A Dynamic Model of Leap-Frogging Investments and Bertrand Price Competition†

Fedor Iskhakov
University Technology Sydney and Frisch Center, University of Oslo

John Rust‡
University of Maryland

Bertel Schjerning
University of Copenhagen

April, 2011

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 and attain, at least temporarily, low-cost leadership. The model resolves a paradox about investing in the presence of Bertrand price competition: if both firms simultaneously invest in the current state-of-the-art production technology and thereby attain the same marginal cost of production, the resulting price competition drives the price down to marginal cost and profits to zero. Thus, it would seem that neither firm can profit from undertaking the cost-reducing investment, so the firms should not have any incentive to undertake cost-reducing investments if they are Bertrand price competitors. We show this simple intuition is incorrect. We formulate a dynamic model of price and investment competition as a Markov-perfect equilibrium to a dynamic game. We show that even when firms start with the same marginal costs of production there are equilibria where one of the firms invests first, and leapfrogs its opponent. In fact, there are many equilibria, with some equilibria exhibiting asymmetries where there are extended periods of time where only one of the firms does most of the investing, and other equilibria where there are alternating investments by the two firms as they vie for temporary low cost leadership. Our model provides a new interpretation of the concept of a "price war". Instead of being a sign of a breakdown of tacit collusion, in our model price wars occur when one firm leapfrogs its opponent to become the new low cost leader.

Keywords: duopoly, Bertrand-Nash price competition, leapfrogging, cost-reducing investments, dynamic models of competition, Markov-perfect equilibrium, tacit collusion, price wars, coordination games, pre-emption

JEL classification: D92, L11, L13

†Preliminary: This version of this paper is not for general distribution and quotation. We thank Joseph E. Harrington, Jr. for helpful comments and suggestions, but he is of course not responsible for any errors or misconceptions in this paper.

‡**Correspondence address:** Department of Economics, University of Maryland, College Park, MD 20742, phone: (301) 405-3489, email: `jrust@gemini.econ.umd.edu`

1 Introduction

This paper provides a dynamic extension of the static textbook Bertrand-Nash duopoly game by allowing firms to make investment decisions as well as pricing decisions. At any point in time, firms are assumed to have the option to replace their current production facilities with a new state-of-the-art production facility. If the state-of-the-art has improved since the time the firm invested in its current production facility, the investing firm will be able to produce at a lower marginal cost — both relative to its own previous costs and potentially also lower than its rival. We use the term "leapfrogging" to describe the longer run competition over investments between the two duopolists when an investment by one firm enables it to produce at a lower cost than its rival and attain, at least temporarily, a position of low cost leadership.

When the competing firms set prices in accordance with the Bertrand equilibrium under constant returns to scale production technologies, then in the absence of capacity constraints, the high cost firm will earn zero profits. The motivation for the high cost firm to undertake a cost-reducing investment is, of course, to obtain a production cost advantage over its rival. The firm that is the low cost leader does earn positive profits by charging a price equal to the marginal cost of production if its higher cost rival. However, if both firms have the same marginal cost of production, both firms set a price equal to their common marginal cost and earn zero profits. Baye and Kovenock describe this as the *Bertrand paradox*.[1]

A new paradox arises when we try to extend the static Bertrand price competition to a dynamic context where the firms are free at any time to invest in the state-of-the-art production technology. Since both firms have the option to acquire the state-of-the-art technology at any time at the same investment cost, there is no guarantee that investment can bring about anything more than a temporary period of low cost leadership. As such, the market we analyze can be regarded as *contestable* (Baumol, Panzar and Willig, 1982), and each firm can deny the other the opportunity to become the low cost leader by investing at the same time. However when investment competition leads both firms to invest at the same time, both will have the same state-of-the-art marginal cost of production, and the resulting Bertrand price competition will ensure that post-investment profits will be zero. This would seem to provide little incentive for either firm to undertake the investment in the first place. As a result, casual reasoning would suggest that *Bertrand duopolists may*

[1] According to Baye and Kovenock, Bertrand did not realize that the perfectly competitive outcome emerges as the equilibrium solution to price competition. The discuss how Bertrand, in his 1883 review of Cournot's 1838 book, "described how, in Cournot's duopoly environment where identical firms produce a homogeneous product under a constant unit cost technology, price competition would lead to price undercutting and a downward spiral of prices. Bertrand erroneously reasoned that this process would continue indefinitely, thereby precluding the existence of an equilibrium." (p. 1).

not have any incentive to undertake cost-reducing investments. Thus, the challenge is to show that there are equilibria where the firms do have an incentive to invest, even when both behave as Bertrand price competitors at every instant. We refer to this as the *Bertrand investment paradox.*

We provide a resolution of the Bertrand investment paradox by solving a dynamic, infinite horizon extension of the Bertrand model of price competition. The extended version of the Bertrand model allows the competing firms to invest in improved technology in addition to setting prices. We solve for Markov-perfect equilibria to this dynamic game, including extensions where where each firm has private information about idiosyncratic adjustment costs/benefits associated with undertaking an investment at any particular point in time. We show that even in complete information versions of this model, and even when firms start with the same marginal costs of production, not investing is never an equilibrium outcome whenever the gains from investing in the new technology are sufficiently high relative to the cost of acquiring it.

However will these "investment equilibria" display leapfrogging behavior? With the exception of the work of Giovannetti (2001), the main result in the previous literature on investment under duopoly with downstream Bertrand price competition is that *leapfrogging investments cannot occur in equilibrium.* Instead, this literature has shown that all equilibria involve *pre-emption* — one of the duopolists undertakes all cost-reducing investments at times determined to deter any leapfrogging investments by the firm's rival. For example a line of work by Gilbert and Newbery (1982), Vickers (1986) and Riordan and Salant (1994) proved that pre-emption is the *only* equilibrium.

Riordan and Salant 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 technological progress is deterministic and thus perfectly predictable. In this framework the equilibrium strategies consist of a sequence of dates at which firms plan to upgrade their production facilities. By paying a fixed upgrade cost at each upgrade date, the investing firm is able to acquire the state-of-the-art production technology, which Riordan and Salant assume is characterized as a constant returns to scale production technology whose marginal cost of production declines deterministically over time.

Riordan and Salant 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: "The leading firm has more to lose from the lagging firm's adoption of a new technology than the lagging firm has to gain. Therefore the leading firm will always preempt the laggard, waiting for the last possible instant to do so." (p. 255).

We will show that leapfrogging is a possible equilibrium outcome in a game where the duopolists make *simultaneous* investment and pricing decisions. Riordan and Salant's analysis assumed the firms make simultaneous choices of prices, but *alternating* choices of whether or not to invest. They 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 and that the timing of investment decisions is crucial to the nature of equilibrium outcomes we observe in these models. We do this by by characterizing *all equilibria* to a fully simultaneous move formulation of the game of pricing and investment, including mixed strategy equilibria. We show that generically there are multiple possible Markov-perfect equilibria in these models and that the choice of equilibria in low cost states affect the set of possible equilibria at higher cost states.

Further, our model differs from Riordan and Salant (1994) and other related work in this literature (e.g. Giovannetti 2001) in that we assume innovations in the state-of-the-art production technology arrive *stochastically.* However it is the simultaneous-move aspect of investments and not the uncertainty over technological progress that drives our results, since Giovannetti (2001) also obtained leapfrogging outcomes in a model of deterministic technological improvement but where firms make simultaneous decisions about investment.

We show that the multiplicity of equilibria in our model has aspects similar to the literature on equilibria of supergames, even though the Markov-perfect framework can be viewed as an attempt to eliminate extraneous equilibria that arise in supergames. We view the multiplicity of equilibria as a consequence of the simultaneity of investment decisions and the fact that the choice of investment strategies by the duopolists can be viewed as a *dynamic coordination game.* Leapfrogging is one of many possible equilibrium solutions to this dynamic coordination game. We prove that when the firms start from a symmetric situation with equal marginal costs of production, the set of all discounted expected equilibrium payoffs for the two duopolists is a *triangle* whose vertices consist of two *asymmetric monopoly pure strategy equilibria* and a third *zero profit symmetric mixed strategy equilibrium.*

Interestingly, we prove that the monopoly equilibria involve efficient technology adoption strategies whereas other duopoly equilibria including leapfrogging and mixed strategy equilibria where there is positive probability of investment by both firms are *inefficient.* In contrast, we show that the duopoly equilibria involving pre-emption in Riordan and Salant's model (the analog of our two efficient monopoly equilibria) are inefficient. In essence, the threat of investment by the high cost firm forces the low cost leader to invest at a rate that is faster than is socially optimal (and faster than the rate monopolist would undertake such investments).

We show that equilibria involving leapfrogging behavior (including the symmetric zero profit mixed strategy equilibrium) that can be efficient or inefficient. In the efficient leapfrogging equilibria, at most one of the firms invests at any node in the game tree, so the leapfrogging equilibrium results in coordination between the two firms that avoids the inefficiency of duplicative investments. However we show that there are also *inefficient* equilibria involving leapfrogging. One source of inefficiency is that in some cases neither firm may invest in a situation where investment it is socially optimal (i.e. reduces the expected discounted value of production and adoption costs) to have invested. We also characterize situations where the inefficiency results from *excessive frequency of investments*, i.e. one or both firms invest in states where no investment would occur in a socially efficient (cost minimizing) investment strategy.

Simulations of numerical solution of the model reveal that the non-monopoly equilibria result in realizations that can involve both simultaneous and alternating investments by the two firms as they vie for temporary production-cost leadership over their opponent. However we show that there are also equilibria where one firm exhibits persistent low cost leadership over its opponent, 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.

Our model also 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 industrial organization literature, where price wars are a punishment device to deter tacitly colluding firms from cheating. The key difference in the prediction of our model compared to the standard model of tacit collusion is that price paths are piece-wise flat and monotonically declining in our model and price wars are very brief, lasting only a single period in our model, whereas in the model of tacit collusion, price wars can extend over multiple periods and prices are predicted to *rise* at the end of a price war.

We present the model in section 2. Our 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 zero cost absorbing state. In section 3 we show that the analysis of this simpler end game solution leads to key insights into the form of the full equilibria of the model which we solve and illustrate in section 4. In particular we prove that a weaker form of leapfrogging must hold in any mixed strategy equilibrium in the end game: *the high cost firm always has a higher probability of investing than its lost cost rival.* In section 5 we formulate and solve the social planner’s problem and characterize the investment strategy that maximizes total surplus. We show that unlike the pre-emption equilibria in Riordan and Salant (1994) the two “monopoly equilibria” are fully efficient and equivalent to the positive profit monopoly outcome under the constraint that the monopolist cannot charge a price higher than the initial marginal cost of production of the firm’s competitor. However we show that other duopoly equilibria can be *inefficient* often as a result of duplicative investments by the competing firms (such as occur in mixed strategy equilibria) but also due to investments that fail to be undertaken by either firm in states where investment is socially optimal, and investments that are undertaken by one of the two firms when it is not socially optimal to invest in new technology.

Thus, unlike the simple static Bertrand model of price competition where competition between even two firms leads to an efficient outcome, in our dynamic generalization of Bertrand competition we find that the duopoly equilibria are often inefficient, even though they do result in the benefits of technological improvement being passed on to consumers in the form of lower prices. We discuss related literature and offer some concluding comments and conjectures in section 6.

2 The Model

Suppose there are two firms producing an identical good. The firms are price setters and have no fixed costs and can produce the good at a constant marginal cost of c_1 (for firm 1), and c_2 (for firm 2). Later we will add time subscripts to these marginal costs, since both firms will have the option of replacing their current production facilities with state-of-the-art production facilities that have a potentially lower marginal cost of production, c. Shortly we will describe dynamics by which the state-of-the-art marginal cost c evolves over time. In this case, the marginal costs of each firm will also depend on time, t, since the firms may choose to replace their current production facilities with a state-of-the-art one.

We assume the production technology is such that neither firm faces capacity constraints, so that effectively, both firms can produce at any given time at what amounts to a constant returns to scale production technology. In the conclusion we will discuss an extension of our model to allow for capacity constraints, where investments can be used both to lower the cost of production and/or to increase the production capacity of the firm. The paper by Kreps 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.

However we believe that it is of interest to start by considering the simplest possible extension of the classic Bertrand price competition model to a multiperiod setting under the assumption that neither firm faces 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 equilibrium outcomes.

We note that in most real markets, firms are rarely capacity constrained. To our thinking, the more problematic aspect of the Bertrand model is not the assumption that firms have no capacity constraints, but rather, the assumption that one of the firms can capture the entire market by slightly undercutting its rival. Real world markets involves switching costs and other idiosyncratic preference factors that lead demand to be more inelastic than the perfectly elastic demand assumed in the standard Bertrand model of price competition. We think that one reason why firms are rarely capacity constrained is that contrary to the assumption underlying the classic Bertrand model, a firm cannot capture all of its opponent's customers by slightly undercutting its price.

Our model does allow for switching costs and idiosyncratic factors to affect consumer demand, so that demand can be less than perfectly elastic in our model. In this case, when one of the firms undercuts its rival's price, it does not succeed in capturing all of its rival's market share. In these versions of the model, leapfrogging behavior does not result in the large swings in market share that occur in the standard Bertrand model when demand is assumed to be infinitely elastic.

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 Bertrand price competition in this case leads directly to the "Bertrand investment paradox" that we noted in the Introduction. The ability of both firms to acquire (at a cost) the current state-of-the-art production technology, combined with the lack of any "loyalty" or inertia in their customers that enables one firm to steal all of its opponent's customers by slighly undercutting its price means that a very strong form of "contestability" holds in this case.

In particular, there is never any permanent advantage to being the low-cost leader: at any time the high cost follower could invest and acquire a state-of-the-art production facility that would enable it to produce at equal or lower cost than the current low cost leader. The only reason the high cost firm may not want to pay the cost necessary to acquire the state-of-the-art production technology is the fear than the rival will also do this and the resulting Bertrand price competition would eliminate or reduce any temporary profits that it would need to justify incurring the fixed costs of purchasing a new state-of-the-art production facility.

In this model, 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, although as we show below, 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 equilibria exist that enable firms to coordinate their investments in ways that are in some respects 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 that low cost leader firm to charge a high 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.

On the other hand, there are 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.

The shortcoming of our analysis is that there is nothing in our analysis to suggest which of these many possible equiibria might be "selected" by firms in any specific situation. To a large extent the problem of equilibrium selection is a topic that is outside the scope of this paper: our main goal is to show how to compute and characterize the set of all MPE in this model.

2.1 Consumers

As is typically done in the industrial organization literature, we extend the usual textbook model of competition between producers of homogeneous goods to allow some degree of monopolistic competition or switching costs. The simplest way to do this is to allow for idiosyncratic benefits or costs that each consumer experiences when they purchase one or the products offered by the two firms. Let the net benefit or payoff to a customer who buys from firm 1 be $u_1 = \sigma\tau_1 - p_1$ and the net benefit from buying from firm 2 be $u_2 = \sigma\tau_2 - p_2$. We can think of the vector (τ_1, τ_2) as denoting the "type" of a particular consumer. Assume there are a continuum of consumers and that the population distribution of (τ_1, τ_2) in the population has a Type 1 extreme value distribution and let $\sigma \geq 0$ be a scaling parameter. Then, as is well known from the literature on discrete choice (see, e.g. Anderson, dePalma and Thisse, 1992), the probability a consumer buys from firm 1 is

$$\Pi_1^\sigma(p_1, p_2) = \frac{\exp\{-p_1/\sigma\}}{\exp\{-p_1/\sigma\} + \exp\{-p_2/\sigma\}}.$$

Now, assuming that the mass (number) of consumers in the market is normalized to 1, we can define Bayesian-Nash equilibrium prices, profits, market shares for firms 1 and 2 in the usual way. That is, we assume that in each period of the dynamic game, the two firms simultaneously choose prices p_1 and p_2 that constitute mutual best responses, in the sense of maximizing each firm's profit taking into account the price set by the firm's opponent.

The Bertrand equilibrium pricing rules are defined by the functions $p_1^\sigma(c_1, c_2)$ and $p_2^\sigma(c_1, c_2)$ that solve the following fixed-point problem

$$\begin{aligned} p_1^\sigma(c_1,c_2) &= \underset{p_1}{argmax}\, \Pi_1^\sigma(p_1, p_2^\sigma(c_1,c_2))(p_1 - c_1) \\ p_2^\sigma(c_1,c_2) &= \underset{p_2}{argmax}\, \Pi_2^\sigma(p_1^\sigma(c_1,c_2), p_2)(p_2 - c_2). \end{aligned} \tag{1}$$

The classic Bertrand equilibrium arises as a special case in the limit as $\sigma \downarrow 0$. Then we have $p_1^\sigma(c_1, c_2) \downarrow p(c_1, c_2)$ and $p_2^\sigma(c_1, c_2) \downarrow p(c_1, c_2)$ where the equilibrium price $p(c_1, c_2)$ is given by

$$p(c_1, c_2) = \max[c_1, c_2]. \tag{2}$$

This is the usual textbook Bertrand equilibrium where the firm with the lower marginal cost of production sets a price equal to the marginal cost of production of the higher cost firm. Thus, the low cost firm can earn positive profits whereas the high cost firm earns zero profits. Only in the case where both firms have the same marginal cost of production do 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.

It is simple to extend this model to the case where there is an *outside good*, i.e. each consumer has the option of not buying the good. In this case we assume that the consumer receives a utility of $u_0 = \sigma\tau_0 - \gamma_0$. For concreteness, We assume that (τ_0, τ_1, τ_2) has a trivariate Type I (standardized) extreme value distribution. We assume these types are independently distributed across consumers, and in the dynamic version of the model, independently distributed over time for any specific consumer (thus, referring to τ as indexing the "type" of a consumer is an abuse of terminology, since the type of the consumer is changing over time in an unpredictable way).

It is not difficult to show that in the presence of the outside good, the probability a consumer buys from firm 1 is given by the classic logit formula:

$$\Pi_1^{\sigma}(p_1, p_2) = \frac{\exp\{-p_1/\sigma\}}{\exp\{-\gamma_0/\sigma\} + \exp\{-p_1/\sigma\} + \exp\{-p_2/\sigma\}}. \tag{3}$$

where γ_0 is a component of the utility of the outside good that does not vary over consumers.

2.2 Production Technology and Technological Progress

We now introduce our dynamic extension of the classical static Bertrand model of price competition by allowing the marginal costs of the two firms vary, endogenously, over time. The evolution of their marginal costs of production will cause the prices charged by the two firms to vary over time as well. 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. We assume that technological progress is an exogenous stochastic process: however the decisions by the firms of whether and when to adopt the state-of-the-art production technology are fully endogenous.

We start with the case where there isn't an outside good option present. It is not difficult to extend the analysis to account for the presence of an outside good, as long as the common component of its utility, γ_0, is time-invariant. If γ_0 evolves over time, it would complicate the analysis, since the value of this time-varying variable would have to be carried as one of the state variables in the game, and we would need to confront questions as to whether consumers have perfect foresight about its evolution, or whether they are uncertain about future values but know the probability law governing its evolution.

Suppose that over time the technology for producing the good improves, decreasing 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 machinery/plant to acquire this state-of-the-art production technology.

We assume that for any value of c, the production technology is such that there are 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 machinery and are embedded as part of the new investment cost $K(c)$. If either one of the firms purchases the state-of-the-art machinery, then after a one period lag (constituting the "time to build" the new production facility), the firm will be able to produce at the marginal cost of 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 cost c. This situation is reflected by choosing K to be a decreasing function of c. However it is also possible that technological improvements lower both the cost of the plant and the marginal cost of production. This situation can be captured by allowing K to be an increasing function of c. Then as c drops over time, so too will the associated fixed costs of investing in the state-of-the-art production technology.

If K is a decreasing function of c, then as c drops over time, the cost of investing in new production facilities increase over time. We can imagine that there can come a point where it is no longer economic to invest in the state-of-the-art because the degree of reduction in the marginal cost of production is insufficient to justify the fixed investment cost of the new plant. We will show below via numerical solution of the model, whether leapfrogging competition will result in steady price declines to consumers, or whether investment competition will eventually stop at some point, depends critically on both the level and slope of $K(c)$.

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 are lower than the cost of purchasing the state-of-the-art production plant at a cost of $K(c)$. 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 simultaneously 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_{jt}. Or firm j can pay the investment cost $K(c)$ in order to acquire the state-of-the-art production plant which will allow it to produce at the marginal cost c_t.

Given that there is a one period lag to build the new production facility, if a firm does invest at the start of period t, it will not be able to produce using its new state-of-the-art production facility until period

$t+1$. If there has been no improvement in the technology since the time firm 1 acquired its production machinery, then $c_{1t} = c_t$, and similarly for firm 2. If there has been a technological innovation since either firm acquired their current production facilities, we have $c_{jt} > c_t$. Thus, in general the state space S for this model is the following polyhedron in R^3, $S = \{(c_1, c_2, c) | c_1 \geq c$ and $c_2 \geq c$ and $c \leq c_0\}$ where $c_0 > 0$ is the initial state of technology.

Suppose that both firms believe that the technology for producing the good evolves stochastically and that the state-of-the-art marginal cost of production c_t evolves according to a Markov process with transition probability $\pi(c_{t+1}|c_t)$. Specifically, suppose that with probability $p(c_t)$ 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 $1 - p(c_t)$ the technology does improve, 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 for the general presentation of the model, making specific functional form assumptions about π is not required. For example, suppose the probability of a technological improvement is

$$p(c_t) = \frac{.01c_t}{1 + .01c_t}. \tag{4}$$

The timing of events in the model is as follows. At the start of time t each firm learns the current value of c_t and simultaneously decide whether or not to invest. Both firms know each others' marginal cost of production, i.e. there is common knowledge of (c_{1t}, c_{2t}, c_t). that each firm has equal access to the new technology after paying the price $K(c_t)$ to acquire the current state-of-the-art technology with marginal cost of production c_t. Each firm i, $\in \{1, 2\}$ also incurs idiosyncratic "disruption costs" $\varepsilon_t^i = (\varepsilon_{0t}^i, \varepsilon_{1t}^i)$ associated with each of the choices of not to invest (ε_{0t}^i) and investing (ε_{1t}^i). These shocks are private information to each firm i.

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 $\eta\varepsilon_t^1$ be the idiosyncratic disruption costs involved in acquiring the state-of-the-art production technology for firm 1, and let $\eta\varepsilon_t^2$ be the corresponding costs for firm 2, where η is a scaling parameter.

For tractability, we assume that it is common knowledge among the two firms that $\{\varepsilon_t^1\}$ and $\{\varepsilon_t^2\}$ are independent *IID* Type 1 bivariate extreme value processes with common scale parameter $\eta \geq 0$. Firm i observes its current and past idiosyncratic investment shocks $\{\varepsilon_t^i\}$, but does does not observe its future shocks or it's opponent's past, present or future idiosyncratic investment cost shocks. After each firm independently and simultaneously decides whether or not to invest in the latest technology, the firms then make a decision of which prices to sell their products at, where production is done in period t with their existing production machinery.

The one period time-to-build assumption implies that even if both firms invest in new production machinery at time t, their marginal cost of production in period t are c_{1t} and c_{2t}, respectively, since they have to wait until period $t+1$ for the new machinery to be installed, and must produce in period t using their old machines that they already have in place. However in period $t+1$ we have $c_{1,t+1} = c_t$ and $c_{2,t+1} = c_t$, since in period $t+1$ the new plants the firms purchased in period t have now become operational. Notice that these new plants reflect the state-of-the-art production cost c_t from period t when they ordered the new machinery. Meanwhile further technological progress could have occurred that drives down c_{t+1} to a value even lower than c_t. That is, continuous technological progress implies the possibility that the new plant(s) may already be out of date by the time they come online.

2.3 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 by now 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. In a simultaneous move formulation of the game, the state is (c_{1t}, c_{2t}, c_t). If 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.

We are interested in exploring how slight variations in the sequencing of investment decisions by the two firms affect the set of MPE. Therefore we introduce a fourth state variable m_t that can assume three possible values, $m_t \in \{0,1,2\}$. When $m_t = 0$ the two firms make their investment decisions simultaneously at time t. When $m_t = 1$ firm 1 moves first and makes its investment decision before firm 2 invests. In this case, it is common knowledge on the part of the two firms that firm 1 invests first and firm 2 invests second, and can condition its investment decision on the investment decision by firm 1. We assume there is no delay in this sequencing of investment decisions: both occur at time t. However once investment decisions are made, there is still a one period "time to build" lag before the new plant becomes operational.

Symmetrically to the case $m_t = 1$, the final case, $m_t = 2$, denotes the case where firm 2 makes its investment decision first and firm 1 makes its investment decision after observing the investment decision by firm 2. We assume that regardless of the value of m_t the firms' *pricing decisions* are made simultaneously. Since we assume that consumer purchase decisions are static and there are no switching costs or other features that could lead to future consequences from current pricing decisions, we assume that in each period t the firms' prices constitute a Nash equilibrium of the simultaneous move game where each firm chooses its price to maximize its profits at t given the firms' marginal costs of production (c_{1t}, c_{2t}), and taking into account the price of the other firm as given in equation (1) above.

Definition: *A Markov perfect equilibrium to the duopoly investment and pricing game consists of a pair of strategies $(\iota_i^{\eta}(c_1,c_2,c,m), p_i^{\sigma}(c_1,c_2))$, $i \in \{1,2\}$ where $\iota_i^{\eta}(c_1,c_2,c,m) \in \{0,1\}$ is firm i's investment decision and $p_i^{\sigma}(c_1,c_2)$ is firm i's pricing decision. The pricing decision is required to be a Nash equilibrium solution the the firm's single period pricing/profit maximization decision given in equation (1) and the investment decision $\iota_i^{\eta}(c_1,c_2,c,m)$ must maximize the expected discounted value of firm i's present and future profit stream taking into account then investment and pricing strategies of its opponent.*

In our formulation of the duopoly game, we allow the state variable m_t to evolve according to an exogenous Markov chain with transition probability $f(m_{t+1}|m_t)$. This allows us to solve versions of the model where the two firms always move simultaneously ($f(0|m_t)=1$ for all t), or where firm 1 or 2 always move first ($f(1|m_t)=1$ or $f(2|m_t)=1$ for all t), or various types of alternating move games, where the firms may sometimes move simultaneously, or one or the other of the firms may move first depending on the value of m_t. One goal of our analysis is to show how the equilibria depend on various assumptions about the timing of the firms' moves.

The σ parameter entering the equilibrium pricing decision rules p_i^{σ} in this definition is the scale parameter of the extreme value distributed person-specific heterogeneity discussed in section 2.1 above. As $\sigma \downarrow 0$, the equilibrium pricing rule converges to the usual static Bertrand-Nash equilibrium solution, $p_i^0(c_1,c_2) = p(c_1,c_2) = \max[c_1,c_2]$. As we discussed in section 2.2, the η parameter in the notation for the equilibrium investment decision rule $\iota_i^{\eta}(c_1,c_2,c,m)$ denotes the scale parameter of extreme value distributed shocks that affect the investment decisions of the two firms. These shocks are private information to each firm, and thus, when $\eta > 0$ the MPE is also a *Bayesian Nash equilibrium* to a dynamic game of incomplete information. As $\eta \downarrow 0$, the investment decision rules will converge to decision rules $\iota_i^0(c_1,c_2,c,m)$ which constitute an MPE for a game of complete information.

To derive the function equations characterizing the Markov-perfect equilibrium, we now drop the time subscripts. We will be focusing initially on a symmetric investment situation where each firm faces the same cost $K(c)$ of investment. However it is straightforward to modify the problem to allow one of the firms to have an *investment cost advantage.* In 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 $c \geq 0$.

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, the marginal cost of production using the current best technology is c, and the ordering of the investment decisions is common knowledge by the two firms and given by m (simultaneous moves if $m = 0$, firm 1 moves first if $m = 1$, and firm 2 moves first if $m = 2$). 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, m), and similarly for firm 2, and these value functions and decision rules are *anonymous* (also called *exchangeable*) in the sense that

$$V^1(c_1, c_2, c, m, \varepsilon_0, \varepsilon_1) = V^2(c_2, c_1, c, m, \varepsilon_0, \varepsilon_1), \tag{5}$$

where $V^1(c_1, c_2, c, m, \varepsilon_0, \varepsilon_1)$ 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, and V^2 is the corresponding value function for firm 2. It is important to note that in both functions V^1 and V^2, the first argument refers to firm 1's marginal cost of production of firm 1, and the second argument to the marginal cost of firm 2.

What the symmetry condition in equation (5) says, is that the value function for the firms only depends on the values of the state variables, not on their identities or the arbitrary labels "firm 1" and "firm 2". Thus if firm 1 has cost of production c_1 and firm 2 has cost of production c_2, and if both firms were to have the same private cost/benefit values of investing/not investing of $(\varepsilon_0, \varepsilon_1)$, respectively, then the expected profits firm 1 would expect would be the same as what firm 2 would expect for the state vector $(c_2, c_1, c, m, \varepsilon_0, \varepsilon_1)$, where we switch the order of the first two arguments c_1 and c_2. Conversely if firm 2 had marginal cost of production c_1 and firm 1 had marginal cost of production c_2, then firm 2's expected discounted profits in this state are the same as the discounted profits firm 1 could expect if these marginal costs were swapped (i.e. if firm 1 had marginal cost of production c_1 and firm 2 had marginal cost of production c_2).

Unfortunately, we will show below that there are interesting equilibria in the game for which the symmetry condition does *not* hold. In these equilibria, the nature of the equilibrium selection rules does confer distinct identities to the two firms, so their "labels" matter and the symmetry condition (5) does not hold. Instead, it is necessary to keep track of the separate value functions V^1 and V^2 in order to correctly compute the equilibria of the game. We will refer to these equilibria as *asymmetric equilibria* to distinguish them from *symmetric equilibria* where the symmetry condition (5) holds. We will show that many "interesting" equilibria of this model, including the two monopoly equilibria and various types of equilibria where there is leapfrogging, are asymmetric.

Now, assume that the cost/benefits from investing or not investing $(\varepsilon^i_{0t}, \varepsilon^i_{1t})$ for each firm $i = 1,2$ are private information to each firm and are *IID* over time and are also *IID* across the two firms, and both firms have common knowledge that these shocks have an extreme value distribution with a common scale parameter η as noted above. Then we can show that the value functions V^i, $i = 1,2$ take the form

$$V^i(c_1,c_2,c,m,\varepsilon^i_0,\varepsilon^i_1) = \max[v^i_0(c_1,c_2,c,m)+\eta\varepsilon^i_0, v^i_1(c_1,c_2,c,m)+\eta\varepsilon^i_1] \tag{6}$$

where $v^i_0(c_1,c_2,c,m)$ is the expected value to firm i if it does not invest in the latest technology, and $v^i_1(c_1,c_2,c,m)$ is the expected value to firm i if it does invest.

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. Note that the static Bertrand-Nash price equilibrium *is* symmetric. That is, firm 2's single period profits when marginal costs of firms 1 and 2 are (c_1,c_2), respectively, is given by $r^2(c_1,c_2) = r^1(c_2,c_1)$. That is, the profits firm 2 can earn in state (c_1,c_2) are the same as what firm 1 can earn in state (c_2,c_1). However in order to maintain notational consistency, we will let $r^i(c_1,c_2)$ denote the profits earned by firm i when the marginal costs of production of firms 1 and 2 are (c_1,c_2), respectively. In the limiting "pure Bertrand" case (i.e. where consumer demand is infinitely elastic) we have

$$r^1(c_1,c_2) = \begin{cases} 0 & \text{if } c_1 \geq c_2 \\ \max[c_1,c_2]-c_1 & \text{otherwise} \end{cases} \tag{7}$$

It is easy to verify directly in this case that the symmetry condition holds for the payoff functions r^1 and r^2, and also it is clear that when $c_1 = c_2$ we have $r^1(c_1,c_2) = r^2(c_1,c_2) = 0$.

The formula for firm i's expected profits associated with not investing, $v^i_0(c_1,c_2,c,m)$, given in equation (6) above (where the 0 subscript denotes the decision not to invest) is given by

$$v^i_0(c_1,c_2,c,m) = r^i(c_1,c_2) + \beta EV^i(c_1,c_2,c,m,0), \tag{8}$$

where $EV^i(c_1,c_2,c,m,0)$ is firm i's conditional expectation of its next period value function V^i given that it does not invest this period, $d_i = 0$, conditional on (c_1, c_2, c, m) (i.e. the last argument, 0, in $EV^i(c_1, c_2, c, m, 0)$ denotes firm i's decision not to invest).

This version of the conditional expectation which conditions only on firm i's investment decision and not on the decision d_{-i} of firm i's rival, is relevant for the case where the two firms move simultaneously, $m = 0$, or in either of the sequential move cases ($m = 1$ or $m = 2$) where firm i is the *first* mover.

However if the firm is the *second mover* (i.e. either firm 2 when $m = 1$ or firm 1 when $m = 2$), then the firm can condition its expectation of V^i both its own decision d_i and the decision of its opponent d_{-i} and we write the expectation in this case as $EV^i(c_1, c_2, c, m, d_i, d_{-i})$. Note that due to our assumption the $\{\varepsilon^i_{0t}, \varepsilon^i_{1t}\}$ are independent (across firms) *IID* processes, $EV^i(c_1, c_2, c, m, 0)$ is also the conditional expectation of V^i given $(c_1, c_2, c, m, \varepsilon^i_{0t}, \varepsilon^i_{1t})$.

The corresponding formula for the conditional expectation of discounted present and future profits given the decision to invest is

$$v^i_1(c_1, c_2, c, m) = r^i(c_1, c_2) - K(c) + \beta EV^i(c_1, c_2, c, m, 1), \tag{9}$$

where $EV^i(c_1, c_2, c, m, 1)$ is firm i's conditional expectation of its next period value function V^i given that it invests, $d_i = 1$, conditional on (c_1, c_2, c, m) (and $(\varepsilon^i_{0t}, \varepsilon^i_{1t})$ as per the note above). Similar to the discussion above in cases where firm i is a second mover, the firm observes the decision of its opponent d_{-i} and therefore conditions both on its own decision $d_i = 1$ and d_{-i} when computing the expectation of V^i, $EV^i(c_1, c_2, c, m, 1, d_{-i})$.

To compute the conditional expectations $EV^i(c_1, c_2, c, m, d_i)$ (and $EV^i(c_1, c_2, c, m, d_i, d_{-i})$ when firm i is a second mover) 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 $(\varepsilon^i_0, \varepsilon^i_1)$ is given by the standard "log-sum" formula when these shocks have the Type-III extreme value distribution. Thus, after taking expectations over $(\varepsilon^i_0, \varepsilon^i_1)$ in the equation for V^i in (6) above, we have

$$\int_{\varepsilon^i_0} \int_{\varepsilon^i_1} V^i(c_1, c_2, c, m, \varepsilon^i_0, \varepsilon^i_1) q(\varepsilon^i_0) q(\varepsilon^i_1) d\varepsilon^i_1 d\varepsilon^i_0 = \eta \log\left[\exp\{v^i_0(c_1, c_2, c, m)/\eta\} + \exp\{v^i_1(c_1, c_2, c, m)/\eta\}\right]. \tag{10}$$

The log-sum formula provides a closed-form expression for the conditional expectation of the value functions $V^i(c_1, c_2, c, m, \varepsilon^i_0, \varepsilon^i_1)$ for each firm i, where V^i is the maximum of the value of not investing or investing as we can see from equation (6) 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 (10) with respect to the next-period values of $(\varepsilon_0^i, \varepsilon_1^i)$. However we do need to compute the two functions $v_0^i(c_1,c_2,c,m)$ and $v_1^i(c_1,c_2,c,m)$ for both firms $i = 1,2$. We will describe one algorithm for doing this below.

To simplify notation, we let $\phi(v_0^i(c_1,c_2,c,m), v_1^i(c_1,c_2,c,m))$ be the log-sum formula given above in equation (10), that is define ϕ as

$$\phi(v_0^i(c_1,c_2,c,m), v_1^i(c_1,c_2,c,m)) \equiv \eta \log\left[\exp\{v_0^i(c_1,c_2,c,m)/\eta\} + \exp\{v_1^i(c_1,c_2,c,m)/\eta\}\right]. \qquad (11)$$

The ϕ function is also sometimes called the "smoothed max" function since we have

$$\lim_{\eta \to 0} \phi(v_0, v_1) = \max[v_0, v_1]. \qquad (12)$$

Further, for any $\eta > 0$ we have $\phi(v_0, v_1) > \max[v_0, v_1]$.

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 $(\varepsilon_0^1, \varepsilon_1^1)$ 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

$$P^1(c_1,c_2,c,m) = \frac{\exp\{v_1^1(c_1,c_2,c,m)/\eta\}}{\exp\{v_1^1(c_1,c_2,c,m)/\eta\} + \exp\{v_0^1(c_1,c_2,c,m)/\eta\}} \qquad (13)$$

Firm 2's belief of firm 1's probability of not investing, $P^1(c_1,c_2,c,m)$ 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 given by

$$P^2(c_1,c_2,c,m) = \frac{\exp\{v_1^2(c_1,c_2,c,m)/\eta\}}{\exp\{v_1^2(c_1,c_2,c,m)/\eta\} + \exp\{v_0^2(c_1,c_2,c,m)/\eta\}} \qquad (14)$$

If the symmetry condition holds, then we have $P^2(c_1,c_2,c,m) = P^1(c_2,c_1,c,m)$.

Now consider the firms' beliefs in the case where firm 1 makes its investment decision first, $m = 1$. In this case firm 2 can condition its investment decision on firm 1's investment choice. Let P_1^2 denote firm 1's belief about the probability firm 2 will invest if firm 1 chooses to invest, and let P_0^2 denote firm 1's belief of firm 2's investment probability if it chooses not to invest. Similarly, let $v_{ij}^2(c_1,c_2,c,m)$ denote firm 2's expected value from taking investment decision $i \in \{0,1\}$ in state (c_1,c_2,c,m) given that firm 1 moved first and took investment decision $j \in \{0,1\}$. Then $P_j^2(c_1,c_2,c,m)$, firm 1's belief of firm 2's probability of investing given firm 1's investment decision j, is given by

$$P_j^2(c_1,c_2,c,m) = \frac{\exp\{v_{1j}^2(c_1,c_2,c,m)/\eta\}}{\exp\{v_{1j}^2(c_1,c_2,c,m)/\eta\} + \exp\{v_{0j}^2(c_1,c_2,c,m)/\eta\}}. \qquad (15)$$

The value functions $v_{ij}^2(c_1,c_2,c,m)$ are given by

$$\begin{aligned}
v_{10}^2(c_1,c_2,c,m) &= r^2(c_1,c_2) - K(c) + \beta EV^2(c_1,c_2,c,m,1,0) \\
v_{11}^2(c_1,c_2,c,m) &= r^2(c_1,c_2) - K(c) + \beta EV^2(c_1,c_2,c,m,1,1) \\
v_{00}^2(c_1,c_2,c,m) &= r^2(c_1,c_2) + \beta EV^2(c_1,c_2,c,m,0,0) \\
v_{01}^2(c_1,c_2,c,m) &= r^2(c_1,c_2) + \beta EV^2(c_1,c_2,c,m,0,1).
\end{aligned} \tag{16}$$

Firm 1 makes its investment decision, $j \in \{0,1\}$ taking the reaction by firm 2 to its decision, embodied by the probabilities P_j^2, as given. The values to firm 1 for not investing and investing are, respectively,

$$\begin{aligned}
v_0^1(c_1,c_2,c,m) &= r^1(c_1,c_2) + \beta P_0^2(c_1,c_2,c,m) EV^1(c_1,c_2,c,m,0,1) + \\
&\quad (1 - P_0^2(c_1,c_2,c,m)) EV^1(c_1,c_2,c,m,0,0) \\
v_1^1(c_1,c_2,c,m) &= r^1(c_1,c_2) - K(c) + \beta P_1^2(c_1,c_2,c,m) EV^1(c_1,c_2,c,m,1,1) + \\
&\quad (1 - P_1^2(c,c_2,c,m)) EV^1(c_1,c_2,c,m,1,0).
\end{aligned} \tag{17}$$

These values enter the logit probability for firm 1's investment decision in equation (13).

The value functions in equation (17) are valid when firm 1 is the first mover, i.e. in state $m = 1$. The equations need to be slightly modified in case firms 1 and 2 move simultaneously, $m = 0$. In that case firm 1's beliefs about firm 2's probability of investing are no longer conditional on firm 1's investment decision d_1. Instead, firm 1 has beliefs about the probability firm 2 will invest in the simultaneous move game of the form $P^2(c_1,c_2,c,m)$ that depend on the current state of the game (c_1,c_2,c,m) but not on firm 1's investment decision d_1. Notationally, all that is involved to reflect this is to replace the probabilities P_0^2 and P_1^2 that condition on firm 1's investment decision $d_1 = 0$ or $d_1 = 1$ via the values in the subscripts by the single probability P^2 that has no subscript, reflecting that firm 1 believes that its *realized* investment decision d_1 cannot affect its beliefs about its opponent's probability of investing if it is common knowledge that the two firms are moving simultaneously ($m = 0$).

In order to compute the conditional expectations, it is necessary to show how these expectations depend on the order of decisions in the next period, which we denote as m'. Below, we provide formulas for the conditional expectations $EV^i(c_1,c_2,c,m,m',d_i,d_{-i})$, and we use the transition probability $f(m'|m)$ to write

$$EV^i(c_1,c_2,c,m,d_i,d_{-i}) = \sum_{m'} EV^i(c_1,c_2,c,m,m',d_i,d_{-i}) f(m'|m). \tag{18}$$

Further, it is not difficult to show that once we condition on the mover order next period, m', the expectation of future discounted profits does not depend on mover order last period, m, we we have

$$EV^i(c_1,c_2,c,m,m',d_i,d_{-i}) = EV^i(c_1,c_2,c,m',d_i,d_{-i}). \tag{19}$$

Now consider the case where $m'=0$ (simultaneous moves by the firms) or the case where firm i is the first mover. Specifically, consider the formulas for firm 1. We have

$$\begin{aligned}
EV^1(c_1,c_2,c,m',0,0) &= \int_0^c \phi(v_0^1(c_1,c_2,c',m'),v_1^1(c_1,c_2,c',m'))\pi(c'|c)dc' \\
EV^1(c_1,c_2,c,m',0,1) &= \int_0^c \phi(v_0^1(c_1,c,c',m'),v_1^1(c_1,c,c',m'))\pi(c'|c)dc' \\
EV^1(c_1,c_2,c,m',1,0) &= \int_0^c \phi(v_0^1(c,c_2,c',m'),v_1^1(c,c_2,c',m'))\pi(c'|c)dc' \\
EV^1(c_1,c_2,c,m',1,1) &= \int_0^c \phi(v_0^1(c,c,c',m'),v_1^1(c,c,c',m'))\pi(c'|c)dc'.
\end{aligned} \tag{20}$$

Equation (20) is valid either in the simultaneous move case, $m'=0$, or when firm 1 moves first, $m'=1$. In the case $m'=2$ then firm 2 moves first in the next period and firm 1 moves second. In this case the formulas for the expected values of discounted profits are different since firm 1 has to anticipate whether firm 2 will invest or not when it makes the first move next period.

$$\begin{aligned}
EV^1(c_1,c_2,c,m',0,0) &= \int_0^c P^2(c_1,c_2,c',m')\phi(v_{01}^1(c_1,c_2,c',m'),v_{11}^1(c_1,c_2,c',m'))+ \\
&\quad (1-P^2(c_1,c_2,c',m'))\phi(v_{00}^1(c_1,c_2,c',m'),v_{10}^1(c_1,c_2,c',m'))\pi(c'|c)dc' \\
EV^1(c_1,c_2,c,m',0,1) &= \int_0^c P^2(c_1,c,c',m')\phi(v_{01}^1(c_1,c,c',m'),v_{11}^1(c_1,c,c',m'))+ \\
&\quad (1-P^2(c_1,c,c',m'))\phi(v_{00}^1(c_1,c,c',m'),v_{10}^1(c_1,c,c',m'))\pi(c'|c)dc' \\
EV^1(c_1,c_2,c,m',1,0) &= \int_0^c P^2(c,c_2,c',m')\phi(v_{01}^1(c,c_2,c',m'),v_{11}^1(c,c_2,c',m'))+ \\
&\quad (1-P^2(c,c_2,c',m'))\phi(v_{00}^1(c,c_2,c',m'),v_{10}^1(c,c_2,c',m'))\pi(c'|c)dc' \\
EV^1(c_1,c_2,c,m',1,1) &= \int_0^c P^2(c,c,c',m')\phi(v_{01}^1(c,c,c',m'),v_{11}^1(c,c,c',m'))+ \\
&\quad (1-P^2(c,c,c',m'))\phi(v_{00}^1(c,c,c',m'),v_{10}^1(c_1,c,c',m'))\pi(c'|c)dc'.
\end{aligned} \tag{21}$$

Consider a version of the duopoly game where the firms move simultaneously in every period, $f(0|m)=1$, $m\in\{0,1,2\}$. Then we can simplify notation by dropping the m argument in the value functions and investment probabilities and write the functional equations for the simultaneous move version of the dynamic duopoly problem as follows

$$\begin{aligned}
v_0^1(c_1,c_2,c) &= r^1(c_1,c_2)+\beta\int_0^c \left[P^2(c_1,c_2,c)\phi(v_0^1(c_1,c,c'),v_1^1(c_1,c,c')) + \right. \\
&\quad \left. (1-P^2(c_1,c_2,c))\phi(v_0^1(c_1,c_2,c'),v_1^1(c_1,c_2,c'))\right]\pi(dc'|c).
\end{aligned}$$

$$\begin{aligned} v_1^1(c_1,c_2,c) &= r^1(c_1,c_2)-K(c)+\beta\int_0^c \big[P^2(c_1,c_2,c)\phi(v_0^1(c,c,c'),v_1^1(c,c,c'))+ \\ &\quad (1-P^2(c_1,c_2,c))\phi(v_0^1(c,c_2,c'),v_1^1(c,c_2,c'))\big]\,\pi(dc'|c). \end{aligned} \tag{22}$$

$$\begin{aligned} v_0^2(c_1,c_2,c) &= r^2(c_1,c_2)+\beta\int_0^c \big[P^1(c_1,c_2,c)\phi(v_0^2(c,c_2,c'),v_1^2(c,c_2,c'))+ \\ &\quad (1-P^1(c_1,c_2,c))\phi(v_0^2(c_1,c_2,c'),v_1^2(c_1,c_2,c'))\big]\,\pi(dc'|c). \\ v_1^2(c_1,c_2,c) &= r^2(c_1,c_2)-K(c)+\beta\int_0^c \big[P^1(c_1,c_2,c)\phi(v_0^2(c,c,c'),v_1^2(c,c,c'))+ \\ &\quad (1-P^1(c_1,c_2,c))\phi(v_0^2(c_1,c,c'),v_1^2(c_1,c,c'))\big]\,\pi(dc'|c). \end{aligned} \tag{23}$$

These are the functional equations that need to be solved to compute a Markov-perfect equilibrium to the simultaneous move version of dynamic duopoly investment problem.

Now consider the special case where firm 1 always moves first, $f(1|m)=1$, $m\in\{0,1,2\}$. In this case the functional equations become (once again dropping the *m* argument in the value functions and choice probabilities to simplify the notation)

$$\begin{aligned} v_0^1(c_1,c_2,c) &= r^1(c_1,c_2)+\beta\int_0^c \big[P_0^2(c_1,c_2,c)\phi(v_0^1(c_1,c,c'),v_1^1(c_1,c,c'))+ \\ &\quad (1-P_0^2(c_1,c_2,c))\phi(v_0^1(c_1,c_2,c'),v_1^1(c_1,c_2,c'))\big]\,\pi(dc'|c). \\ v_1^1(c_1,c_2,c) &= r^1(c_1,c_2)-K(c)+\beta\int_0^c \big[P_1^2(c_1,c_2,c)\phi(v_0^1(c,c,c'),v_1^1(c,c,c'))+ \\ &\quad (1-P_1^2(c_1,c_2,c))\phi(v_0^1(c,c_2,c'),v_1^1(c,c_2,c'))\big]\,\pi(dc'|c). \end{aligned} \tag{24}$$

Notice that the main difference between the functional equations for (v_0^1,v_1^1) in the simultaneous move case in equation (22) and the case where firm 1 always moves first in equation (24) is that firm 2's investment probability P^2 in the simultaneous move case is replaced by the two investment probabilities P_0^2 and P_1^2 reflecting that when firm 1 moves first, firm 2 is able to observe firm 1's *realized* investment decision and condition on it in making its own investment decision. Firm 1 therefore needs to take this into account when it makes its investment decisions.

When it comes to firm 2 there are now *four* rather than only 2 functional equations to be solved when firm 2 is always the second mover, since firm 2 has the extra information on firm 1's realized investment decision d_1 to take into account when making its investment decisions. Adapting the recursion equations (16) above to remove the *m* argument and substitute the formulas for the EV^2 (adapting the corresponding equations for EV^1 when we assumed it is the second mover given in (21) above) we obtain

$$v_{00}^2(c_1,c_2,c) = r^2(c_1,c_2)+\beta\int_0^c P^1(c_1,c_2,c')\phi(v_{01}^2(c_1,c_2,c'),v_{11}^2(c_1,c_2,c'))+$$

$$
\begin{aligned}
& (1-P^1(c_1,c_2,c'))\phi(v^2_{00}(c_1,c_2,c'),v^2_{10}(c_1,c_2,c'))\pi(c'|c)dc' \\
v^2_{01}(c_1,c_2,c) \quad = \quad & r^2(c_1,c_2)+\beta\int_0^c P^1(c_1,c,c')\phi(v^2_{01}(c_1,c,c'),v^2_{11}(c_1,c,c'))+ \\
& (1-P^1(c_1,c,c'))\phi(v^2_{00}(c_1,c,c'),v^2_{10}(c_1,c,c')\pi(c'|c)dc' \\
v^2_{10}(c_1,c_2,c) \quad = \quad & r^2(c_1,c_2)-K(c)+\beta\int_0^c P^1(c,c_2,c')\phi(v^2_{01}(c,c_2,c'),v^2_{11}(c_1,c_2,c'))+ \\
& (1-P^1(c,c_2,c'))\phi(v^2_{00}(c,c_2,c'),v^2_{10}(c,c_2,c'))\pi(c'|c)dc' \\
v^2_{11}(c_1,c_2,c) \quad = \quad & r^2(c_1,c_2)-K(c)+\beta\int_0^c P^1(c,c,c')\phi(v^2_{01}(c,c,c'),v^2_{11}(c_1,c,c'))+ \\
& (1-P^1(c,c,c'))\phi(v^2_{00}(c,c,c'),v^2_{10}(c,c,c'))\pi(c'|c)dc'. \qquad (25)
\end{aligned}
$$

In the most general case where $f(m'|m)$ admits positive probability on any of the 3 possible move orderings, there will be 2 functional equations for firm 1's value function (v^1_0, v^1_1) analagous to (24) and 4 functional equations for firm 2's value functions $(v^2_{00}, v^2_{01}, v^2_{10}, v^2_{11})$ analagous to (25) when $m = 1$, and additional set of 6 functional equations (2 for firm 2 as the first mover and 4 for firm 1 as the second mover) when $m = 2$, and 4 additional functiional equations (2 each for firms 1 and 2) analagous to equations (22) and (23) when $m = 0$, or a system of 16 functional equations to be solved in the most general case where the designation of mover evolves randomly according to a general Markov transition probability matrix over time.

Although we believe there are analytic closed-form solutions to this system of functional equations in certain polar cases (e.g. when $\eta = 0$ and $\sigma = 0$, the "pure Bertrand case"), our approach so far has been to solve the set of functional equations (22) and (23) numerically, showing that its is possible using a recursive algorithm we call a *state space recursion algorithm* to find all possible solutions to the system of functional equations (22) and (23), and thus all MPE of the dynamic duopoly game.

Although the system is a pair of "Bellman equations" (one for firm 1 and one for firm 2) and a single firm Bellman equation typically has a unique solution, in this case the resemblance is only superficial. We will show below that the set of functional equations (22) and (23) are *not* contraction mappings due to the interdependence of the best response probabilities P^1_1 and P^2_1. When the contract property fails, far from having a unique solution, we show that there can be a continuum of different solutions to equations (22) and (23). The various solutions to these equations correspond to the set of possible equilibria of the dynamic duopoly game.

Another implication of the fact that equations (22) and (23) do not define the equilibrium values of the two firms as a fixed point to a contraction mapping is that the usual method of *successive approximations* (also known as backward induction) — is not guaranteed to converge. For example in the specialization of the game where the two firms make simultaneous investment decisions in every period ($f(0|m) = 1$, $m \in \{0,1,2\}$), we can represent the method of successive approximations as sequence of four functions generated by iterating on the functional equations (22) and (23). This results in a sequence $\{v_t\}$ where $v_t = \Gamma(v_{t-1})$ and $v_t = (v^1_{t0}, v^1_{t1}, v^2_{t0}, v^2_{1t})$ and $\Gamma : B^4(S) \to B^4(S)$ where $B^4(S)$ is the Banach space of 4-tuples of continuous, bounded real functions from $S \to R$ and $\Gamma = (\Gamma_1, \Gamma_2, \Gamma_3, \Gamma_4)$ is the 4-tuple of operators mapping $v_t \in B^(S)$ to $v_{t+1} \in B^4(S)$ via the 4 functional equations in (22) and (23) above.

That is, in the case of simultaneous investment decisions, we can write the equilibria as the solution to the four functional equations

$$\begin{aligned} v^1_0 &= \Gamma_1(v^1_0, v^1_1, v^2_0, v^2_1) \\ v^1_1 &= \Gamma_2(v^1_0, v^1_1, v^2_0, v^2_1) \\ v^2_0 &= \Gamma_3(v^1_0, v^1_1, v^2_0, v^2_1) \\ v^2_1 &= \Gamma_4(v^1_0, v^1_1, v^2_0, v^2_1), \end{aligned} \tag{26}$$

where each of the Γ_i operators depends implicitly on all four value functions $(v^1_0, v^1_1, v^2_0, v^2_1)$ due to the fact that the probabilities entering the MPE fixed point equations in (22) and (23) depend on the probabilities that firms 1 and 2 invest in the new technology, P^1_1 and P^2_1, and these probabilites depend in turn on (v^1_0, v^1_1) and (v^2_0, v^2_1), respectively as we can see in equations (13) and (14) above.

Unless the Γ_i are contrnction mappings or have some other type of general structture, there is generally no guarantee that the method of successive approximations will converge. Indeed, we have found in numerical experiments that it frequently doesn't converge. However it is not hard to show that *if* successive approximations does converge, it converges a fixed point of the functional equations (22) and (23), and thus to a particular equilibrium of the dynamic game.

We have found that the convergence of successive approximations is very dependent on the starting value v_0, and as we show in the next section, there are various versions of the Γ operator corresponding to different choices for the state-specific equilibria involved in evaluating possible values of the response probabilities $P^1_1(c_1, c_2, c)$ and $P^2_1(c_1, c_2, c)$ at different points in the state space $(c_1, c_2, c) \in S$.

Thus while it may not be immediately apparent from inspection of the system of functional equations (22) and (23) and the corresponding choice probabilities, (13) and (14), the solutions to this system will depend on the *equilibrium selection rule* that chooses one of several possible *state-specific equilibria* in the simultaneous move investment stage game at each possible state $(c_1, c_2, c) \in S$. In the next section we introduce a different recursive algorithm (similar to but different in important ways to standard backward induction) that enables us to compute and fully characterize *all* equilibria of this game, and thereby bound the possible set of payoffs to consumers and the two firms. The wide set of possible equilibria that can emerge from this simple model is reminiscent of the literature on the *Folk Theorem* in repeated games, where there are also theorems characterizing the set of possible equilibria and bounds on the set of equilibrium payoffs.

3 Solving the "End Game"

Under our assumptions the exogenously specified Markov process governing improvements in production technology has an absorbing state, which without loss of generality we assume equals the minimum possible production cost equal to $c = 0$. This will also turn out to be the absorbing state of the game, so that once costs of the firms reach zero, they can go no lower and we will show that unless costs of investment are sufficiently high, in every possible end game state one or both of the firms will invest to attain this zero cost marginal cost of production. Since there is no forgetting or physical or knowledge depreciation in our model that would ever cause costs to rise, it follows that once costs attain the zero cost absorbing state, they will remain at this value in all future periods. The game will have effectively ended at that point, even though we have forumlated as an infinite horizon problem.

We now proceed to analyze the equilibria of this endgame. We show that once this absorbing state is reached, once one of the firms has invested in the state-of-the-art, technology that allows it to produce at zero marginal cost, the firm will no longer have any incentive to undertake any further cost-reducing investments since costs are as low as they can go. If any further 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* and *forward looking* motivation for undertaking further investments. However there are states of the form $(c_1, c_2, 0)$ where $c_1 > 0$ and $c_2 > 0$ where there are multiple possible investment equilibria. We will show that the analysis of these equilibria is similar in many respects to the analysis of equilibria of coordination games. In this case the firms want to coordinate on "good equilibria" that avoids the chance of both of them investing simultaneously, thereby earning negative profits from their decision to invest. The various solutions to this coordination game constitute the first hint of our resolution of the "Bertrand investment paradox" that we discussed in the introduction.

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 this section by discussing the easiest cases first, showing how we derive the payoffs (which are value functions that are solutions to the functional equations (22) and (23) given in section 2) simultaneously with determining the equilibrium decision rules in the endgame.

3.1 The $(0,0,0)$ End Game

The simplest "end game" corresponds to the state $(0,0,0)$, i.e. when the zero cost absorbing has been reached and both firms have adopted this state-of-the-art production technology. In the absence of random *IID* shocks $(\varepsilon_0^i, \varepsilon_1^i)$ corresponding to investing or not investing, respectively, neither of the firms would have any further incentive to invest since we assume there is is no depreciation in their capital stock, and they have both already achieved the lowest possible state-of-the-art production technology.

In the absence of privately observed idiosyncratic shocks, $(\varepsilon_0^i, \varepsilon_1^i)$, $i = 1,2$ (i.e. when $\eta = 0$), the $(0,0,0)$ end game would simply reduce to an infinite repetition of the zero-price, zero-profit Bertrand equilibrium outcome. No further investment would occur. Thus if this state were ever reached via the equilibrium path, the Bertrand investment paradox will hold, but in a rather trivial sense. There is no point in investing any further once technology has attained the lowest possible marginal cost of production, $c = 0$ since in this absorbing state the investment cannot enable one of the firms to gain a production cost advantage over its opponent.

When there are idiosyncratic shocks affecting investment decisions, there may be some short term reason (e.g. a temporary investment tax credit) that would induce one or both of the firms to invest, but such investments would be purely idiosyncratic unpredictable events with no real strategic consequence to their opponent, since the opponent has already achieved the minimum cost of production and thus, there is no further possibility of leapfrogging its opponent. In this zero-cost absorbing state the equations for the value functions (v_0^i, v_1^i) can be solved "almost" analytically.

$$\begin{aligned} v_0^i(0,0,0) &= r^i(0,0) + \beta P_1^{\sim i}(0,0,0)\phi(v_0^i(0,0,0), v_1^i(0,0,0)) \\ &+ \beta[1 - P_1^{\sim i}(0,0,0)]\phi(v_0^i(0,0,0), v_1^i(0,0,0)) \\ &= r^i(0,0) + \beta\phi(v_0^i(0,0,0), v_1^i(0,0,0)) \end{aligned} \tag{27}$$

where $P_1^{\sim i}(0,0,0)$ is a shorthand for firm i's opponent's probability of investing,

$$P_1^{\sim i}(0,0,0) = \frac{\exp\{v_1^{\sim i}(0,0,0)/\eta\}}{\exp\{v_0^{\sim i}(0,0,0)/\eta\} + \exp\{v_1^{\sim i}(0,0,0)/\eta\}} \tag{28}$$

Due to the fact that $(0,0,0)$ is an absorbing state, it can be easily shown that the value of investing, $v_1^i(0,0,0)$, is given by

$$v_1^i(0,0,0) = v_0^i(0,0,0) - K(0), \tag{29}$$

which implies via equation (28) that

$$P_1^{\sim i}(0,0,0) = \frac{\exp\{-K(0)/\eta\}}{1 + \exp\{-K(0)/\eta\}}. \tag{30}$$

Thus, as $\eta \to 0$, we have $P_1^{\sim i}(0,0,0) \to 0$ and $v_0^i(0,0,0) = r^i(0,0)/(1-\beta)$, and in the limiting case where the two firms are producing perfect substitutes, then $r^i(0,0) = 0$ and $v_0^i(0,0,0) = 0$. For positive values of η we have

$$v_0^i(0,0,0) = r^i(0,0) + \beta\phi(v_0^i(0,0,0), v_0^i(0,0,0) - K(0)). \tag{31}$$

This is a single non-linear equation for the single solution $v_0^i(0,0,0)$. The derivative of the right hand side of this equation with respect to $v_0^i(0,0,0)$ is 1 whereas the derivative of the right hand side is strictly less than 1, so if $r^i(0,0) > 0$, this equation has a unique solution $v_0^i(0,0,0)$ that can be computed by Newton's method.

Note that symmetry property for $r^i(0,0)$ implies that symmetry also holds in the $(0,0,0)$ end game: $v_0^1(0,0,0) = v_0^2(0,0,0)$ and $v_1^1(0,0,0) = v_1^2(0,0,0)$.

3.2 The $(c,0,0)$ End Game

The next simplest end game state is $(c,0,0)$. This is where firm 1 has not yet invested to attain the state-of-the-art zero cost plant, and instead has an older plant with a positive marginal cost of production c. However firm 2 has invested and has attained the lowest possible marginal cost of production 0. In the absence of stochastic shocks, in the limiting Bertrand case, it is clear that firm 1 would not have any incentive to invest since the investment would not allow it to leapfrog its opponent, but only to match its opponent's marginal cost of production. But doing this would unleash Bertrand price competition and zero profits for both firms. Therefore for any positive cost of investment $K(0)$ firm 1 would choose not to invest, leaving firm 2 to have a permanent low cost leader position in the market and charge a price of $p = c$.

In the case with stochastic shocks, just as in the $(0,0,0)$ endgame analyzed above, there may be transitory shocks that would induce firm 1 to invest and thereby match the 0 marginal cost of production of its opponent. However this investment is driven only by stochastic *IID* shocks and not by any strategic considerations, given that once the firm invests, it will generally not be in much better situation than if it had not invested (that is, even though $r^1(0,0) > r^1(c,0)$, both of these will be close to zero and will approach zero as $\eta \downarrow 0$). In the general case where $\eta > 0$ we have

$$\begin{aligned} v_0^1(c,0,0) &= r^1(c,0) + \beta\phi(v_0^1(c,0,0), v_1^1(c,0,0)) \\ v_1^1(c,0,0) &= r^1(c,0) - K(0) + \beta\phi(v_0^1(0,0,0), v_1^1(0,0,0)). \end{aligned} \tag{32}$$

Note that the solution for $v_1^1(c,0,0)$ in equation (32) is determined from the solutions $(v_0^1(0,0,0), v_1^1(0,0,0))$ to the $(0,0,0)$ endgame in equations (31) and (29) above. Substituting the resulting solution for $v_1^1(c,0,0)$ into the first equation in (32) results in another nonlinear equation with a single unique solution $v_0^1(c,0,0)$ that can be computed by Newton's method. Note that, as we show below, the probability that firm 2 invests in this case, $P_1^2(c,0,0)$ is given by

$$P_1^2(c,0,0) = \frac{\exp\{-K(0)/\eta\}}{1+\exp\{-K(0)/\eta\}} \tag{33}$$

since firm 2 has achieved the lowest possible cost of production and its decisions about investment are governed by the same idiosyncratic temporary shocks, and result in the same formula for the probability of investment as we derived above in equation (30) for the $(0,0,0)$ endgame.

It is not hard to see that the symmetry condition holds in the $(c,0,0)$ end game as well: $v_0^2(c,0,0) = v_0^1(0,c,0)$, and $v_1^2(c,0,0) = v_1^1(0,c,0)$, where the solutions for the latter functions are presented below.

3.3 The $(0,c,0)$ End Game

In this end game, firm 1 has achieved the lowest possible cost of production $c = 0$ but firm 2 hasn't yet. Its marginal cost of production is $c > 0$. Clearly firm 1 has no further incentive to invest since it has achieved the lowest possible cost of production. However in the presence of random cost shocks (i.e. in the case where $\eta > 0$), firm 1 will invest if there are idiosyncratic shocks that constitute unpredictable short term benefits from investing that outweigh the cost of investment $K(0)$. But since this investment confers no long term strategic advantage in this case, the equations for firm 1's values of not investing and investing, respectively, differ only by the cost of investment $K(0)$. That is,

$$v_1^1(0,c,0) = v_0^1(0,c,0) - K(0). \tag{34}$$

The equation for $v_0^1(0,c,0)$ is more complicated however, due to the chance that firm 2 might invest, $P_1^2(0,c,0)$. We have

$$\begin{aligned} v_0^1(0,c,0) = r^1(0,c) &+ \beta P_1^2(0,c,0)\phi(v_0^1(0,0,0), v_0^1(0,0,0) - K(0)) \\ &+ \beta[1-P_1^2(0,c,0)]\phi(v_0^1(0,c,0), v_0^1(0,c,0) - K(0)). \end{aligned} \tag{35}$$

The probability that firm 2 will invest, $P_1^2(0,c,0)$ is given by

$$\begin{aligned} P_1^2(0,c,0) &= \frac{\exp\{v_1^2(0,c,0)/\eta\}}{\exp\{v_1^2(0,c,0)/\eta\} + \exp\{v_0^2(0,c,0)/\eta\}} \\ &= \frac{\exp\{v_1^1(c,0,0)/\eta\}}{\exp\{v_1^1(c,0,0)/\eta\} + \exp\{v_0^1(c,0,0)/\eta\}}, \end{aligned} \tag{36}$$

where we used the symmetry condition that $v_j^2(0,c,0) = v_j^1(c,0,0)$, $j = 0,1$. Using the solution for $v_0^1(c,0,0)$ and $v_1^1(c,0,0))$ in the $(c,0,0)$ end game in equation (32) above, these solutions can be substituted into equation (36) to obtain the probability that firm 2 invests, and then this probability can be substituted into equation (35) to obtain a unique solution for $v_0^1(0,c,0)$, and finally the value of investing $v_1^1(0,c,0)$ is given by equation (34).

Once again, it is not hard to see that the symmetry condition holds in the $(0,c,0)$ end game: $v_0^2(0,c,0) = v_0^1(c,0,0)$ and $v_1^2(0,c,0) = v_1^1(c,0,0)$.

3.4 The $(c_1,c_2,0)$ End Game

The final case to consider is the end game where both firms have positive marginal costs of production, c_1 and c_2, respectively. We will show that in this end game, asymmetric equilibrium solutions are possible. We begin by showing how to solve the equations for the values to firm 1 of not investing and investing, respectively, which reduce to

$$\begin{aligned} v_0^1(c_1,c_2,0) = r^1(c_1,c_2) &+ \beta P_1^2(c_1,c_2,0)\phi(v_0^1(c_1,0,0), v_1^1(c_1,0,0)) \\ &+ \beta[1-P_1^2(c_1,c_2,0)]\phi(v_0^1(c_1,c_2,0), v_1^1(c_1,c_2,0)) \\ v_1^1(c_1,c_2,0) = r^1(c_1,c_2) - K(0) &+ \beta P_1^2(c_1,c_2,0)\phi(v_0^1(0,0,0), v_1^1(0,0,0)) \\ &+ \beta[1-P_1^2(c_1,c_2,0)]\phi(v_0^1(0,c_2,0), v_1^1(0,c_2,0)). \end{aligned} \tag{37}$$

Given the equation for $v_1^1(c_1,c_2,0)$ in equation (37) depends on known quantities on the right hand side (the values for v_0^1 and v_1^1 inside the ϕ functions can be computed in the $(0,0,0)$ and $(0,c,0)$ end games already covered above), we can treat $v_1^1(c_1,c_2,0)$ as a linear function of P_1^2 which is not yet "known" because it depends on $(v_0^2(c_1,c_2,0), v_1^2(c_1,c_2,0))$ via the identity:

$$P_1^2(c_1,c_2,0) = \frac{\exp\{v_1^2(c_1,c_2,0)/\eta\}}{\exp\{v_0^2(c_1,c_2,0)/\eta\} + \exp\{v_1^2(c_1,c_2,0)/\eta\}}. \tag{38}$$

We write $v_1^1(c_1,c_2,0,P_1^2)$ to remind the reader that it can be viewed as an implicit function of P_1^2: this is the value of v_1^1 that satisfies equation (37) for an arbitrary value of $P_1^2 \in [0,1]$. Substituting this into the equation for v_0^1, the top equation in (37), there will be a unique solution $v_0^1(c_1,c_2,0,P_1^2)$ for any $P_2 \in [0,1]$ since we have already solved for the values $(v_0^1(c_1,0,0), v_1^1(c_1,0,0))$ in the $(c,0,0)$ end game (see equation (32) above). Using these values, we can write firm 1's probability of investing $P_1^1(c_1,c_2,0)$ as

$$P_1^1(c_1,c_2,0,P_1^2) = \frac{\exp\{v_1^1(c_1,c_2,0,P_1^2)/\eta\}}{\exp\{v_0^1(c_1,c_2,0,P_1^2)/\eta\} + \exp\{v_1^1(c_1,c_2,0,P_1^2)/\eta\}}. \tag{39}$$

Now, the values for firm 2 $(v_0^2(c_1,c_2,0), v_1^2(c_1,c_2,0))$ that determine firm 2's probability of investing in equation (38) can also be written as functions of P_1^1 for any $P_1^1 \in [0,1]$. This implies that we can write firm 2's probability of investing as a function of its perceptions of firm 1's probability of investing, or as $P_1^2(c_1,c_2,0,P_1^1)$. Substituting this formula for P_1^2 into equation (39) we obtain the following fixed point equation for firm 1's probability of investing

$$P_1^1 = \frac{\exp\{v_1^1(c_1,c_2,0,P_1^2(c_1,c_2,0,P_1^1))/\eta\}}{\exp\{v_0^1(c_1,c_2,0,P_1^2(c_1,c_2,0,P_1^1))/\eta\} + \exp\{v_1^1(c_1,c_2,0,P_1^2(c_1,c_2,0,P_1^1))/\eta\}}. \tag{40}$$

3.5 End Game Equilibrium Solutions

By Brouwer's fixed point theorem, at least one solution to the fixed point equation (40) exists. Further, when $\eta > 0$, the objects entering this equation (i.e. the value functions $v_0^1(c_1,c_2,0,P_1^2)$, $v_1^1(c_1,c_2,0,P_1^2)$, $v_0^2(c_1,c_2,0,P_1^1)$, and $v_1^2(c_1,c_2,0,P_1^1)$ and the logit choice probability function P_1^2) are all C^∞ functions of P_1^2 and P_1^1, standard topological index theorems be applied to show that for almost all values of the underlying parameters, there will be an odd number of separated equilibria. Further, as $\eta \to 0$, the results of Harsanyi (1973) as extended to dynamic Markovian games by Doraszelski and Escobar (2009) show that η serves as a "homotopy parameter" and for sufficiently small η the set of equilibria to the "perturbed" game of incomplete information converge to the limiting game of complete information.

However rather than using the homotopy approach, we found we were able to directly solve for equilibria of the problem in the limiting pure Bertrand case where $\eta = 0$ and $\sigma = 0$. The case $\sigma = 0$ corresponds to the case where demand is perfectly elastic and all consumers buy from the firm with the lower price, and the case $\eta = 0$ corresponds to the situation where there are no random shocks affecting the returns to investing or not investing in the state-of-the-art production technology.

In the remainder of sections 3 and 4, we will focus our analysis on this limiting version of the model where there are no random unobservable shocks. In this case, limiting game is one of complete information between the two firms. In subsequent versions of this paper we will also solve and characterize the solutions to the incomplete information versions of the model and game. The resuts of Harsanyi and Doraszelski and Escobar show that for sufficiently small σ and η the set of equilibria to the incomplete-information "purified" games will be close to the set of equilibria we calculate below for the limit game (which will include mixed strategy equilibria that are not present in the incomplete-information, "purified" versions of this game).

We find that there are either 1 or 3 equilibria in the $(c_1, c_2, 0)$ end game, depending on the values of the parameters. The trivial equilibrium is a no-investment equilibrium that occurs when the cost of investment $K(0)$ is too high relative to the expected cost savings, and neither firm invests in this situation. However whenever $K(0)$ is below a critical threshold, there will be 3 equilibria to the end game: two pure strategy equilibria and an intermediate mixed strategy equilibrium.

It turns out that the investment game is isomorphic to a *coordination game.* 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. The mixed strategy equilibrium corresponds to the situation where firm 1 invests with probability π_1 and firm 2 invests with probability π_2. It is not hard to see that when $c_1 = c_2$ the game is fully symmetric and we have $\pi_1 = \pi_2$. However when $c_1 \neq c_2$, then the game is asymmetric and $\pi_1 \neq \pi_2$. In general, we can show that $c_1 > c_2$ implies that $\pi_1 > \pi_2$, i.e. *the cost-follower has a greater probability of investing and leapfrogging the low-cost leader.* Further, from the standpoint of the firms, the mixed strategy equilibrium is the "bad" equilibrium. In the symmetric case, $c_1 = c_2$, the mixed strategy 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.

Figure 1 plots the equilibria computed by plotting the best response function in equation (40) against the 45 degree line. In this example firm 1 is the low-cost leader with a substantially lower marginal cost of production than firm 2, $c_1 = 0.714$ vs. $c_2 = 2.14$. In the mixed strategy equilibrium, the low cost leader, firm 1, invests with probability 0.484, whereas the firm 2, the high cost follower, invests with probability

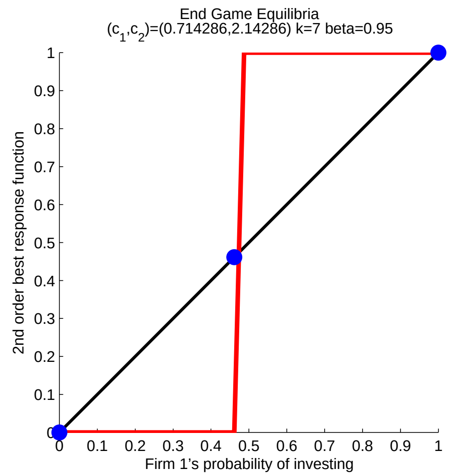

Figure 1 End Game Equilibria

0.82. Thus, the high cost follower has a significantly higher chance of leapfrogging its rival to attain the position of low cost leadership. This leadership is permanent (unless the firms happen to simultaneously invest) since by assumption, the production technology has reached the zero marginal cost absorbing state and there can be no further future improvements in production cost.

To get further insight into the potentially counterintuitive finding that the low cost leader has a *lower* probability of investing than the high cost follower, consider the payoff matrix for the simultaneous move game in investment decisions by firms 1 and 2 in state (c_1, c_2, c) below. This matrix is for the special case of the pure Bertrand case where the two firms produce perfect substitutes ($\sigma = 0$) and there are no unobserved shocks to the investment decisions ($\eta = 0$). Further, we show the payoff matrix in the asymmetric equilibrium case where $c_1 > c_2$, i.e. firm 2 is the low cost leader and firm 1 is the high cost follower.

		Firm 2	
		Invest	Don't Invest
Firm 1	Invest	$-K, c_1 - c_2 - K$	$\beta c_2/(1-\beta) - K, c_1 - c_2$
	Don't Invest	$0, c_1 - c_2 + \beta c_1/(1-\beta) - K$	$\beta V_1, c_1 - c_2 + \beta V_2$

Figure 1: End Game Payoff Matrix in state $(c_1, c_2, 0)$ with $c_1 > c_2$

To understand the formulas for the payoffs, it is easiest to start with the upper left hand corner of the payoff matrix when both firms decide to invest. In this case, since both firms attain the state-of-the-art marginal cost of $c = 0$, Bertrand competition insures that both firms earn zero profits following the investment, which costs K today. Since firm 2 is the low cost leader, it earns a profit of $c_1 - c_2$ in the current period, less its investment cost K, and zero profits thereafter, so its payoff is $c_1 - c_2 - K$. Firm 1 is the high cost follower so it earns zero profits in the current period, incurs the investment cost K, and earns zero profits thereafter, so its payoff is just $-K$.

In the upper right hand corner, we have the payoffs in the event firm 1 invests and firm 2 doesn't. In this case, once firm 1 has acquired the zero marginal cost state-of-the-art production technology, it can charge a price of c_2, the marginal cost of production of its rival. Once firm 1 has attained this position, firm 2 will clearly never have an incentive to try to invest in the future, so this investment will result in firm 1 having leapfrogged firm 2 to attain *permanent* low-cost leadership. Since the profits it will earn come with a one period delay (due to the time to install the new production machinery), firm 1's discounted profits after the investment cost are $\beta c_2/(1-\beta) - K$. Firm 2 will earn profits of $c_2 - c_1$ in the current period but zero profits thereafter.

In the lower left hand corner are the payoffs when firm 2 invests and firm 1 doesn't. In this case firm 2 invests and pre-empts firm 1 from undertaking any future investments and thereby improves its profitability and ensures that it has permanent low cost leadership. Its profits are given by $c_2 - c_1 + \beta c_1/(1-\beta) - K$, since firm 2 will be able to set a price equal to the marginal cost of its rival, c_1 and will have 0 marginal costs of production following its investment. However in the current period, while the new machinery is being installed and firm 2 is still producing with its existing machinery with marginal cost c_2, firm 2 will earn profits of $c_1 - c_2$ and will have to pay the investment cost K. Firm 1 will earn zero profits in the current period and 0 profits in every future period after firm 2 invests, so its payoff is 0.

The remaining case to consider is the lower right hand square of the payoff matrix, covering the case where neither firm invests. While it is tempting to write the payoffs as simply 0 for firm 1 (since it is the high cost follower and earns zero profits in the current period), and $c_1 - c_2$ for firm 2, this calculation of the payoffs would be incorrect since it ignores the value of the *future option to invest.* If both firms are playing a stationary, mixed strategy equilibrium, then in any future period where neither of the two firms have invested yet, the firms will continue to have the same strategy of investing with probability π_1 for firm 1 and π_2 for firm 2. Let $V_1(\pi_1, \pi_2)$ denote the expected present value of profits of firm 1 under this stationary mixed strategy equilibrium and $V_2(\pi_1, \pi_2)$ be the corresponding expected present value of profits for firm 2, *in the event that neither firm invests.* For firm 1 we have

$$V_1 = 0 + \beta V_1 \tag{41}$$

which implies that $V_1 = 0$. Since firm 1's expected payoffs are zero when it doesn't invest regardless of whether firm 2 invests or not, this implies that if firm 2 invests with probability π_2, the expected payoff to firm 1 from investing must also be 0, so we have

$$-K\pi_2 + (1-\pi_2)[\beta c_2/(1-\beta) - K] = 0, \tag{42}$$

or

$$\pi_2 = \frac{\beta c_2/(1-\beta) - K}{\beta c_2/(1-\beta)}. \tag{43}$$

From this formula we see that firm probability of investing is an increasing function of its own marginal cost c_2 and a decreasing function of the cost of investment, K, which seems eminently reasonable.

For firm 2 we have the following equation for V_2

$$V_2 = \pi_1(c_1 - c_2) + (1-\pi_1)(c_1 - c_2 + \beta V_2) \tag{44}$$

which implies that

$$V_2 = \frac{c_1 - c_2}{1 - \beta(1-\pi_1)}. \tag{45}$$

In order for firm 2 to be willing to pay a mixed investment strategy, its expected return from investing must also be equal to V_2, so we have

$$V_2 = \pi_1(c_1 - c_2 - K) + (1-\pi_1)(c_1 - c_2 + \beta c_1/(1-\beta) - K). \tag{46}$$

Combining equations (45) and (46) into a single equation for the unknown π_1, we can solve this quadratic equation, taking the positive root and ignoring the negative one.

Lemma 3.1. *Suppose* $\eta = 0$. *If* $c_1 > c_2 > 0$ *and* $K < \frac{\beta c_2}{1-\beta}$, *then in the unique mixed strategy equilibrium of the pure Bertrand dynamic investment and pricing game in state* $(c_1, c_2, 0)$ *we have* $\pi_1 > \pi_2$.

The proof of Lemma 3.1 is provided in the appendix. This result provides a first taste of the possibility of leapfrogging since the high cost leader has a higher probability of investing to become the (permanent) low cost leader with the state-of-the-art plant with zero marginal costs of production. However the co-ordination between the two firms in the mixed strategy equilibrium is far from desirable, since it implies a positive probability of inefficient simultaneous investment by the two firms. The question is, can more efficient coordination mechanisms be established as equilibria to the full game?

The desirable solutions, of course, are the two pure strategy equilibria since they involve the desirable coordination where only one firm invests and the other firm does not invest with probability 1. However which of the two equilibria would the firms choose? One pure strategy equilibrium gives all surplus to firm 1 and none to firm 2 and the opposite for the other pure strategy equilibrium. Since we assume the firms are risk neutral, a "fair" resolution to this problem would be to play a *correlated equilibrium* (Aumann, 1987) involving choosing one of the pure strategy equilibria with probability $1/2$. The expected payoff to firm 1 under this correlated equilibrium is $(1/2)(\beta c_2/(1-\beta)-K)$ and the expected payoff to firm 2 is $c_1 - c_2 + (1/2)(\beta c_1/(1-\beta)-K)$. These payoffs Pareto-dominate the expected payoffs to the two firms under the mixed strategy equilibrium which are 0 and $c_1 - c_2$, respectively.

Thus, if we view correlated equilibria in the end game as a *stochastic, state-specific equilibrium selection rule* then in the state $(c_1, c_2, 0) = (0.714, 2.14, 0)$, the set of possible equilibrium payoffs for firms 1 and 2 (with firm 1 payoffs on the x-axis and firm 2 payoffs on the y-axis) consists of the triangle with vertices at the points $(0, c_1 - c_2)$, $(0, c_1 - c_2 + \beta c_1/(1-\beta) - K)$ and $(\beta c_2/(1-\beta) - K, c_1 - c_2)$. Using this approach, we can form equilibrium payoff sets by taking convex combinations of the set of equilibria payoffs at all equilibria at each end game state, and these expected equilibrium payoffs become the continuation values for calculating equilibrium payoffs at higher cost states, i.e. at states of the form $(c_1, c_2, c) \in S$ where $c > 0$. We turn to this calculation in the next section.

4 Solving the Full Game

With the end game solutions in hand, we are now ready to proceed 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 an infinite horizon game.

We will assume initially *deterministic* equilibrium selection rules, i.e. a function that picks out one of the set of equilibria in each possible state of the game, (c_1, c_2, c). We now wish to analyze how different state-contingent equilibrium selection rules can support a wider range of equilibria in the full game, including a pattern of dynamic coordination between alternating pure strategy equilibria that we have referred to as leapfrogging.

Specifically, the classical notion of "leapfrogging" corresponds to the following class of *alternating investment equilibria* in the full game: *the high cost follower invests whenever the state-of-the-art production cost c falls sufficiently below the marginal cost of the low cost leader to justify the investment cost* $K(c)$*, otherwise no investment occurs.* Thus, in these equilibria, the two firms invest in "turn" with the investment, when it occurs, always being done by the high cost follower and never by the low cost leader. We will show that it is possible to "enforce" this equilibrium without relying on the types of "incredible threats" that have been discussed in the repeated game literature, and in the literature on tacit collusion in IO. In these tacitly collusive equilibria, various types of coordination between the two duopolists are supported by threats of a "price war" if one of the firms deviates from the collusive "agreement." One example might be if the low cost leader should ever become too "greedy" and invest when it is not "its turn", then firm 2 will respond by investing. One way of punishing such a deviation is to posit that in the event of such a deviation, the two firms would then engage in a "price war" or some other "bad equilibrium" that would be triggered whenever an investment that occurs "out of turn" moves the firms to the symmetric state $c_1 = c_2 = c$.

However in a Markovian game it is not possible to use this sort of punishment device since Markovian strategies do not carry enough memory to determine which of the two firms invests "out of turn". For example, consider an equilibrium where the firms play mixed investment strategies whenever $c_1 = c_2 = c$. Suppose the current state-of-the-art cost of production is c and consider two different situations: a) where $c_1 > c_2$ (so that firm 1 is the follower) and b) where $c_1 < c_2$ (so that firm 2 is the follower). If firm 1 "moves out of turn" by adopting the state-of-the-art production technology with cost c in case a) then the new state becomes (c, c_2, c). However if firm 1 does move "in turn" in case b) and invests in the state-of-the-art technology, the state is also (c, c_2, c). Thus, the ensuing behavior must be the same for both cases in a Markovian equilibrium. This implies that firm 1 cannot be punished for moving "out of turn" in case a) without also being punished for doing the correct thing and investing when it is its turn in case b).[2]

It follows that the traditional sorts of state dependent punishment strategies that have been used to support tacitly collusive equilibria in the repeated games literature cannot be used given the much more

[2]We thank Joseph E. Harrington, Jr. for pointing this out to us.

constrained level of history dependence allowed in a *Markovian* equilibrium. Nevertheless, we show that the set of equilibria in the full game is very large, in a way that is reminiscent of the Folk Theorem sort of "anything can happen" (i.e. any feasible payoff higher than the players' minimax payoffs) result from the repeated game literature.

In order to solve the full game, i.e. the pair of functional equations (22) and (23), it is helpful to rewrite them in the following way,

$$v_0^1(c_1,c_2,c) = r^1(c_1,c_2) + \beta\left[P_1^2(c_1,c_2,c)H^1(c_1,c,c) + (1-P_1^2(c_1,c_2,c))H^1(c_1,c_2,c)\right] \tag{47}$$

$$v_1^1(c_1,c_2,c) = r^1(c_1,c_2) - K(c) + \beta\left[P_1^2(c_1,c_2,c)H^1(c,c,c) + (1-P_1^2(c_1,c_2,c))H^1(c,c_2,c)\right] \tag{48}$$

where the function H^1 is given by

$$H^1(c_1,c_2,c) = p(c)\int_0^c \phi(v_0^1(c_1,c_2,c'),v_1^1(c_1,c_2,c'))f(c'|c)dc' + (1-p(c))\phi(v_0^1(c_1,c_2,c),v_1^1(c_1,c_2,c)), \tag{49}$$

where $p(c)$ is the probability that a cost-reducing innovation will occur, and $f(c'|c)$ is the conditional density of the new (lower) state-of-the-art marginal cost of production conditional on an innovation having occurred. We assume that the support of $f(c'|c)$ is in the interval $[0,c)$, as indicated also by the the interval of integration in equation (49).

For completeness, we present the corresponding equation for firm 2 below.

$$v_0^2(c_1,c_2,c) = r^1(c_2,c_1) + \beta\left[P_1^1(c_1,c_2,c)H^2(c,c_2,c) + (1-P_1^1(c_1,c_2,c))H^2(c_1,c_2,c)\right] \tag{50}$$

$$v_1^2(c_1,c_2,c) = r^1(c_2,c_1) - K(c) + \beta\left[P_1^1(c_1,c_2,c)H^2(c,c,c) + (1-P_1^1(c_1,c_2,c))H^2(c_1,c,c)\right] \tag{51}$$

where the function H^2 is given by

$$H^2(c_1,c_2,c) = p(c)\int_0^c \phi(v_0^2(c_1,c_2,c'),v_1^2(c_1,c_2,c'))f(c'|c)dc' + (1-p(c))\phi(v_0^2(c_1,c_2,c),v_1^2(c_1,c_2,c)), \tag{52}$$

If we set the arguments (c_1,c_2,c) to v_0 in equation (47) to (c,c,c), and similarly in equation (48) for v_1, we deduce that

$$v_1^1(c,c,c) = v_0^1(c,c,c) - K(c). \tag{53}$$

Clearly, if the firms have all invested and have in place the state-of-the-art production technology, there is no further incentive for either firm to invest. For the same reasons we have

$$v_1^1(c,c_2,c) = v_0^1(c,c_2,c) - K(c). \tag{54}$$

Similar to the strategy we used to solve the value functions (v_0^i, v_1^i) $i = 1,2$ in the end game, we can substitute equation (53) into equation (47) and use Newton's method to compute the unique fixed point $v_0^1(c,c,c)$. Similarly, we can solve for $v_0^1(c,c_2,c)$ by substituting equation (54) into equation (47) and solving the latter by Newton's method. Finally, to solve for $v_0^1(c_1,c_2,c)$ we note that using the solutions for $v_0^1(c,c,c)$ and $v_0^1(c,c_2,c)$ and equations (53) and (54) to obtain $v_1^1(c,c,c)$ and $v_1^1(c,c_2,c)$, we can compute $v_1^1(c_1,c_2,c)$ by substituting these values into equation (48). Then we substitute $v_1^1(c_1,c_2,c)$ into equation (47) and use Newton's method to compute $v_0^1(c_1,c_2,c)$.

Note that we can assume that the integral term in equation (49) is "known". This is because we can structure a recursive algorithm for solving the game by starting with the end game solution and recursively solving the equilibria and value functions for positive values c' that are less than the current value c that we are computing. Then for each $c' < c$, the value functions $v_0^i(c_1,c_2,c')$ and $v_1^i(c_1,c_2,c')$ will be "known" for all (c_1,c_2) in the rectangle $R(c') = \{(c_1,c_2)|c' \leq c_1 \leq \overline{c}, c' \leq c_2 \leq \overline{c}'\}$. This is how the equilibrium selection rule at "lower cost nodes" of the game tree (i.e. at states (c_1,c_2,c') with $c' < c$) affect the set of possible equilibria at each node (c_1,c_2,c).

More specifically, following the procedure we used to solve for equilibria in the $(c_1,c_2,0)$ end game, the set of all equilibria for the investment "stage game" at state (c_1,c_2,c) can be computed by finding all fixed points to the following "second order best response function" for firm 1:

$$P_1^1 = \frac{\exp\{v_1^1(c_1,c_2,c,P_1^2(c_1,c_2,c,P_1^1))/\eta\}}{\exp\{v_0^1(c_1,c_2,c,P_1^2(c_1,c_2,c,P_1^1))/\eta\} + \exp\{v_1^1(c_1,c_2,c,P_1^2(c_1,c_2,c,P_1^1))/\eta\}}. \tag{55}$$

Depending on the rule we choose to select among the possible equilibria in each state (c_1,c_2,c) (and similarly the selection rule for equilibria at all feasible points in the state space (c_1,c_2,c') with $c' < c$) we can construct a wide variety of equilibria for the overall game. The restriction is that any equilibrium selection rule must be such that the functional equations for equilibrium (see equations (47) and (48) above) are satisfied. The following steps are used to solve for the set of all equilibria at each state point (c_1,c_2,c) in the full Bertrand/investment game.

1. For each $P_1^1 \in [0,1]$ we compute the value functions $(v_0^2(c_2,c_1,c,P_1^1), v_1^2(c_2,c_1,c,P_1^1)$ representing *firm 2's* values of not investing and investing in state (c_1,c_2,c), respectively, by solving the system (50) and (51) for each $P_1^1 \in [0,1]$.

2. Compute firm 2's "best response", i.e. its probability of investing, $P_1^2(c_1,c_2,c,P_1^1)$, in response to its

perception of firm 1's probability of investing, P_1, via the equation

$$P_1^2(c_1,c_2,c,P_1^1) = \frac{\exp\{v_1^2(c_1,c_2,c,P_1^1)/\eta\}}{\exp\{v_0^2(c_1,c_2,c,P_1^1)/\eta\}+\exp\{v_1^2(c_1,c_2,c,P_1^1)/\eta\}}. \tag{56}$$

using the value functions for firm 2 computed in step 1 above.

3. Using firm 2's best response probability, P_1^2, calculate the value functions $v_0^1(c_1,c_2,c,P_1^2)$ and $v_1^1(c_1,c_2,c,P_1^2)$ representing *firm 1's* values of not investing and investing in state (c_1,c_2,c), respectively, by solving the system (47) and (48).

4. Using the values for firm 1, compute firm 1's probability of investing, *the second order best response function* for firm 1, and search for all fixed points in equation (55).

We refer to the recursive algorithm defined in steps [1] to [4] above as a *state space recursion* since it 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" not in time, but in terms of the movement of the states in the game from lower values to c to higher values. The state-space recursion computes an equilibrium of the dynamic game in a single iteration, and as such, there is no issue about "convergence" that is present in the traditional time-based backward induction/successive approximations approach to solving the pair of functional equations (22) and (23). As we noted above, there is no guarantee that a time-domain successive approximations algorithm will converge and we have found that it frequently *does not* converge. However the algorithm described in steps 1 to 4 above is *guaranteed to converge* because it is fundamentally *not an iterative algorithm.* Instead, it is a type of "backward induction algorithm" that works via the state space, using the state-of-the-art marginal cost of production c as the analog of the time index in a traditional backward induction (successive approximations) algorithm.

The advantage of the state space recursion algorithm is that it enables us to solve for all equilibria of the game. When there are a continuum of states, in principle there are a continuum of possible equilibria since we are free to choose different equilibria at each point in the state space (c_1,c_2,c) and there are a continuum of such points when we assume that is the polyhedron $S \subset R^3$. There are issues we have glossed over that restrict the choice of equilibria, such as guaranteeing that the resulting value functions $(v_0^i(c_1,c_2,c), v_1^i(c_1,c_2,c))$, $i \in \{1,2\}$ are integrable functions (otherwise we cannot guarantee that these value functions satisfy the functional equations defining the equilibrium, (22) and (23) in section 2 above).

However consider the specialization of the game were we restrict c to a finite grid consisting of n points in the interval $[0,c_0]$ (with 0 and c_0 included). Then the state space becomes a *lattice* that is a finite subset of S and the integrability issues disappear. In the end game there are $N_0 = (n-1)^2$ points where possible multiple equilibria can occur and as we noted in the previous section, there are 3 possible equilibria at each $(c_1, c_2, 0)$ end game state assuming that it is optimal to invest at all in that state (i.e. provided $K(0)$ is not too large). As a result there are 3^{N_0} possible equilibria in the end game alone. Then at the next grid point c above the $c = 0$ grid point there are $(n-2)^2$ interior grid points (c_1, c_2, c) where $c_1 > c$ and $c_2 > c$. If there are 3 possible equilibria at each of these states, then there are a total of $3^{N_0+N_c}$ possible equilibria. Continuing, we find that in the game as a whole there are N^* possible equilibria, where

$$N^* = 3^{n(n-1)(2n-1)/6}. \tag{57}$$

Thus, for $n = 50$ there are $N^* = 3^{285} = 9.5402 \times 10^{135}$ possible equilibria to this game! Morever this is the number when we restrict attention to *deterministic* equilibrium selection rules. When we allow stochastic equilibrium selection rules, then there are a continuum of possible equilibria to the game.

Obviously we cannot literally compute *all* equilibria when the number of possible equilibria expand at this exponential rate, but we can compute any arbitrarily selected equilibrium in the set of all possible equilibria using our state space recursion algorithm. It turns out that this capability if enough for us to be able to deduce a considerable amount about the nature of all possible equilibria of the dynamic duopoly game.

For example, provided that the investment cost function $K(c)$ is not "too large" we find that there are always two "pure monopoly equilibria" i.e. pure strategy equilibria where one of the firms undertakes all investments, the other never invests, and this enables the investing firm to capture all of the benefits from the cost-reducing investments. By virtue of the way we calculate the equilibria, we can show that it is never optimal for the high cost laggard to challenge the low cost leader by undertaking a cost-reducing investment of its own. Consumers never benefit from price reductions in these equilibria and all of the benefit from the cost reducing investments flows to the firm that undertakes them, in the form of successively lower costs of production. It is not difficult to show that the investment decision rule in these equilibria are identical to the optimal investment strategy of an actual monopolist whose pricing is constrained by the existence of an "outside good" whose price is the same as the initial marginal cost of production of the passive, non-investing firm in the duopoly equilibrium. We will return to this topic and establish this result in section 5.

However our calculations also reveal that there are equilibria involving leapfrogging behavior where the firms do compete dynamically by undertaking competing cost reducing investments. This causes

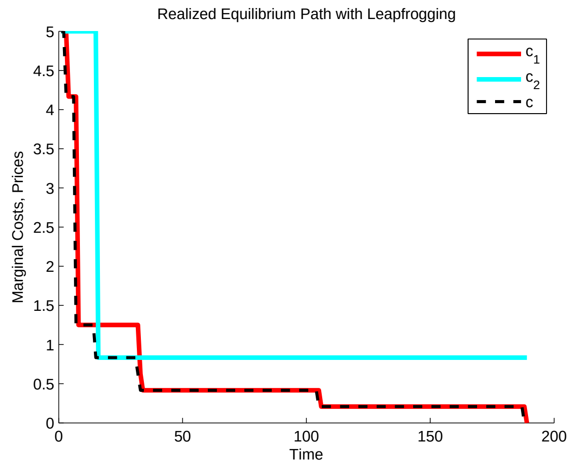

Figure 2 Equilibrium realization with leapfrogging

prices to fall over time so consumers do benefit from declining prices in these equilibria. Figure 2 plots a realization of the equilibrium play in one such game, where both firms 1 and 2 undertake cost reducing investments leapfrogging each other, though not in a pure alternating move fashion with leapfrogging occurring in every period as in the equilibrium of Giovannetti (2001).

We can see from figure 2 that firm 1 is a *dominant firm* and it undertakes cost-reducing investments most of the time. Starting from a symmetric situation where $(c_1, c_2, c) = (5, 5, 5)$, firm 1 undertakes the first two cost-reducing investments, one at time period 3 of the simulation after the state-of-the-art c falls from 5 to 4.1667, and a second investment at time period 7 when c falls again from 4.1667 to 1.25. During this entire time, the prices to the consumer are equal to the initial price, 5, since the low cost leader, firm 1, sets a price equal to the marginal cost of its rival, which remains at its initial value of 5. It is not until period 15, when there is a further technological innovation that decreases c from 1.25 to 0.8333 that firm 2 finally invests, leapfrogging firm 1 to become the low-cost leader. When firm 2 does this, the prices to the consumer finally drop — to $p = 1.25$ — since firm 2 now sets a price equal to the marginal cost of production of firm 1, its higher cost rival. The large price drop in period 15, from $p = 5$ to $p = 1.25$ constitutes a price war caused by firm 2 when it invested and leapfrogged firm 1 to become the new low cost leader.

Prices remain at $p = 1.25$ until period $t = 32$ when c drops again to a value of 0.625. Now firm 1 leapfrogs firm 2 to regain the position of low cost leader, and the price to the consumer falls to $p = 0.8333$. In period $t = 33$ c falls again to 0.4167 and firm 1 invests again to acquire this technology, but the price to

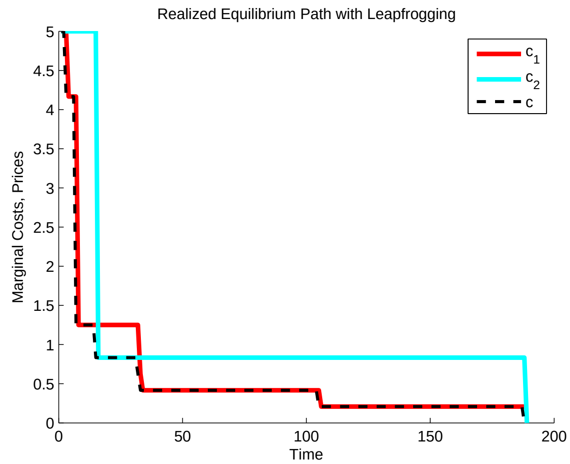

Figure 3 equilibrium realization with leapfrogging

the consumer remains at $p = 0.8333$. Then there is a long interval where there are no further technological innovations and the price remains at this level until period $t = 105$ when c drops to 0.2083 and firm 1 invests once again. Finally, by period $t = 188$ there is a last technological innovation that decreases c to its lowest possible value of $c = 0$, where it remains forever after.[3] Firm 1 decides to invest one more time and attain the best possible marginal cost of production of $c_1 = c = 0$, and secure a position of *permanent* low cost leadership over firm 2. The game then "ends" in an absorbing state where firm 1 can produce at 0 marginal cost and sell to consumers at a price of $p = 0.8333$, which equals the marginal cost of production of firm 2, the high cost "loser".

Figure 3 illustrates a slightly different equilibrium of the model. To isolate the effect of the different equilibrium on the simulated outcomes, we use the same realized path of $\{c_t\}$ in figure 3 as we used in figure 2. This equilibrium realization is almost the same as the one shown in figure 2, except that in period 190, when c_t falls from $c_{189} = .2$ to $c_{190} = 0$, firm 2 *does* invest and leapfrogs firm 1 one final time to become the permanent low cost leader. This means that prices converge to $p = 0.2$ in this equilibrium simulation rather than $p = 0.8333$ in the equilibrium simulation illustrated in figure 2.

[3]Note that for these simulations we discretized the possible values that c could take on 50 possible values over the interval $[0,5]$. When a simulated value of c_t was off of this grid, we used the closest grid point instead. Thus, this discretized simulation process for the Markov process for $\{c_t\}$ can yield the absorbing state $c_t = 0$ in a finite time t, whereas if the actual Markov process governing technological progress has continuously distributed improvements over the current state-of-the-art (such as if c_{t+1} is drawn from a Beta distribution on the interval $(0, c_t)$), then technological progress will only *asymptotically approach* a zero marginal cost of production rather than reaching it in finite time with probability 1. However as noted, when c_t becomes sufficiently small and $K(c)$ does not also tend to zero sufficiently rapidly, there will come a point where the firms no longer have further incentive to invest. Thus, a more accurate simulation of the process would reveal that investments continue until a small but positive value of c_t is reached, after which further investment stops. So in the figures,

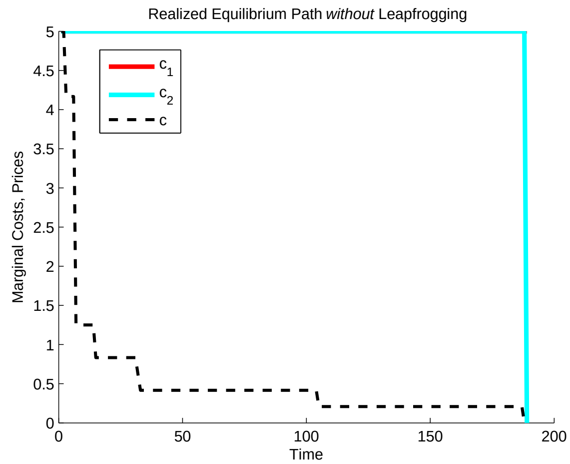

Figure 4 equilibrium realization *without* **leapfrogging**

Figure 4 illustrates a very different equilibrium, again using the same realized path of $\{c_t\}$ as in figures 2 and 3 above. In this equilibrium there is no leapfrogging and no investment, except for a single pre-emptory investment by firm 2 in period $t = 190$ when it invests, acquires the zero marginal cost production technology, and achieves permanent low cost leadership over firm 1. Notice that firm 1 never invests in this equilibrium realization, and so prices to the consumer never fall, and remain at the initial value of $p = 5$ forever. For the first 190 periods of the game, both firms are symmetric Bertrand price competitors and therefore both earn profits of zero. However firm 2 invests in period 190, and starting in period $t = 191$ onward, firm 1 earns profits of 5 by charging a price of $p = 5$. It has thus attained an outcome that is very similar to *limit pricing* by a monopolist. Recall that in limit pricing, a monopolist charges the maximum price it can get away with, subject to the constraint that this price is not too high to induce entry. In this case, the limit price is determined by the marginal cost of production of firm 1, since this firm plays the same role as a new entrant in the limit pricing model: if firm 2 tried to charge more than firm 1's marginal cost of production, there would be room for firm 1 to undercut firm 2, take the entire market, and still earn a profit. Note that there is also a mirror-image equilibrium outcome when we select another equilibrium where firm 1 invests at $t = 190$ instead of firm 2.

Figure 5 illustrates another equilibrium where firm 1 undertakes nearly all of the cost-reducing investments and therefore attains a highly persistent role of low cost leader in this equilibrium realization. Howver in period $t = 190$ firm 2 does finally invest, leapfrogging firm 1 to attain a permanent position of

the reader should interpret $c_t = 0$ as this small positive value of c_t at which further investment is no longer economic.

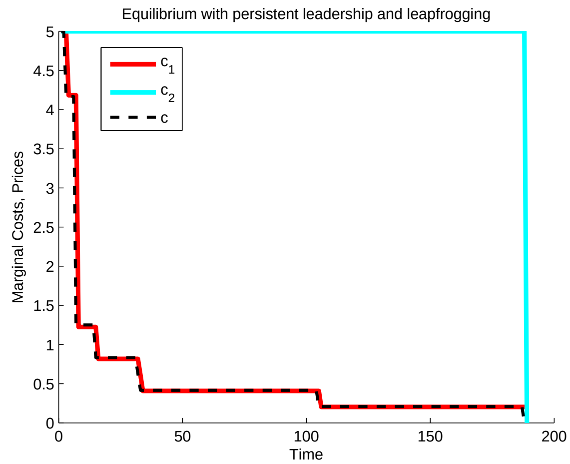

Figure 5 equilibrium realization with leapfrogging and persistent leadership

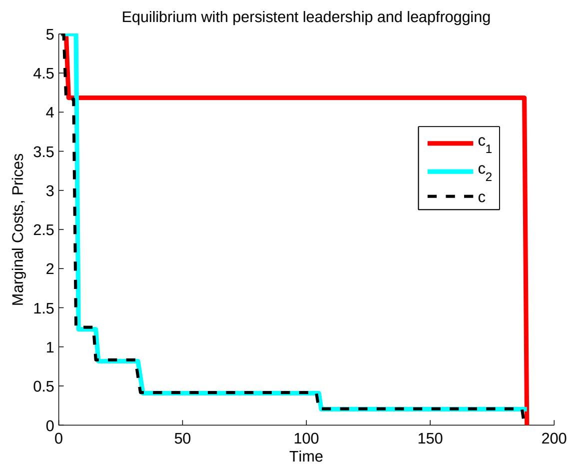

Figure 6 equilibrium realization with leapfrogging and alternating leadership

low cost leadership. From the standpoint of consumers, the equilibrium outcome in fiure 5 is identical to the one displayed in figure 4 for the first 190 periods: the price is $p = 5$ in both cases. All of the cost-reducing investments undertaken by the low cost leader, firm 1, in the first 190 periods accrue entirely to firm 1 and not consumers. However unlike figure 4, when firm 2 finally invests and leapfrogs firm 1 to become the new (permanent) low cost leader in period $t = 190$, a price war breaks out that drives prices from $p = 5$ down to $p = 0.2$, where they remain ever after. Firm 1's profits fall to zero starting in period $t = 191$ and firm 2 is able to earn a small per profit of 0.2 for all $t \geq 191$.

Figure 6 illustrates yet another equilibrium where there is leapfrogging and an alternating pattern of low cost leadership that results in more of the benefits of cost-reducing investments being passed on to

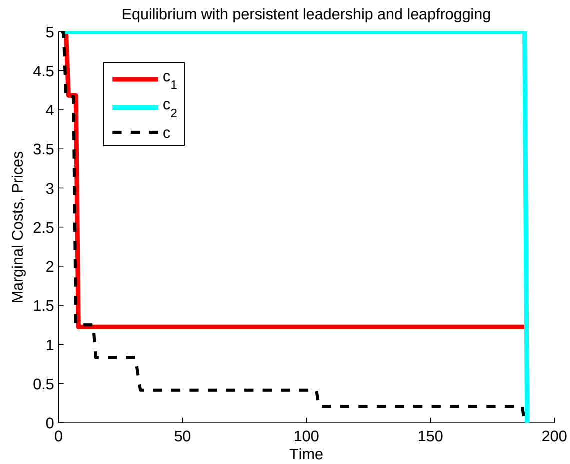

Figure 7 equilibrium realization with leapfrogging and alternating leadership

consumers. Starting from the symmetric situation where $c_1 = c_2 = c = 5$ in period $t = 1$, firm 1 moves first and invests in a new plant that produces at the new lower state-of-the-art marginal cost $c = 4.16667$ in period $t = 3$. Then in period $t = 7$ another large technological innovation occurs that reduces the marginal cost of production under the state-of-the-art from $c = 4.16667$ to $c = 1.25$. This large drop induces firm 2 to invest and leapfrog firm1 to become the new low cost leader, but ths does not ignite a serious price war since prices only fall from $p = 5$ to $p = 4.16667$. Firm 2 remains a persistent low cost leader, undertaking all subsequent cost-reducing investments until period $t = 190$ when firm 1 invests and replaces its high cost plant with a new state-of-the-art plant with a marginal cost of production of $c = 0$. At this point a major price war erupts that drives down prices from $p = 4.16667$ to $p = 0.2$.

Figure 7 provides a final illustration of another equilibrium with leapfrogging and persistent leadership, but where the low cost leader, firm 1, stops investing and "coasts" for an extended period of time after aggressively investing early on in periods $t = 3$ and $t = 7$, where it drove down its marginal cost of production successively from $c = 5$ to $c_1 = 4.1667$ and then to $c_1 = 1.25$. However firm 1 decided not to undertake any further cost reducing investments after that, until in period $t = 190$ firm 2 invested and leapfrogged firm 1 to become the permanent low cost leader. This move ignited a price war that reduced the price from $p = 5$ to $p = 1.25$.

Each of the equilibrium simulations illustrated above correspond to different equilibria of the dynamic game. These are just a few of the many different ones we could have shown. It should be clear that there are many equilibria with a wide range of investment outcomes and prices to consumers. It may be surprising that such complexity can be obtained in such a simple extension of the classical static Bertrand model of price competition, which of course has a very simple, unique solution.

Although we have yet to systematically characterize the set of all equilibria to this model, we have noted that generically (except for a measure zero set of states $(c_1, c_2, c) \in S$) there will be a odd number of equilibria. Generally we have found that there are either 1 or 3 possible equilibria, but we do not know yet whether it is possible to obtain equilibria where there are more than 3 equilibria in any states (c_1, c_2, c). However we can state the following propositions based on our work so far.

Proposition 1: *If the fixed costs of investing $K(c)$ is not too high, no investment by either firm is never an equilibrium outcome. There always exists a value of $c < c_0$ (where c_0 is the initial value for the state-of-the-art marginal cost of production) at which one or both of the firms has a positive probability of investing.*

The proof of this proposition follows from the existence of at least 3 possible equilibria in the end game states $(c_1, c_2, 0)$ at least when $K(0)$ is not too large. Of course, if $K(0)$ is too large relative to the maximum gain from investing in an end game state $(c_1, c_2, 0)$ (i.e. if $\beta c_2/(1-\beta) - K(0) < 0$ and $\beta c_1/(1-\beta) - K(0) < 0$), then no investment will occur in those states. However if we assume that $\beta c_0/(1-\beta) - K(0) > 0$, then there are states in which investment by at least one of the firms will occur with positive probability. Thus, Proposition 1 constitutes our resolution to the "Bertrand investment paradox."

Proposition 2: *If investment in a new technology is not prohibitively costly (so that none of the firms will invest with positive probability in any possible equilibrium of the game), then there will also exist two monopoly equilibria, where one of the firms always invests and the other never invests.*

Proposition 3: *Under the same conditions as Proposition 2, there will also exist a symmetric mixed strategy equilibrium that results in* 0 *ex ante payoffs to the two firms when the game starts in the initial state* (c_0, c_0, c_0).

Propositions 1 to 3, combined with the use of stochastic equilibrium select rules, enable us to characterize the set of possible expected discounted equilibrium payoffs to the two firms in the initial state of the game (c_0, c_0, c_0).

Proposition 4: *Under the same conditions as Propositions 2 and 3, the set of possible expected discounted equilibrium payoffs to the two firms in all possible equilibria of the game is a triangle with vertices at the points* $(0,0)$, $(0, V_M)$ *and* $(V_M, 0)$ *where* $V_M = \max(v_0^i(c_0, c_0, c_0), v_1^i(c_0, c_0, c_0))$ *is the expected discounted monopoly payoff to firm i in the monopoly equilibrium where firm i is the monopolist investor and low cost* *leader.*

The three vertices of the set of equilibrium payoffs in Proposition 4 represent what we might refer to as "extreme equilibria" of the game. As we show in the next section, the two monopoly equilibria constitute the two efficient equilibria of the Bertrand duopoly game, since of course they minimize the occurrence of inefficient duplicative investments by the two firms. Of course, these two monopoly equilibria are the worst from the standpoint of consumers since only the monopolist and not any of the consumers benefits from the cost-reducing investments that the monopolist undertakes.

The worst equilibria for the two firms is the symmetric mixed strategy equilibrium. This equilibrium involves a high probability of inefficient duplicate investment by the two firms in nearly every period, and this intense investment competition results in "rent dissipation" similar to the type of rent dissipation described in Riordan and Salant (1994). Clearly, the firms would be best off in terms of expected payoff if they can agree on a correlated equilibrium that selects one or the other of the monopoly equilibria with some probability $\alpha \in [0,1]$. This equilibrium selection rule can be viewed as a correlated equilibrium to the overall game that solves the underlying investment coordination problem by designating one or the other of the two firms to play the role of monopolist investor.

However it may strain credulity to expect the firms to agree on an efficient equilibrium selection rule. In practice it does not appear that such outcomes are very common. If we assume instead that the firms have somehow agreed on some other equilibrium in this game, then generally all that we can say is that the initial equilibrium payoffs will lie somewhere in the payoff triangle described in proposition 4, and that when there is leapfrogging, the payoffs to the two firms will be in the interior of this payoff triangle, for the simple reason that by its nature, leapfrogging competition results in price reductions to consumers over time, and thus some degree of "rent dissipation".

From figures 2 to 7 above, we have shown that there is a very wide range of profits and prices that are consistent with the Markov perfect equilibria of this model. Some equilibria result in very high prices to consumers, little investment, and high profits for one of the firms, other equilibria can results in high prices, little investment and no profits to either firm, whereas still other equilibria result in active investment by both firms that gives both modest profits while passing the majority of the benefits from these cost reducing investments on to consumers in the form of lower prices.

We have also seen that even when cost-reducing investments occur, they do not always result in price reductions to consumers. Only those investments that result in one firm leapfrogging *over* its opponent to become the new low cost leader result in price reductions to consumers. However there are instances where one firm undertakes a cost-reducing investment starting from a situation where both firms have the same marginal cost of production. In these situations the cost-reducing investment generates no benefit to consumers, similar to the situation where cost-reducing investments are undertaken by the firm that is already the low cost leader. Although these investments do not immediately benefit consumers in the form of lower prices, they can eventually benefit consumers if the other firm eventually does invest and leapfrogs its opponent. This point is illustrated most dramatically in figure 5 where firm 1 undertakes a large number of cost-reducing investments that it captures entirely in increased profits for the first 190 periods of the game, but when firm 2 finally invests and leapfrogs firm 1 in period $t = 190$, the price war that erupts results in a new permanent low price regime for consumers that was only possible due to aggressive prior investments by firm 1. Compare this to figure 4, where absence of cost-reducing investments by either firm in the first 190 periods implies that even when firm 2 finally invested at $t = 190$, the prices would remain forever at $p = 5$.

A final point to note is that behavior reminiscent of "sniping" appears in the equilibrium simulations (see figure 7). By this we mean a situation where one of the firms remains passive and takes the role of the high cost follower for extended periods of time, but the follower does eventually "jump in" by investing at a point when technology improves sufficiently that the firm can invest in a plant that has a sufficiently low marginal cost of production that it deters its opponent from any further attempt to leapfrog to regain the low cost leadership position in the future. These cases illustrate the contestable nature of competition in this model. Being a high cost follower for an extended period of time does not necessarily impair the firm's ability to jump in and leapfrog its opponent at any point in the future, provided that the low cost leader's own investments have not driven down its costs of production too low in the interim. This propensity of the high cost follower to "come from behind" is, we believe, related to our result from section 3 (Lemma 3.1) that in the mixed strategy equilibrium of the $(c_1, c_2, 0)$ end game, the high cost follower has a greater probability of investing than the low cost leader.

5 Socially Optimal Investment

It is of interest to compare investment outcomes from duopoly competition in pricing and investment to those that would emerge under the social planning solution where the social planner is charged with maximizing total expected discounted surplus. In the simple static model of Bertrand price competition, the duopoly solution is well known to be efficient and coincide with the social planning solution: the firm with lower cost produces the good, resulting in consumer demand being satisfied at the smallest possible cost of production.

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, and it is typically ignored in the social planning calculation. However in a dynamic model, the social planner does/should 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 only a single plant, and it would never be optimal to operate two plants as occurs in the duopoly equilibria (except for the two "monopoly" outcomes where one or the other of the firms does all of the investing). Thus, the duopoly equilibria are typically *inefficient* in the sense that there is redundant investment costs that would not be incurred by a social planner.

If we assume that consumers have quasi-linear preferences so that the surplus they receive from consuming the good at a price of p is $u - p$, then the social planning solution involves selling the good at 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 can do is 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. The social planner merely solves and adopts the optimal investment strategy that determines when the current plant should be replaced by a new, cheaper state-of-the-art plant. The goods produced by this optimal plant are provided to consumers in each period at a price equal to the plant's marginal cost of production.

Let $V(c_1, c)$ be the present discounted value of costs of production when the plant operated by the social planner has marginal cost c_1 and the state-of-the-art technology (which is available with one period delay after incurring an investment cost of $K(c)$ just as in the duopoly problem above) has a marginal cost of $c \leq c_1$. We have

$$V(c_1,c) = \min\left[c_1 + \beta\int_0^c V(c_1,c')\pi(dc'|c), c_1 + K(c) + \beta\int_0^c V(c,c')\pi(dc'|c)\right]. \tag{58}$$

The optimal investment strategy can be easily seen to take the form of a *cutoff rule* where the firm invests in the state-of-the-art technology when the current state-of-the-art c falls below a cutoff threshold $\overline{c}(c_1)$, and keeps producing using its existing plant with marginal cost c_1 otherwise. The cutoff rule $\overline{c}(c_1)$ is the solution to the following equation

$$K(\overline{c}(c_1)) = \beta\int_0^{\overline{c}(c_1)} \left[V(c_1,c') - V(\overline{c}(c_1),c')\right]\pi(dc'|\overline{c}(c_1)). \tag{59}$$

This equation tells us that at the optimal cutoff $\overline{c}(c_1)$ 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 $\overline{c}(c_1)$. This implies that the decrease in expected discounted production costs is exactly equal to the cost of the investment when c is equal to the cutoff threshold $\overline{c}(c_1)$. When c is above the threshold, the drop in operating costs is insufficiently large to justify undertaking the investment, and when c is below the threshold, there is a strictly positive net benefit from investing.

Proposition 5: *The socially optimal investment rule $\overline{c}(c_1)$ is also the same at the profit maximizing investment decision rule of a monopolist with discount factor β that can charge a price of c_0 in every period, where c_0 is the initial value of the state-of-the-art production technology.*

The proof of this proposition is straightforward, since if the monopolist is constrained to charge a price of c_0 every period, it follows directly that the monopolist will maximize the expected discounted value of profits if and only if the monopolist follows a cost-minimizing technology adoption strategy. The optimal adoption strategy is to invest in the new technology, reducing the firm's marginal cost of production from c_1 to c whenever $c > \overline{c}(c_1)$ where the threshold $\overline{c}(c_1)$ is given in equation (59). Proposition 5 now leads to Proposition 6, which establishes the efficiency of the monopoly equilibria in the duopoly game.

Proposition 6: *Both of the monopoly equilibria of the duopoly investment and pricing game involve the monopolist adopting an efficient investment policy for investment in new technology.*

The proof of Proposition 6 follows from Proposition 5 and the functional equations defining the Markov perfect equilibria to the duopoly game in equations (22) (23) of section 2. Consider the MPE where firm 1 is the monopolist investor. Then in equation (22) we have $P_1^2(c_1,c_2,c) = 0$ (since it is never optimal for firm 2 to invest in this equilibrium) and the equations for v_0^1 and v_1^1 in equation (22) are equivalent to the Bellman equation for the optimal monopoly investment strategy for firm 1 when it cannot charge a price higher than c_0, which is its opponent's marginal cost of production. From Proposition 5, it follows that the equilibrium investment strategy for firm 1, $\iota_1(c_1, c_0, c)$, is equal to the optimal investment cutoff rule in Proposition 5.

While the result in Proposition 6 may seem trivial, we remind the reader that the existing economics literature, particularly the work of Riordan and Salant (1994), obtained a very different result — namely that in the equilibrium where only one firm does all of the new investment in technology the fear of possible investment by its rival at all points along the equilibrium path causes it to make pre-emptive investments that completely dissipate all monopoly rents. That is, the "monopolist investor" in the equilibrium in Riordan and Salant's model makes *zero expected discounted profits.* Furthermore we now show that the sequence of investments made by the monopolist in their model is *inefficient.*

To show this, recall that Riordan and Salant used a continuous time game formulation where the two duopolists share a common discount rate r. Let K be the cost of adopting a new technology (assumed to be independent of time) and let $c(t)$ be a deterministically decreasing marginal cost of production under the state-of-the-art production technology where t is time. Riordan and Salant's main result is that if the two firms are Bertrand duopolists, there are only two possible equilibria, with one of the firms (the "incumbent") making all of the investments to adopt new technology at a sequence of m dates $(t_1, t_2, \ldots, t_m)$ satisfying

$$\begin{aligned} c(0) - c(t_1) &= rK \\ c(t_1) - c(t_2) &= rK \\ &\vdots \\ c(t_{m-1}) - c(t_m) &= rK. \end{aligned} \tag{60}$$

Now consider the optimal investment strategy of a monopolist who can charge a maximum price per period of $c(0)$ who can also choose a sequence of dates $(s_1, \ldots, s_n)$ to maximize its expected discounted profits. Similar to Proposition 5, it is easy to show that the monopolist's optimal strategy is to choose a sequence of dates for adopting new technologies $(s_1, \ldots, s_n)$ that satisfy the following first order conditions for discounted cost minimization

$$\begin{aligned} c(0) - c(s_1) &= rK - c'(s_1)[1 - \exp\{-r(s_2 - s_1)\}]/r \\ c(s_1) - c(s_2) &= rK - c'(s_2)[1 - \exp\{-r(s_3 - s_2)\}]/r \\ &\vdots \end{aligned}$$

$$
\begin{aligned}
c(s_{n-2}) - c(s_{n-1}) &= rK - c'(s_{n-2})[1 - \exp\{-r(s_n - s_{n-1})\}]/r \\
c(s_{n-1}) - c(s_n) &= rK - c'(s_n)/r. \qquad (61)
\end{aligned}
$$

Comparing the equations defining the optimal investment strategy for a firm in a duopoly equilibrium in equation (60) with the first order condition for a monopolist's optimal investment strategy in equation (61) we find

Proposition 7 *The equilibrium pre-emptive investment strategy in Riordan and Salant's model of Bertrand duopoly pricing and investment is inefficient: it entails the firm investing more frequently and making a larger number of investments (i.e. $m > n$) than occurs under the socially optimal investment rule that a monopolist producer would adopt if there was never any threat of investment by its competitor.*

It follows from Proposition 7 and $s_1 > t_1$ and $s_2 > t_2$ and so on, and generally the total number of times the monopolist would invest, n, is less than the total number of investments mX that occur in the pre-emptive duopoly equilibrium of Riordan and Salant (1994). Intuitively, the threat of entry by the rival forces the incumbent firm to invest more frequently than optimal in order to pre-empt its competitor.

We conclude by analyzing the inefficiency of the duopoly equilibria in our model. To illustrate the nature of the inefficiencies, we solved numerically for the optimal investment strategy using the same parameters that we used to solve for the various duopoly equilibria that we presented in our numerical analysis and simulation section 4. It turns out for the particular realization of the state-of-the-art technology that we used to illustrate the various duopoly equilibria in section 4, it is always optimal to invest: in this case the optimal investment threshold $\overline{c}(c_1)$ is only slightly below the 45 degree line — i.e. $\overline{c}(c_1) < c_1$ but the difference $c_1 - \overline{c}(c_1)$ is small. Thus, in every instance where the state-of-the-art cost improved in figures 2 to 7 of section 4 investment is socially optimal.

We note that in the leapfrogging equilibria displayed in figures 2 and 3 investments were undertaken by one of the firms in every instance where an improvement in the state-of-the-art occurred, so investments were *ex post* optimal in these equilibria. However in figure 4 we showed that no investment occurred until the state-of-the-art cost had finally dropped to its absorbing state of 0. Thus, in this equilibrium there was an *ex post* inefficiency involving *too little investment.*

Similarly figures 5 and 6 display equilibria where the pattern of investments are *ex post* optimal, whereas figure 7 displays another equilibrium where there is *ex post* inefficiency — in the case again there are states where it is optimal to invest but where neither of the duopolists invest in the particular leapfrogging equilibrium we illustrated.

It is also possible to generate examples where the duopoly equilibrium involves *excessive investment* relative to the social optimum, but in an *ex post* sense and in an *ex ante sense* (i.e. where there is a positive probability of investments occurring in states where it is not socially optimal for investment to occur).

6 Conclusions

We have developed a model of Bertrand price competition and sequential investments in a stochastically improving technology and resolved the *Bertrand investment paradox* — that is, we showed that it is possible and indeed generally the case that investment occurs in equilibrium in this model even though by the nature of the problem the market is contestable and both firms have the option at any point in time to invest in the state of technology to reduce their marginal cost of production. Casual reasoning would suggest that if both firms realize that if they both invest in the state-of-the-art technology with marginal cost of production c, Bertrand price competition will erupt after their simultaneous adoption of the new technology and ensure that neither firm will make positive profits from this investment in new technology.

We have resolved the Bertrand investment paradox by showing that the equilibria of the duopoly pricing and investment game is isomorphic to a coordination game, and that there are equilibria involving leapfrogging investments that enable the firms to implicitly coordinate, avoiding the potential losses from duplicative simultaneous investments.

We have also related our model to the existing literature on technology investment adoption under Bertrand duopoly that argues that the only equilibrium outcomes are ones that involve a single firm doing all of the investing, with investments strategically timed to pre-empt the other firm from investing and leapfrogging the low cost leader. We have showed that this result depends critically on the assumption that technology adoption is an alternating move game. Our framework models technology adoption as a simultaneous move game and we show that this results in a huge multiplicity of equilibria.

We do not take a stand on whether outcomes “in reality” are better approximated as simultaneous move or alternating move games. In many respects the model in this paper is too oversimplified to be treated as a serious theoretical model of dynamic competition. As we discuss further below, we find it disconcerting that even the simplest dynamic models such as the one we studied here can have so many equilibria and details about whether firms move simultaneously or alternately can have such an important bearing on the equilibrium behavior predicted by these models.

Nevertheless we attempted to characterize the set of all equilibria of this game and we developed a state-space recursion algorithm that can calculate all equilibria of the game. We showed that among the equilibria are two “monopoly equilibria” where one of the firms always does the investing in new technology and the other firm never invests. However unlike the equilibrium in the Riordan and Salant model, we showed that the firm doing the investing in these equilibria earns fully monopoly rents and adopts a socially optimal (expected discounted cost minimizing) adoption strategy. This suggests if the firms could coordinate, the optimal coordination would involve their playing a correlated equilibria that makes one or the other firm the monopolist with some probability $\alpha \in [0,1]$.

We believe a separate theoretical contribution 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 that there are equilibria of our model where cost-reducing investments are relatively infrequent and leapfrogging rarely occurs, so that consumers obtain little or no benefit 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 curiosum, or whether this model can provide insights into a variety of possible competitive behaviors in actual markets.

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 her model, with the cost of investing in the state-of-the-art production facility declining geometrically in each period. She 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 sufficiently many new consumers can be drawn away from consuming the outside good, 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 can continue indefinitely in our model as well.

Giovannetti also found there were equilibria with "persistent leadership" an outcome she termed *increasing asymmetry*. These equilibria are the analogs of the equilibria we find in our model where one of the firms takes the role of "low cost leader" for extended periods of time and does all of the investing at every point in time where there is a sufficiently large reduction in the marginal cost of production in the state-of-the-art technology, relative to its fixed investment cost. However 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.

We also refer the reader to the very important paper by Goettler and Gordon (2010) that studies leapfrogging R&D and pricing decisions by the duopolists Intel and AMD. This model is considerably more complex than our model in that AMD and Intel leapfrog each other by undertaking R&D investments to produce faster microprocessors rather than by simply investing in a cost reducing production technology that evolves exogenously as in our model. In addition, the Goettler and Gordon model has consumers that make *dynamic* rather than static choices about whether to purchase a new computer with the latest microprocessor, or keep their existing computer with a prior-generation microprocessor. This creates considerable complexity and added interesting dynamics, since the duopolists must consider as a relevant state variable *the entire distributions of holdings of microprocessors in the consumer population.* When a sufficiently large fraction of consumers have sufficiently outdated microprocessors, conditions are more opportune for gaining a large market share by introducing a newer, faster microprocessor.

The Goettler and Gordon analysis shows that the dynamic R&D competition results in a form of "leapfrogging" that results in lower prices and better computers for consumers. However, similar to our finding of investment inefficiency in our simpler framework (with innovation occurring less frequently than the socially optimal level in several equilibria we analyzed in section 5), Goettler and Gordon find that "innovation would be 8.2 percent higher if Intel were a monopoly. Consumer surplus, however, would be 2.5 percent ($9 billion per year) lower without AMD since prices would be higher. To evaluate the effect of Intel's alleged anti-competitive practices, we perform counterfactual simulations in which we vary the share of the market from which AMD is foreclosed. As the foreclosure share rises, prices and industry innovation both increase, lending support to the Schumpterian hypothesis. Moreover, consumer surplus peaks when AMD is excluded from half the market since the higher innovation more than fully oiffsets the effect of higher prices on surplus. This finding supports the FTC's recent consideration of the dynamic trade-off between lower current consumer surplus from higher prices and higher future surplus from more innovation." (p. 3).

Our study is not intended to present an empirically realistic analysis of a specific industry such as the microprocessor industry that Goettler and Gordon analyzed. While the comments below are not a specific criticism of the Goettler and Gordon study (which we think is an important contribution both theoretically and empirically), our analysis can serve as a cautionary note on the potential problems that are created when there are a vast multiplicity of Markov-perfect equilibria in dynamic games. This has important consequences for how these models ought to be applied and interpreted in the empirical IO literature. We believe there may be a misimpression in the empirical IO literature that the "Markov-perfect" restriction can be relied on to guarantee that the models will have unique equilibria, facilitating empirical analysis and comparative dynamic policy analysis. However if in fact there a vast number of possible equilibria in many of these dynamic models, the implications for empirical IO become much cloudier since it is not clear that economic theory has much to say about which equilibrium will be "selected" in an particular empirical context.

Although there have been attempts to empirically select the most likely of several possible equilibria in recent empirical work, our findings suggest that even in our very simple extension of the basic Bertrand price competition, the variety of equilibria are so great that essentially "anything can happen" — a result that is reminiscent of the Folk Theorems in the older repeated game literature, and one that has largely limited the empirical relevance and number of empirical applications of this literature to IO.

The new, dynamic, Markov-perfect equilibrium approaches that are transforming empirical work in the new structural IO literature may be subject to similar criticisms and limitations. One of the important strategies for structural estimation in this new literature involves "nested fixed point algorithms" that consist of an inner fixed point algorithm that repeatedly recalculates the equilibrium to a dynamic game for different trial values of the unknown parameters of the model, while an outer optimization algorithm searches for values of these parameters that maximize a likelihood function, or minimize some measure of the distance between the predictions of the dynamic model and the data. If the inner fixed point algorithm is based on an iterative successive approximations type of algorithm, our results suggest that these algorithms, when they converge, can be acting as implicit *equilibrium selection devices.* It can happen that small details in how these algorithms are initialized and specified can affect which particular equilibrium of the game they converge to.

If the empirical analyst is not aware of these problems, and if the inner fixed point algorithm is not converging to the "same" equilibrium for different trial values of the parameters being estimated, then the underlying "equicontinuity" and econometric properties that researchers in the structural IO literature rely on to establish the statistical properties of their estimators may be called into question. Further, "policy analysis" of these models, including predicting how the equilibrium would change as a result of a counterfactual technological change or policy innovation are called into question if the model solution algorithm is calculating one possible equilibrium prior to the policy change and a different one after the change. Then, some of the imputed "behavioral change" may be an artifact of an inadvertent selection of a different equilibrium of the model, not a true impact resulting from the change in the "actual" equilibrium (if we assume that in reality the players/firms are in fact able to select and coordinate on particular equilibria before and after the policy change).

Overall, we were quite surprised by how complex are the various types of equilibrium behavior that can emerge from such a simple model. We do not regard this multiplicity as a good thing, and are currently looking for reasonable alternative formulations of the model that may succeed in limiting the number of equilibria — ideally to result in a unique equilibrium. One possibility is to go back to the alternating move structure that Riordan and Salant (1994) analyzed. However we do not regard the much more limited set of equilibria in their model to be particularly plausible, and are considering specifications with randomly alternating moves to see if it is possible to vastly reduce the number of equilibria in our model but while still obtaining leapfrogging behavior, which we believe is realistic and occurs commonly in reality.

A final direction we are interested in exploring is add switching cost frictions and dynamic consumer choice as well as capacity constraints to the model and use it to understand whether it would lead to results in an infinite horizon setting similar to those found by Kreps and Scheinkman (1983) in a two-stage game framework. Namely, whether capacity investment followed by Bertrand price competition yields an outcome identical to the Cournot-Nash equilibrium in a model where firms chooses quantities only.

7 Appendix: Proof of Lemma 1

Lemma 3.1 *Suppose $\eta = 0$. If $c_1 > c_2 > 0$ and $K < \frac{\beta c_2}{1-\beta}$, then in the unique mixed strategy equilibrium of the pure Bertrand dynamic investment and pricing game in state $(c_1, c_2, 0)$ we have $\pi_1 > \pi_2$.*

Proof. First, note that the condition $K < \frac{\beta c_2}{1-\beta}$ in Lemma 3.1 ensures that investment is profitable in the long term even for firm 1 whose potential pay-off is smaller ($\frac{\beta c_2}{1-\beta} < \frac{\beta c_1}{1-\beta}$). In other words, this condition ensures that for both firms' investment decisions are economically justified. Next, observe that when $\beta = 0$ in the $(c_1, c_2, 0)$ end game there is unique pure strategy equilibrium where neither of the companies invests. Thus, we only consider the case $\beta > 0$.

The value functions of the two firms in the $(c_1, c_2, 0)$ end game when $c_1 > c_2$ are

$$\begin{aligned} V_1 &= \pi_1 \times \left(\pi_2 \cdot (-K) + (1-\pi_2) \cdot \left(\frac{\beta c_2}{1-\beta} - K \right) \right) + \\ &\quad + (1-\pi_1) \times (\pi_2 \cdot 0 + (1-\pi_2) \cdot \beta V_1) \\ V_2 &= \pi_2 \times \left(\pi_1 \cdot (c_1 - c_2 - K) + (1-\pi_1) \cdot \left(c_1 - c_2 + \frac{\beta c_1}{1-\beta} - K \right) \right) + \\ &\quad + (1-\pi_2) \times (\pi_1 \cdot (c_1 - c_2) + (1-\pi_1) \cdot (c_1 - c_2 + \beta V_2)) \end{aligned}$$

where the definition of the probability π_1 of investment by firm 1 in the mixed strategy equilibrium gives

$$\pi_2 \cdot (-K) + (1-\pi_2) \cdot \left(\frac{\beta c_2}{1-\beta} - K \right) = \pi_2 \cdot 0 + (1-\pi_2) \cdot \beta V_1$$

and thus the value function itself becomes the weighted sum of equal parts, leading to

$$V_1 = \pi_2 \cdot (-K) + (1-\pi_2) \cdot \left(\frac{\beta c_2}{1-\beta} - K \right) = \pi_2 \cdot 0 + (1-\pi_2) \cdot \beta V_1$$

Using the second equality in the last expression, we find $V_1 = 0$, and then using the first equality in the same expression, we find $1 - \pi_2 = \frac{K(1-\beta)}{\beta c_2}$.

The definition of the probability π_2 of investment by firm 2 in the mixed strategy equilibrium, similarly gives

$$\begin{aligned} V_2 &= \pi_1 \cdot (c_1 - c_2 - K) + (1-\pi_1) \cdot \left(c_1 - c_2 + \frac{\beta c_1}{1-\beta} - K \right) \\ &= \pi_1 \cdot (c_1 - c_2) + (1-\pi_1) \cdot (c_1 - c_2 + \beta V_2) \end{aligned}$$

Using the second equality in the last expression, we find $V_2 = \frac{c_1 - c_2}{(1-\beta \cdot (1-\pi_1))}$, and using the it once again we get

$$\begin{aligned}
\pi_1(c_1-c_2-K)+(1-\pi_1)\left(c_1-c_2+\frac{\beta c_1}{1-\beta}-K\right) &= \pi_1(c_1-c_2)+(1-\pi_1)(c_1-c_2+\beta V_2) \\
(1-\pi_1)\left(\frac{\beta c_1}{1-\beta}-K\right)-\pi_1 K &= (1-\pi_1)\beta V_2 \\
\frac{c_1}{1-\beta}-\frac{K}{\beta\cdot(1-\pi_1)} &= V_2
\end{aligned}$$

Combining the two expressions for the value function V_2, we get the following equation

$$\frac{c_1-c_2}{1-\beta\cdot(1-\pi_1)} = \frac{c_1}{1-\beta}-\frac{K}{\beta\cdot(1-\pi_1)}$$

Multiplying by $1-\beta$ and incerting the expression for $1-\pi_2$, we have

$$\begin{aligned}
\frac{c_1-c_2}{1+\frac{\beta}{1-\beta}\pi_1} &= c_1-\frac{1-\pi_2}{1-\pi_1}c_2 \\
\frac{c_1-\frac{1-\pi_2}{1-\pi_1}c_2}{c_1-c_2} &= \frac{1}{1+\frac{\beta}{1-\beta}\pi_1}\leqslant 1 \\
c_1-\frac{1-\pi_2}{1-\pi_1}c_2 &\leqslant c_1-c_2 \\
\frac{1-\pi_2}{1-\pi_1} &\geqslant 1 \\
\pi_1 &\geqslant \pi_2
\end{aligned}$$

The inequalities are due to the fact that $0\leqslant\pi_1\leqslant 1$, $\frac{\beta}{1-\beta}>0$, $c_1-c_2>0$, $c_2>0$. The final inequality is strict unless $\pi_1=\pi_2=0$, which implies $K=\frac{\beta c_2}{1-\beta}$ thus leading to a contradiction. We conclude then that $\pi_1>\pi_2$. $\square$

References

[1] Anderson, S., A. DePalma and J. Thisse (1992) *Discrete Choice Theory of Product Differentiation* MIT Press.

[2] Aumann, R. (1987) "Correlated Equilibrium as an Expression of Bayesian Rationality" *Econometrica* **55-1** 1–18.

[3] Baumol, W.J. and Panzar, J.C. and Willig, R.D (1982) *Contestable Markets and the Theory of Industry Structure* Harcourt, Brace Jovanovich.

[4] Baye, M. and D. Kovenock (2008) "Bertrand Competition" in *New Palgrave Dictionary of Economics* 2nd Edition.

[5] 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.

[6] Cournot, A. (1838) *Recherches sur les Principes de la Théorie du Richesses* Paris: Hachette.

[7] Doraszelski, Ulrich and Juan Escobar (2010) "A Theory of Regular Markov Perfect Equilibria in Dynamic Stochastic Games: Genericity, Stability, and Purification" forthcoming, *Theoretical Economics.*

[8] Doraszelski, Ulrich and Mark Satterthwaite (2010) "Computable Markov-Perfect Industry Dynamics" *Rand Journal of Economics* **41-2** 215–243.

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

[10] Gilbert, R and D. Newbery (1982) "Pre-emptive Patenting and the Persistence of Monopoly" *American Economic Review* **74** 514–526.

[11] Giovannetti, Emanuelle (2001) "Perpetual Leapfrogging in Bertrand Duopoly" *International Economic Review* **42-3** 671–696.

[12] Goettler, Ronald and Brett Gordon (2009) "Does AMD spur Intel to innovate more?" working paper, University of Chicago Booth School of Business.

[13] Hall, Robert E. (2008) "Potential Competition, Limit Pricing, and Price Elevation from Exclusionary Conduct" in *Issue in Competition Law and Policy* **433** (ABA Section of Antitrust Law).

[14] Harsanyi, J. (1973a) "Games with randomly disturbed payoffs: A new rationale for mixed strategy equilibrium points" *International Journal of Game Theory* **2-1** 1–23.

[15] Harsanyi, J. (1973b) "Oddness of the Number of Equilibrium Points: A New Proof" *International Journal of Game Theory* **2-4** 235–250.

[16] Kreps, D. and J. Scheinkman (1983) "Quantity Precommitment and Bertrand Competition Yield Cournot Outcomes" *Bell Journal of Economics* **14-2** 326–337.

[17] Reinganum, J. (1981) “On the Diffusion of New Technology — A Game-Theoretic Approach” *Review of Economic Studies* **153** 395–406.

[18] Riordan, M. and D. Salant (1994) “Preemptive Adoptions of an Emerging Technology” *Journal of Industrial Economics* **42-3** 247–261.

[19] Riordan, M. (2010) “Quality Competition and Multiple Equilibria” working paper, Department of Economics, Columbia University.

[20] Rosenkranz, Stephanie (1996) “Quality Improvements and the Incentive to Leapfrog” *International Journal of Industrial Organization* **15** 243–261.

[21] Rust, John (1986) “When Is it Optimal to Kill Off the Market for Used Durable Goods?” *Econometrica* **54-1** 65–86.

[22] Scherer, F. (1967) “Research and Development Resource Allocation Under Rivalry” *Quarterly Journal of Economics* **81** 359–394.

[23] Vickers. M. (1986) “The Evolution of Market Structure When There is a Sequence of Innovations” *Journal of Industrial Economics* **35** 1–12.