy equilibrium at all points $(c_1, c_2, c') \in S$ with $c' \leq c$. Further, the same inequality on the mixed strategy as was proved in the end game in Lemma 3 continues to hold, i.e. at any point $(c_1, c_2, c) \in S$ where investment occurs with positive probability under the mixed strategy equilibrium, we have

$$c_1 > c_2 \implies P_1(c_1, c_2, c) > P_2(c_1, c_2, c). \quad (30)$$

5.4 Efficiency of Equilibria

We evaluated the efficiency of duopoly equilibria by calculating the *efficiency index* $\xi(c_1, c_2, c) \in [0, 1]$ defined in equation (24) of section 4. Recall this index is the ratio of total surplus (i.e. the sum of discounted consumer surplus plus total discounted profits) under the duopoly equilibrium to the maximum total surplus achieved under the social planning solution. In order to calculate the social

surplus under the duopoly equilibrium we had to recursively calculate the total expected discounted costs of production, accounting for potential redundancy in investment costs in situations where the two firms invest at the same time along the equilibrium path. The total expected discounted costs for the two firms at state (c_1, c_2, c) is given by a function $C_d(c_1, c_2, c)$ that has a recursive definition similar to the Bellman equations for the firms' value functions in equation (14) (for the simultaneous move game) and (15) (for the alternating move game). We present the full set of functional equations for $C_d(c_1, c_2, c)$ in appendix A2 of the unabridged version of this paper.


We note that the calculation of C_d is *equilibrium specific* and thus its value depends on the particular equilibrium of the overall game that we select. For example, we have already proved that monopoly investment by one of the firms is an equilibrium in the simultaneous move game, provided the cost of investment is not prohibitively high. This implies that $C_d(c_0, c_0, c_0) = C(c_0, c_0)$ where the right hand side is the minimized discounted cost of production and investment from the social planning solution, given by the Bellman equation (17) in section 4. However we now show the non-monopoly equilibria of either the simultaneous or alternating move investment games are generally *inefficient* and this inefficiency is typically due to two sources a) duplicative investments (valid only in mixed strategy equilibria in the simultaneous move investment game), and b) excessively frequent investments.

Note that it is logically possible that inefficiency could arise from *excessively infrequent investments* and the logic of the Bertrand investment paradox might lead us to conjecture that we should see investments that are too *infrequent* in equilibrium relative to what the social planner would do. However surprisingly, we find that duopoly investments are generally excessively frequent to the social optimum, with preemptive investments (when they arise) representing the most extreme form of inefficient excessively frequent investment in new technology.

The two panels in the left column in Figure 4 illustrate the set of equilibrium payoffs from all MPE equilibria computed by the RLS algorithm of Iskhakov, Rust and Schjerning (2013). We also compute efficiency index $\xi(c_0, c_0, c_0)$ for each of the equilibria, and treating there calculated efficiency indices as “data”, plot their empirical distribution in the corresponding panels in the right column in Figure 4.

Panels (a) and (c) in Figure 4 represent an alternating move investment game with deterministic alternations of the right to move and the technological progress which is not strictly monotonic, i.e.

Figure 4: Payoff maps and efficiency of MPE in three specifications of the game



Notes: Panel (a)-(b) plots payoff maps and panel (c)-(d) cdf plots of efficiency by equilibrium type for two versions of the Bertrand investment pricing game. In panel (a) and (c) the case of deterministic alternating moves and non-strictly monotonic one step stochastic technological progress. Parameters in this case are $\eta = 0$, $\beta = 0.9592$, $k_1 = 5$, $k_2 = 0$, $f(1|1) = f(2|2) = 0$, $f(2|1) = f(1|2) = 1$, $c_{tr} = 1$, $n_c = 4$. In panel (b) and (d) we plot the payoffs and the distribution of efficiency for the simultaneous move game with deterministic one step technology. Leapfrog equilibria are defined as having positive probability to invest by the cost follower along the equilibrium path, mixed strategy equilibria are defined as involving at least one mixed strategy stage equilibrium along the equilibrium path.

$\pi(c|c) > 0$ for some c . The opposite of the latter condition ensures unique equilibrium in this game according to Theorem 5, but multiple equilibria is a typical outcome in the alternative move game with “sticky” state of the art technology. Consistent with Theorem 10 the set of equilibrium payoffs is a strict subset of the triangle, showing that it is not possible to achieve the monopoly payoffs (corners) or the zero profit mixed strategy equilibrium payoff (origin) in this case. As before, we have used the size of the plotted payoff points to indicate the number of repetitions of the payoff

points, but now we use the color of plotted equilibrium payoffs to indicate the efficiency index. Red (hot) indicates high efficiency payoffs, and blue (cool) indicates lower efficiency payoffs.

We see a clear positive correlation between payoff and efficiency in panel (a) — there is a tendency for the points with the highest total payoffs (i.e. points closest to the line connecting the monopoly outcomes) to have higher efficiency indices. The CDFs of efficiency levels in panel (c) shows that 1) overall efficiency is reasonably high, with the median equilibrium having an efficiency index in excess of 97%, and 2) the maximum efficiency of the equilibria involving mixed strategies along the equilibrium path is strictly less than 100%.

In panels (b) and (d) of Figure 4 we plot the set of equilibrium payoffs and distribution of equilibrium efficiency indices for a simultaneous move investment under the deterministic technology process. In accordance with Theorems 7 and 8 now the monopoly and zero profit outcomes are present among the computed MPE equilibria of the model. Overall, the equilibria in this game are less efficient compared to the equilibria in the alternating move game displayed in the top row panels, but the tendency of more efficient equilibria to be located closer to the “monopoly” frontier remains. The additional source of inefficiency in this game is redundancy of simultaneous investments, which appear in the mixed strategy equilibria. It is clearly seen in the cumulative distribution plot in panel (d) that even though more than 30% of the equilibria are approaching full efficiency¹⁷, the mixed strategy equilibria are not among them. Instead, the distribution of their efficiency indices is stochastically dominated by the distribution of efficiencies in all the equilibria of the game.


We formalize the above discussion in the following theorem.

Theorem 13 (Inefficiency of mixed strategy equilibria). *A necessary condition for efficiency in the dynamic Bertrand investment and pricing game is that along MPE path only pure strategy stage equilibria are played.*

Figure 5 establishes the existence of *fully efficient leapfrogging equilibria*. Panel (a) of figure 5 plots the set of equilibrium payoffs in a simultaneous move investment game where there are four possible values for state of the art costs $\{0, 1.67, 3.33, 5\}$ and technology improves deterministically. Recall that the payoff points colored in dark red are 100% efficient, so we see that there

¹⁷To be exact, 15.22% have efficiency index 0.9878 and the same fraction of equilibria is fully efficient.

Figure 5: Efficiency of equilibria



Notes: Panel (a) and panel (b) plots the apex payoff map and distribution of efficiency indices for the simultaneous move game. 25.88% of all equilibria are fully efficient. The most efficient mixed strategy equilibrium has the efficiency index 0.99998 but does not violate Theorem 13. Panel (c) displays the simulated investment profile from a fully efficient “monopoly” equilibrium, while panel (d) displays the example of fully efficient equilibrium that involves leapfrogging.

are a number of other *non-monopoly equilibria* that can achieve *full efficiency*. The significance of this finding is that we have shown that it is possible to obtain competitive equilibria where leapfrogging by the firms ensures that consumers receive some of the surplus and benefits from technological progress without a cost in terms of inefficient investment such as we have observed occurs in mixed strategy equilibria of the game where socially inefficient excessive investment results in lower prices to consumers but at the cost of zero expected profits to firms. Notice, however, that even the least efficient mixed strategy equilibrium still has an efficiency of 96%, so that in this

particular example the inefficiency of various equilibria may not be a huge concern.

Panels (c) and (d) of Figure 5 plot the simulated investment profiles of two different equilibria. Panel (c) shows the monopoly equilibria where firm 2 is the monopolist investor. The socially optimal investment policy is to make exactly two investments: the first when costs have fallen from 5 to 1.67, and the second when costs have fallen to the absorbing value of 0. Panel (d) shows the equilibrium realization from a pure strategy equilibrium that involves leapfrogging, yet the investments are made at exact same time as the social planner would do. After firm 1 invests when costs reach 1.67 (consumers continue to pay the price $p_1 = 5$), in time period 5 it is leapfrogged by firm 2 who becomes the permanent low cost producer. At this point a “price war” brings the price down from 5 to 1.67, which becomes new permanent level.

We conclude that the leapfrogging equilibria may be fully efficient if the investments are made in the same moments of time as the social planner or monopoly would invest, but in these equilibria consumers also benefit from the investments because the price decreases in a series of permanent drops.

Theorem 14 (Existence of efficient non-monopoly equilibria). *In both the simultaneous move and alternating move investment games, there exist fully efficient non-monopoly equilibria.*

Proof. The proof is by example shown in Figure 5. An example of a fully efficient non-monopoly equilibrium when the firms move alternately (in deterministic fashion) can be constructed as well¹⁸.

□

While we find that efficient leapfrogging occur generically as equilibria in the simultaneous move investment game, the result that there exist efficient leapfrogging equilibria in the alternating move investment game should be viewed as a special counterexample, and that we typically do not get fully efficient leapfrogging equilibria in alternating move games with sufficient numbers of states and when investment costs are “reasonable” in relation to production costs (i.e. where the cost of building a new plant $K(c)$ is not too low). However due to the vast multiplicity of

¹⁸Let the possible cost states be $\{0, 5, 10\}$, assume deterministic technological progress, the cost of investing $K = 4$, and the discount factor $\beta = 0.95$. Then the socially optimal investment strategy is for investments to occur when $c = 5$ and $c = 10$, and these investments will occur at those states in the unique equilibrium of the game, but where one firm makes the first investment at $c = 5$ and the opponent makes the other investment when $c = 0$. These investments clearly involve leapfrogging that is also fully efficient.

equilibria in the simultaneous move investment game, we have no basis for asserting that efficient leapfrogging equilibria are any more likely to arise than other more inefficient equilibria.

We conclude this subsection by commenting that our numerical calculations of equilibria and the efficiency of equilibria lead to a general finding is that for the equilibria which are inefficient, the inefficiency is caused by *excessive frequency of investment* rather than *underinvestment*. In simultaneous move games we already noted that another source of inefficiency is *redundant, duplicative investments* that occur only in mixed strategy equilibria. We noted that while mixed strategy equilibria also exist in the alternating move investment game, duplicative simultaneous investments cannot occur by the assumption that only one firm can invest at any given time. Thus, the inefficiency of the mixed strategy equilibria of the alternating move games is generally a result of excessively frequent investment under the mixed strategy equilibrium. However it is important to point out that we have constructed examples of inefficient equilibria where there is *underinvestment* relative to the social optimum. Such an example is provided in panel (b) of Figure 6 in the next section.

5.5 Leapfrogging, Rent-dissipating Preemption and the Riordan and Salant's Conjecture

In this section we consider the Riordan and Salant conjecture that was discussed in section 2. Riordan and Salant conjectured that regardless of whether the firms move simultaneously or alternately, or whether technological progress is deterministic or stochastic, the general outcome in all of these environments should be that of *rent-dissipating preemption*, a situation where only one firm invests and does so sufficiently frequently in order to deter its opponent from investing. These frequent preemptive investments fully dissipate any profits the investing firm can expect to earn from preempting its rival (and hence also dissipating all social surplus as well). We first confirm their main result stated in terms of our model.


Theorem 15 (Riordan and Salant, 1994). *Consider a continuous time investment game with deterministic alternating moves. Assume that the cost of investment is independent of c , $K(c) = K$ and is not prohibitively high in the sense of inequality (21). Further, assume that technological progress is deterministic with state of the art costs at time $t \geq 0$ given by the continuous, non-decreasing function $c(t)$ and continuous time interest rate $r > 0$. Assume that the continuous time*

analog of the condition that investment costs are not too high holds, i.e. $C(0) > rK$. Then there exists a unique MPE of the continuous time investment game (modulo relabeling of the firms) that involve preemptive investments by one or the other of the two firms and no investment in equilibrium by its opponent. The discounted payoffs of both firms in equilibrium is 0, so the entire surplus is wasted on excessively frequent investments by the preempting firm.

Corollary 15.1 (Riordan and Salant, 1994). *The continuous time equilibrium in Theorem 15 is a limit of the unique equilibria of a sequence of discrete time games where $\beta = \exp\{-r\Delta t\}$ and per period profits of the firms, $r_i(c_1, c_2)$, are proportional to Δt and the order of moves alternates deterministically, for a deterministic sequence of state of the art costs given in equation (2) as $\Delta t \rightarrow 0$.*

The proof of Theorem 15 and Corollary 15.1 is given in Riordan and Salant (1994) who used a mathematical induction argument to establish the existence of the continuous time equilibrium as the limit of the equilibria of a sequence of discrete time alternating move investment games. We restate it here to show their result emerges as a limiting case of the dynamic Bertrand duopoly game of pricing and investment that we analyze here. In Figure 6 we plot simulated MPE for three versions of the Bertrand investment pricing game with deterministic alternating move and strictly monotonic technological progress. In the panel (a) we let the length of the time periods be relatively small to provide a good discrete time approximation to Riordan and Salant's model in continuous time. In panel (b) we decrease the number points of support of the marginal cost and increase the length of the time period. In panel(c) in addition we lower investment cost. These three examples demonstrate that preemptive rent-dissipating investments indeed can happen in discrete time when the cost of investing in the new technology $K(c)$ is large enough relative to per period profits, but fails when the opposite is true as shown in panels (b) and (c). In discrete time, both duopolist have temporary monopoly power that can lead to inefficient under-investment as shown in the equilibrium realization in panel (b) or leapfrogging as shown in panel (c). Since per period profits are proportional to the length of the time period, the latter increases the value of the temporary cost advantage a firm gains after investment in the state of the art technology. If investment costs are sufficiently low relative to per period profits, it can actually be optimal for the cost follower to leapfrog the cost leader, in the limiting case even for a one period cost leadership.

Figure 6: Production and state of the art costs in simulated MPE: continuous. vs. discrete time




Notes: The figure plots simulated MPE equilibria for three versions of the Bertrand investment pricing game with deterministic alternating move and strictly monotonic technological progress. In panel (a) we present a discrete time approximation to Riordan and Salant's model in continuous time, with parameters $\eta = 0$, $\beta = 0.9512$, $k_1 = 2$, $k_2 = 0$, $\pi(c|c) = 0$, $f(1|1) = f(2|2) = 0$, $f(2|1) = f(1|2) = 1$, $n_c = 100$, $\Delta t = 0.25$. In panel (b) we decrease the number of discrete support points for c to $n_c = 25$ and increase the length of the time period such that $\Delta t = 1$ adjusting per period values. In panel(c) we in addition lower investment costs by setting $k_1 = 0.5$.

While the Riordan and Salant result of strategic preemption with *full rent dissipation* only holds in the continuous time limit $\Delta t \rightarrow 0$, their conclusion that investment preemption will occur is robust to discreteness of time in the sense that we find investment preemption as the only equilibria in our discrete time numerical solutions when Δt is sufficiently small. Thus, there is a “neighborhood” of Δt about the limit value 0 for which their unique preemption equilibrium also holds in a discrete time framework.

However what is not robust to modeling the game in discrete versus continuous time is the conclusion that preemption is fully inefficient and rent dissipating. In discrete time the preempting firm does earn positive profits and efficiency of equilibrium is positive, though profits and efficiency both converge to 0 as $\Delta t \rightarrow 0$. The intuitive reason that full rent dissipation result only merges in the continuous time limit is similar to the reasoning underlying the Coase conjecture: the shorter the time interval between moves the less exogenous “precommitment” there is for the preempting firm, so the more frequent alternation of the right to invest as $\Delta t \rightarrow 0$ makes the game increasingly “contestable” resulting asymptotically in the elimination of all rents to the preempting firm. In discrete time, the low cost leader is willing and able to undertake excessive preemptive

Figure 7: Production and state of the art costs in simulated MPE under uncertainty



Notes: The figure plots simulated MPE by type for four stochastic generalizations of the model illustrated in Figure 6.b. In panel (a) we consider random alternating moves where $f(1|1) = f(2|2) = 0.2$ and $f(2|1) = f(1|2) = 0.8$. In panel (b) we allow for non-strictly monotonic one step random technological improvement. In panel (c) we allow technological progress to follow a beta distribution over the interval $[c, 0]$ where c is the current best technology marginal cost of production. The scale parameters of this distribution is $a = 1.8$ and $b = 0.4$ so that the expected cost, given an innovation, is $c * a / (a + b)$. Panel (d) plots an equilibrium path from the simultaneous move game. Unless mentioned specifically remaining parameters are as in panel (b) of Figure (6).

investments to protect its position as the low cost leader, but only if investment costs $K(c)$ are high enough relative to within period profits.

Allowing for *random alternation* in the right to move, we obtain a unique pure strategy equilibrium, since random alternations does not violate the sufficient conditions for uniqueness given in Theorem 5. Random alternation of the right to move destroys the ability to engage in strategic preemption and creates the opportunity for leapfrogging, since firms cannot perfectly control when they have the right to move and thereby they cannot time their preemptive investments exactly. Fig-

ure 7, panel (a) gives an example of a simulated equilibrium path when the right to move alternates randomly. While this equilibrium path depicts a unique pure stately equilibrium, we clearly see the leapfrogging pattern.

From Theorem 5 it follows that if there is positive probability of remaining with the same state of the art cost c for more than one period of time, i.e. $\pi(c|c) > 0$, the main results of Riordan and Salant (1994) will no longer hold. We may have multiple equilibria, there will be leapfrogging, and full rent dissipation fails.

Figure 7 presents simulated equilibrium paths when we introduce randomness in the evolution of the state of the art technology, the order of moves in the alternating move game, or possibility for simultaneous investment. All panels exhibit leapfrogging, reflecting the statement that stochasticity in the model presents the cost follower with more opportunities to leapfrog its opponent and makes it harder for the cost leader to preempt leapfrogging. Overall, in presence of uncertainty, the game becomes much more contestable.

Corollary 15.2. *(Limits to Riordan and Salant result) Preemption does not hold when (1) cost of investment $K(c)$ is sufficiently small relative to per period profits, (2) investment decisions are made simultaneously, (3) the right to move alternates randomly, (4) $\pi(c|c) > 0$, i.e. under other than strictly monotonic technological progress.*

Proof. The proof is by counter examples which are shown in Figure 6 and 7. □

The vast majority of MPE equilibria in the many specifications of the game we have solved using the RLS algorithm exhibited leapfrogging. We conclude that the Riordan and Salant conjecture that rent dissipation and full preemption are *general* results, is incorrect. Moreover, it appears that their results are not robust to any of the mentioned assumptions, at least in the discrete time analog of their model. However with the exception of the full rent dissipation result, we believe that there is a *neighborhood* about the limiting set of parameter values that Riordan and Salant used to prove Theorem 15 for which their conjectured preemption equilibrium will continue to hold, at least with high probability. That is, as long as the time between moves Δt is sufficiently small, that the order of moves is sufficiently close to deterministically alternating (in the sense that the transition probability matrix for which firm moves has diagonal elements that are sufficiently close to zero), and

the stochastic process governing technological improvements is sufficiently close to a deterministic sequence, then there will still be a unique equilibrium of the game, and with probability tending to 1, the equilibrium will involve strategic preemption.

6 Conclusions

We have developed a model of Bertrand price competition that allows for sequential investments in a stochastically improving technology. We have resolved the *Bertrand investment paradox* by showing that it is not only possible that Bertrand price competitors will invest in cost reducing technologies, their attempts to leapfrog each other to gain temporary low cost leadership will generally lead to inefficient *overinvestment* relative to the social optimum.

Our analysis provides new insights into a conjecture of Riordan and Salant (1994) that leapfrogging will not occur in equilibrium when firms are Bertrand price competitors. In a continuous time alternating move framework, they proved that *investment preemption* is the only possible equilibrium. While our analysis confirms that this result is correct, we have showed that their conjecture that strategic preemption is the *only* possible equilibrium is incorrect. Instead we have shown that leapfrogging investments is the dominant mode of competition in the equilibria of these models, particularly when investment decisions are made simultaneously. We have also shown that leapfrogging and not strategic preemption emerges in alternating move versions of the game provided either a) technology evolves stochastically, or b) the order of moves is stochastic.

We believe a separate insight from our analysis is a new interpretation for price wars. In our model price wars occur when a high cost firm leapfrogs its opponent to become the new low cost leader. It is via these periodic price wars that consumers benefit from technological progress and the competition between the duopolists. However, what we find surprising is complex set of price and investment dynamics that are possible from such a simple model. We find a large number of equilibria ranging from monopoly outcomes to highly competitive equilibria where the firms compete each others' profits to zero. In between are equilibria where leapfrogging investments are relatively infrequent so that consumers see fewer benefits from technological progress in the form of lower prices. It remains an open question as to whether these sorts of equilibrium outcomes are a theoretical curiozum, or whether this model can be extended to provide insights into a variety of

possible competitive behaviors in actual markets.

As we noted in the introduction, our paper is not the first to establish the possibility of leapfrogging equilibria in a dynamic extension of the classic Bertrand model of price competition. After we completed our analysis, we became aware of the work of Giovannetti (2001), who appears to have provided the first analysis of Bertrand competition with cost-reducing investments in a framework similar to ours'. The main differences between our setup and Giovannetti's is that improvements in technology occur deterministically in his model, with the cost of investing in the state-of-the-art production facility declining geometrically in each period. He established in this environment that there are leapfrogging equilibria in which investments occur in every period, but with the two firms undertaking leapfrogging investments alternately in every period. Under a constant elasticity of demand formulation where the demand elasticity is greater than 1, Giovannetti showed that these alternating leapfrogging investments by the two firms will continue *forever*.

In our model, in the absence of an outside good from which "new customers" can be drawn, the leapfrogging will generally not occur forever, but will end after a finite span of time with probability one. This result, however, is dependent on assumptions about how the cost of adopting new technology changes over time. If this adoption costs also decreases over time at a sufficiently rapid rate and if technological progress results in costs only asymptoting to 0 rather than reaching 0 in a finite amount of time with probability 1, then we expect it would be possible to show that leapfrogging investments could continue indefinitely in our model as well.

Giovannetti's analysis did not trace out the rich set of possible equilibria that we have found in our model, including the possibility of "sniping" where a firm that has been the high cost follower for extended periods of time suddenly invests at the "last minute" (i.e. when the state-of-the-art marginal cost is sufficiently low that any further investments are no longer economic), thereby displacing its rival to attain a permanent low cost leadership position. This is one of the benefits of being able to solve the model numerically, which facilitates the study of possible equilibrium outcomes.

Our analysis also contributes to the long-standing debate about the relationship of market structure and innovation. Schumpeter (1939) argued a monopolist will innovate more rapidly than a competitive industry since the monopolist can fully appropriate the benefits of R&D or other cost-reducing investments, whereas some of these investments would be dissipated in a competitive

market. However Arrow (1962) argued that innovation (or new technology adoption) under a monopolist will be slower than would occur in a competitive market which is in turn lower than the rate of innovation that would be chosen by a social planner. Both types of results have appeared in the subsequent literature. For example, in the R&D investment model analyzed by Goettler and Gordon (2011), the rate of innovation under monopoly is higher than under duopoly but still below the rate of innovation that would be chosen by a social planner. These inefficiencies are driven in part by the existence of externalities such as *knowledge spillovers* that are more commonly associated with R&D investments, but which are not present in our framework.

We have shown that the rate of adoption of new cost-reducing technologies under the duopoly equilibrium is generally *higher* than the monopoly or socially optimal solution. As we noted in the introduction, this result may be counter-intuitive given the Bertrand investment paradox, where we might expect that the lower rents that accrue to the investing duopolists as a result of Bertrand price competition should *lower* their incentive to make new investments in cost reducing technologies. However we showed that the desire to *leapfrog* or to *preempt* causes the duopolists to collectively invest more in cost reducing technologies than a social planner. At the same time, we view our analysis as rather specialized, and would caution against using it to draw any general empirical conclusions about the effect of market structure on innovation. In particular, it is important to extend the model to allow for entry and exit of firms.¹⁹

A disturbing aspect of our findings from a methodological standpoint is the plethora of Markov perfect equilibria present in a very simple extension of the standard static textbook model of Bertrand price competition, which is reminiscent of the “Folk theorem” for repeated games. In fact, Hörner *et. al.* (2011) have already proved that a version of the Folk theorem holds for a class of discounted stochastic games as $\beta \rightarrow 1$ (recall β is the common discount factor of the players in the game) where the limiting set of payoffs is independent of the initial state. While this assumption does not hold in our game (the limiting set of payoffs does depend on the starting point of the game) and our results on the multiplicity of equilibria hold for any $\beta \in [0, 1)$ not just in the limit, it is still the case that there is an essential similarity in their results and our finding that the con-

¹⁹We refer readers to the original work by Reinganum (1985) as well as recent work by Acemoglu and Cao (2011) and the large literature they build on. It is an example of promising new models of endogenous innovation by incumbents *and* new entrants. In their model entrants are responsible for more “radical” innovations that tend to replace incumbents, who focus on less radical innovations that improve their existing products.

vex hull of the set of equilibrium payoffs in the simultaneous move game is the triangle. Though we have shown that the set of payoffs shrinks dramatically to a strict subset of the triangle under the alternating move game, there will generally still be a huge multiplicity of equilibria since we showed that uniqueness of equilibrium is only obtained under a fairly special assumption of strictly monotonic technological progress.

Thus, from the standpoint of empirical IO, though we have demonstrated how leapfrogging can be viewed as an endogenous solution to the “anti-coordination problem” that solves the Bertrand investment paradox by enabling firms to avoid inefficient and unprofitable simultaneous investments, our paper leaves unsolved the more general problem of how it is that firms, without any explicit communication, can coordinate on a single equilibrium in these games when there is generally such a vast multiplicity of possible equilibria.

References

- [1] Acemoglu, D. and D. Cao (2011) “Innovation by Entrants and Incumbents” working paper, Georgetown University.
- [2] Anderson, S., A. DePalma and J. Thisse (1992) *Discrete Choice Theory of Product Differentiation* MIT Press.
- [3] Arrow, K.J. (1962) “Economic Welfare and the Allocation of Resources for Inventions” in Nelson, R.R. (ed.), *The Rate and Direction of Inventive Activity: Economic and Social Factors* Princeton University Press. Princeton.
- [4] Baumol, W.J. and Panzar, J.C. and Willig, R.D (1982) *Contestable Markets and the Theory of Industry Structure* Harcourt, Brace Jovanovich.
- [5] Baye, M. and D. Kovenock (2008) “Bertrand Competition” in *New Palgrave Dictionary of Economics* 2nd Edition.
- [6] Bertrand, J. (1883) [Review of] “*Théorie Mathématique de la Richesse Social* par Léon Walras: *Recherches sur les Principes de la Théorie du Richesses* par Augustin Cournot” *Journal des Savants* **67** 499–508.
- [7] Bowlus, Audra J. and Shannon Seitz (2006) “Domestic Violence, Employment, and Divorce” *International Economic Review* **47-4** 1113–1149.
- [8] Cournot, A. (1838) *Recherches sur les Principes de la Théorie du Richesses* Paris: Hachette.
- [9] Doraszelski, Ulrich and Juan Escobar (2010) “A Theory of Regular Markov Perfect Equilibria in Dynamic Stochastic Games: Genericity, Stability, and Purification” *Theoretical Economics* **5** 369–402.
- [10] Doraszelski, Ulrich and Kenneth Judd (2011) “Avoiding the Curse of Dimensionality in Dynamic Stochastic Games” *Quantitative Economics* **3-1** 53–93.
- [11] Doraszelski, Ulrich and Mark Satterthwaite (2010) “Computable Markov-Perfect Industry Dynamics” *Rand Journal of Economics* **41-2** 215–243.
- [12] Dubé, Jean-Pierre, Güntter Hitsch and Pradeep K. Chintagunta (2010) “Tipping and Concentration in Markets with Indirect Network Effects” *Marketing Science* **29-2** 216–249.
- [13] Fudenberg, D. R. Gilbert, J Stiglitz and J. Tirole (1983) “Preemption, leapfrogging and competition in patent races” *European Economic Review* **22-1** 3–31.
- [14] Fudenberg, D. and J. Tirole (1985) “Premption and Rent Equalization in the Adoption of New Technology” *Review of Economic Studies* **52** 383–401.
- [15] Gilbert, R and D. Newbery (1982) “Pre-emptive Patenting and the Persistence of Monopoly” *American Economic Review* **74** 514–526.
- [16] Giovannetti, Emanuelle (2001) “Perpetual Leapfrogging in Bertrand Duopoly” *International Economic Review* **42-3** 671–696.

- [17] Goettler, Ronald and Brett Gordon (2011) “Does AMD spur Intel to innovate more?” *Journal of Political Economy* **119-6** 1141–2000.
- [18] Hörner, J., Sugaya, T., Takahashi, S. and Vieille, N. (2011) “Recursive Methods for Discounted Stochastic Games: An Algorithm for $\delta \rightarrow 1$ and a Folk Theorem” *Econometrica* **79-4** 1277–1318.
- [19] Iskhakov, F., Rust, J. and B. Schjerning (2013) “Recursive Lexicographical Search: A Solution Method for Finding All Equilibria of Dynamic Games”
- [20] Kovenock, D. and R. J. Deneckere (1996) “Bertrand-Edgeworth Duopoly with Unit-Cost Asymmetry” *Economic Theory* **8-1** 1–25.
- [21] Kreps, D. and J. Scheinkman (1983) “Quantity Precommitment and Bertrand Competition Yield Cournot Outcomes” *Bell Journal of Economics* **14-2** 326–337.
- [22] Kushner, Harold J. and Paul G. Dupuis (2001) *Numerical Methods for Stochastic Control Problems in Continuous Time* second edition, Springer Verlag.
- [23] Lagunoff, R. and A. Matsui (1997) “Asynchronous Choice in Repeated Coordination Games” *Econometrica* **65-6** 1467–1477.
- [24] Reinganum, J. (1981) “On the Diffusion of New Technology — A Game-Theoretic Approach” *Review of Economic Studies* **153** 395–406.
- [25] Reginanum, J. (1985) “Innovation and Industry Evolution” *Quarterly Journal of Economics* **100-1** 81–99.
- [26] Riordan, M. and D. Salant (1994) “Preemptive Adoptions of an Emerging Technology” *Journal of Industrial Economics* **42-3** 247–261.
- [27] Riordan, M. (2010) “Quality Competition and Multiple Equilibria” working paper, Department of Economics, Columbia University.
- [28] Rosenkranz, Stephanie (1996) “Quality Improvements and the Incentive to Leapfrog” *International Journal of Industrial Organization* **15** 243–261.
- [29] Routledge, Robert R. (2010) “Bertrand competition with cost uncertainty” *Economics Letters* **107** 356–359.
- [30] Rust, John (1986) “When Is it Optimal to Kill Off the Market for Used Durable Goods?” *Econometrica* **54-1** 65–86.
- [31] Scherer, F. (1967) “Research and Development Resource Allocation Under Rivalry” *Quarterly Journal of Economics* **81** 359–394.
- [32] Schumpeter, J.A. (1939) *Business Cycles: A Theoretical, Historical and Statistical Analysis of the Capitalist Process*. New York: McGraw Hill.
- [33] Swan, P. (1970) “Market Structure and Technological Progress: The Influence of Monopoly on Product Innovation” *Quarterly Journal of Economics* **84** 627–638.
- [34] Vickers. M. (1986) “The Evolution of Market Structure When There is a Sequence of Innovations” *Journal of Industrial Economics* **35** 1–12.