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

Joseph E. Harrington, Jr.
Johns Hopkins University

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

John Rust[‡]
University of Maryland

Bertel Schjerning
University of Copenhagen

July, 2010

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 "leap frog" 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 leap frogs 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 staggered or alternating investments by the two firms as they vie for temporary production-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 leap frogs its opponent to become the new low cost leader.

Keywords: Bertrand-Nash price competition, leap frogging, cost-reducing investments, dynamic models of competition, Markov-perfect equilibrium, tacit collusion, price wars

JEL classification: D92, L11, L13

[†]Preliminary: This version of this paper is not for general distribution and quotation.

[‡]**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 "leap frogging" 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. Thus, the motivation for the high cost firm to undertake a leap frogging 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. Whenever both firms invest the same time, then both will have the same marginal cost of production, and the resulting Bertrand price competition will ensure that neither firm can earn positive profits that would justify undertaking the investment in the first place. Therefore, casual reasoning would suggest that *Bertrand duopolists would not have any incentive to undertake cost-reducing investments.* We refer to this as the *Bertrand investment paradox.*

We show that this simplistic reasoning is incorrect and provide a resolution of the Bertrand investment paradox by solving a dynamic, infinite horizion extension of the Bertrand model of price competition. The extended version of the Bertrand model allows the competing firms to invest in improved technology

[1]According to Baye and Kovenock, Bertrand did not realize that the perfectly competitive outcome emerges as the equilibrium solution to price competition. In Bertrand's 1883 review of Cournot's 1838 book, "Bertrand 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).

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. However we show that even in complete information versions of this model, and even when firms start with the same marginal costs of production, there are many equilibria that result in various types of leap frogging behavior by the two firms.

Simulations of the solution of the model reveal that the equilibrium realizations typically involve staggered or 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 yields a new interpretation of 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, and solve what we refer to as the “end game” in section 3 and show that this leads to key insights into the form of the full equilibria of the model which we solve and illustrate in section 4. In section 5 we formulate and solve the social planner’s problem and characterize the investment strategy that maximizes total surplus. We show that generally, the equilibria of our model result in inefficient investments relative to the social optimum. So unlike the simple static Bertrand model of price competition, a simple dynamic generalization to allow for investments shows that oligopolistic equilibria are generally inefficient. 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 famous 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 leap frogging 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, leap frogging 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 leap frogging is possible even in the limiting "pure Bertrand" case where consumer demand is perfectly elastic. This represents the most challenging case for leap frogging, 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.

That is, neither firm has any inherent advantage in being the low cost leader other than the profits they can earn by virtue of their temporarily lower marginal cost of production. At any time the high cost follower can to acquire the current state-of-the-art production technology, and thereby assure itself of a marginal cost of production that is at least as low as the 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 investing in the state-of-the-art production technology.

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 leap frogging equilibrium, since the entry of a new competitor provides another mechanism by which high cost firms can be leap frogged 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 equilibrium 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, but in other respects very different. For example, we show there are equilibria where 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. This outcome is similar to the behavior of a monopolist where there is an "outside good" with a price equal to the marginal cost of prodution of the high cost follower firm. The low cost leader undertakes periodic investments to reduce its cost of production, but the consumers never benefit from these investments. Instead, the benefits flow exclusively to the low cost leader during these long "leadership" epochs. Unlike a tacitly collusive outcome, however, the high cost follower firm does not benefit in this equilibria either: it earns zero profits in all periods and the low cost leader receives all the profits and benefits from cost-reducing investments.

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.

In section 5 we formulate the social planner's problem and characterize the optimal investment strategy. We then compare the optimal investment strategy to the investments that occur in the Bertrand duopoly equilibrium and show that generally the equilibria result in inefficient investments and a lower total surplus. This result stands in contrast to the outcome of the static Bertrand equilibrium, which results in a fully efficient solution that delivers full surplus to the consumers (and no profits to the competing firms).

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 (or price) payoff for 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

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

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 $p_1^*(c_1, c_2)$ and $p_2^*(c_1, c_2)$ that solve the following fixed-point problem

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

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

$$p(c_1, c_2) = \max[c_1, c_2]. \tag{1}$$

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:

$$\Pi_1(p_1, p_2) = \frac{\exp\{-p_1/\sigma\}}{\exp\{-\gamma_0/\sigma\} + \exp\{-p_1/\sigma\} + \exp\{-p_2/\sigma\}}. \tag{2}$$

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 is 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 leap frogging 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 leap frogging 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 a cone in R^3, $S = \{(c_1, c_2, c) | c_1 \geq c \text{ and } c_2 \geq c\}$.

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

$$p(c_t) = \frac{.01c_t}{1 + .01c_t}. \tag{3}$$

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 decides whether or not to invest. Both firms know each other's marginal cost of production, i.e. there is common knowledge of (c_{1t}, c_{2t}). Each firm also knows that the cost of buying the current state-of-the-art technology $K(c_t)$, but each firm also incurs idiosyncratic "disruption costs" $\varepsilon_t = (\varepsilon_{0t}, \varepsilon_{1t})$ associated with each of the choices of not to invest (ε_{0t}) and investing (ε_{1t}).

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 $(\varepsilon_t^1, \varepsilon_t^2)$ is an *IID* Type 1 bivariate extreme value process, and that each firm knows its own idiosyncratic cost to investment shock, but does does not know it's opponent's idiosyncratic investment cost shocks. However we assume both firms have common knowledge of the stochastic process for these idiosyncratic investment costs, and believe that they evolve as an *IID* Type 1 extreme value process, and both know η. 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.

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, which 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.

To derive the 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), i.e. firm 1 has a marginal cost of production c_1, firm 2 has a marginal cost of production c_2, and the marginal cost of production using the current best technology is c. Since we have assumed that the two firms can both invest in the current best technology at the same cost $K(c)$, it is tempting to conjecture that there should be a "symmetric equilibrium" where by "symmetric" we mean an equilibrium where the decision rule and value function for firm 1 depends on the state (c_1,c_2,c), and similarly for firm 2, and these value functions and decision rules are *anonymous* (also called *exchangeable*) in the sense that

$$V^1(c_1,c_2,c,\varepsilon_0,\varepsilon_1) = V^2(c_2,c_1,c,\varepsilon_0,\varepsilon_1) \tag{4}$$

where $V^1(c_1,c_2,c,\varepsilon_0,\varepsilon_1)$ is the value function for firm 1 when the mutually observed state is (c_1,c_2,c), 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 (4) 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 $(c_2,c_1,c,\varepsilon_0,\varepsilon_1)$, 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 (4) 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 (4) holds. We will show that the "interesting" equilibria of this model, including the various types of equilibria with leap frogging, 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

$$V^i(c_1, c_2, c, \varepsilon_0^i, \varepsilon_1^i) = \max[v_0^i(c_1, c_2, c) + \eta\varepsilon_0^i, v_1^i(c_1, c_2, c) + \eta\varepsilon_1^i] \tag{5}$$

where $v_0^i(c_1, c_2, c)$ is the expected value to firm i if it does not invest in the latest technology, and $v_1^i(c_1, c_2, c)$ 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

$$r^1(c_1, c_2) = \begin{cases} 0 & \text{if } c_1 \geq c_2 \\ \max[c_1, c_2] - c_1 & \text{otherwise} \end{cases} \tag{6}$$

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 the expected profits associated with *not* investing (after taking expectations over player i's privately observed idiosyncratic shocks $(\varepsilon_0, \varepsilon_1)$ but conditional on the publicly observed state variables (c_1, c_2, c)) is given by:

$$v_0^i(c_1, c_2, c) = r^i(c_1, c_2) + \beta EV^i(c_1, c_2, c, 0). \tag{7}$$

$EV^i(c_1, c_2, c, 0)$ denotes firm i's conditional expectation of its next period value function $V^i(c_1, c_2, c, \varepsilon_0^i, \varepsilon_1^i)$ given that it does not invest this period, conditional on (c_1, c_2, c).

The formula for the expected profits associated with investing (after taking conditional expectations over firm i's privately observed idiosyncratic shocks $(\varepsilon_0^i, \varepsilon_1^i)$ but conditional on the publicly observed state variables (c_1, c_2, c))) is given by

$$v_1^i(c_1, c_2, c) = r^i(c_1, c_2) - K(c) + \beta EV^i(c_1, c_2, c, 1) \tag{8}$$

where $EV^i(c_1, c_2, c, 1)$ is firm i's conditional expectation of its next period value function $V^i(c_1, c_2, c, \varepsilon_0^i, \varepsilon_1^i)$ given that it invests, conditional on (c_1, c_2, c).

To compute the conditional expectations $EV^i(c_1,c_2,c,0)$ and $EV^i(c_1,c_2,c,1)$ we invoke a well known trick and 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 $(\varepsilon_0^i,\varepsilon_1^i)$ is given by the standard "log-sum" formula when these shocks have the Type-III extreme value distribution. Thus, after taking expectations over $(\varepsilon_0^i,\varepsilon_1^i)$ in the equation for V^i in (5) above, we have

$$\int_{\varepsilon_0^i}\int_{\varepsilon_1^i} V^i(c_1,c_2,c,\varepsilon_0^i,\varepsilon_1^i,j)q(\varepsilon_0^i)q(\varepsilon_1^i)d\varepsilon_1^i d\varepsilon_0^i = \eta\log\left[\exp\{v_0^i(c_1,c_2,c)/\eta\}+\exp\{v_1^i(c_1,c_2,c)/\eta\}\right], \quad (9)$$

where $j\in\{0,1\}$.

The log-sum formula provides a closed form expression for the conditional expectation $V^i(c_1,c_2,c,\varepsilon_0^i,\varepsilon_1^i)$ for each firm i, where V^i is the maximum of the value of not investing $j=0$ or investing $j=1$ as we can see from equation (5) 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 (9) with respect to the next-period values of $(\varepsilon_0^i,\varepsilon_1^i)$. However we do need to compute the two functions $v_0^i(c_1,c_2,c)$ and $v_1^i(c_1,c_2,c)$ for both firms $i=1,2$. We will describe one algorithm for doing this below.

To simplify notation, we let $\phi(v_0^i(c_1,c_2,c),v_1^i(c_1,c_2,c))$ be the log-sum formula given above in equation (9), that is define ϕ as

$$\phi(v_0^i(c_1,c_2,c),v_1^i(c_1,c_2,c)) \equiv \eta\log\left[\exp\{v_0^i(c_1,c_2,c)/\eta\}+\exp\{v_1^i(c_1,c_2,c)/\eta\}\right]. \quad (10)$$

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

$$\lim_{\eta\to 0}\phi(v_0,v_1) = \max[v_0,v_1]. \quad (11)$$

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

Let $P_1^1(c_1,c_2,c)$ be firm 2's belief about the probability that firm 1 will invest if the mutually observed state is (c_1,c_2,c). Firm 1's investment decision is probablilistic 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_1^1 as the following binary logit formula

$$P_1^1(c_1,c_2,c) = \frac{\exp\{v_1^1(c_1,c_2,c)/\eta\}}{\exp\{v_1^1(c_1,c_2,c)/\eta\}+\exp\{v_0^1(c_1,c_2,c)/\eta\}} \quad (12)$$

Firm 2's belief of firm 1's probability of not investing, $P_0^1(c_1,c_2,c)$ is of course simply $1-P_1^1(c_1,c_2,c)$.

Firm 1's belief of the probability that firm 2 will invest, $P_1^2(c_1,c_2,c)$, is given by

$$P_1^2(c_1,c_2,c) = \frac{\exp\{v_1^2(c_1,c_2,c)/\eta\}}{\exp\{v_1^2(c_1,c_2,c)/\eta\} + \exp\{v_0^2(c_1,c_2,c)/\eta\}} \tag{13}$$

If the symmetry condition holds, then we have $P_1^2(c_1,c_2,c) = P_1^1(c_2,c_1,c)$.

Now we are in position to write the recursion formulas for the conditional expectations $EV^i(c_1,c_2,c,0)$ and $EV^i(c_1,c_2,c,1)$, corresponding to firm i not investing and investing, respectively. For firm 1 we have:

$$\begin{aligned}
EV^1(c_1,c_2,c,0) &= \int_0^c \left[P_1^2(c_1,c_2,c)H^1(c_1,c,c') + (1-P_1^2(c_1,c_2,c))H^1(c_1,c_2,c')\right]\pi(dc'|c) \\
EV^1(c_1,c_2,c,1) &= \int_0^c \left[P_1^2(c_1,c_2,c)H^1(c,c,c') + (1-P_1^2(c_1,c_2,c))H^1(c,c_2,c')\right]\pi(dc'|c) \\
EV^2(c_1,c_2,c,0) &= \int_0^c \left[P_1^1(c_1,c_2,c)H^2(c,c_2,c') + (1-P_1^1(c_1,c_2,c))H^2(c_1,c_2,c')\right]\pi(dc'|c) \\
EV^2(c_1,c_2,c,1) &= \int_0^c \left[P_1^1(c_1,c_2,c)H^2(c,c,c') + (1-P_1^1(c_1,c_2,c))H^2(c,c_2,c')\right]\pi(dc'|c),
\end{aligned} \tag{14}$$

where H^1 and H^2 are given by

$$\begin{aligned}
H^1(c_1,c_2,c) &= \phi(v_0^1(c_1,c_2,c), v_1^1(c_1,c_2,c)) \\
H^2(c_1,c_2,c) &= \phi(v_0^2(c_1,c_2,c), v_1^2(c_1,c_2,c)).
\end{aligned} \tag{15}$$

Substituting these expressions into the equations (7) and (8) defining v_0^i and v_1^i respectively, results in the following set of recursive equations for the equilibrium

$$\begin{aligned}
v_0^1(c_1,c_2,c) &= r^1(c_1,c_2) + \beta\int_0^c \left[P_1^2(c_1,c_2,c)\phi(v_0^1(c_1,c,c'), v_1^1(c_1,c,c'))\right. \\
&\quad \left.(1-P_1^2(c_1,c_2,c))\phi(v_0^1(c_1,c_2,c'), v_1^1(c_1,c_2,c'))\right]\pi(dc'|c). \\
v_1^1(c_1,c_2,c) &= r^1(c_1,c_2) - K(c) + \beta\int_0^c \left[P_1^2(c_1,c_2,c)\phi(v_0^1(c,c,c'), v_1^1(c,c,c'))\right. \\
&\quad \left.(1-P_1^2(c_1,c_2,c))\phi(v_0^1(c,c_2,c'), v_1^1(c,c_2,c'))\right]\pi(dc'|c).
\end{aligned} \tag{16}$$

$$\begin{aligned}
v_0^2(c_1,c_2,c) &= r^1(c_2,c_1) + \beta\int_0^c \left[P_1^1(c_1,c_2,c)\phi(v_0^2(c,c_2,c'), v_1^2(c,c_2,c'))\right. \\
&\quad \left.(1-P_1^1(c_1,c_2,c))\phi(v_0^2(c_1,c_2,c'), v_1^2(c_1,c_2,c'))\right]\pi(dc'|c). \\
v_1^2(c_1,c_2,c) &= r^1(c_2,c_1) - K(c) + \beta\int_0^c \left[P_1^1(c_1,c_2,c)\phi(v_0^2(c,c,c'), v_1^2(c,c,c'))\right. \\
&\quad \left.(1-P_1^1(c_1,c_2,c))\phi(v_0^2(c_1,c,c'), v_1^2(c_1,c,c'))\right]\pi(dc'|c).
\end{aligned} \tag{17}$$

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

To our knowledge, there is no analytic closed-form solution to the set of functional equations (16) and (17). Instead, in the remainder of this paper we attempt to solve them by a modified method of successive approximations using numerical interpolation and quadrature to undertake the calculations in equations (16) and (17). Once solutions are calculated, the model can be simulated to reveal the behavior implied by this dynamic extension of the Bertrand model.

Although the system appears to resemble a pair of "Bellman equations" (one for firm 1 and one for firm 2) and the 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 (16) and (17) is *not* a type of contraction mapping. So far from having a unique solution, there can be many different solutions to equations (16) and (17). The various solutions to these equations correspond to different equilibria of the dynamic duopoly game.

Another implication of the fact that equations (16) and (17) do not define the equilibrium values of the two firms as a fixed point to a contraction mapping is that the method of *successive approximations* (also known as backward induction) — is not guaranteed to converge. However it is easy to see that if successive aproximations does converge, it converges a fixed point of the functional equations (16) and (17), and thus to a particular equilibrium of the dynamic game. We will show that the successive approximations algorithm can converge, but it will converge to different equilibria depending on the *equilibrium selection rule* we use to select an equilibrium in the investment "stage game" and also on the values from which the algorithm is intialized. We do not yet have a way to fully characterize *all* equilibria of this game, or bound the possible set of payoffs to consumers and the two firms. In contrast, the literature on the *Folk Theorem* in repeated games has succeeded in characterizing the set of possible equilibria and bounds on the set of equilibrium payoffs. We hope that eventually bounds and better characterization theorems can be estasblished for the class of dynamic games we consider here.

3 Solving the "End Game"

Under our assumptions the Markov process governing exogenous improvements in production technology has an absorbing state, where we assume (without loss of generality) that the minimum possible production cost is $c = 0$. This is also the absorbing state of the game, so that once costs of the firms reach zero, they can go no lower, and there is no forgetting or knowledge depreciation in our model that would ever cause them to go back up in the future.

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 leap frog 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 leap frogging 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.

$$\begin{aligned} v_0^i(0,0,0) &= r^i(0,0) + \beta P_1^{\sim i}(0,0,0)\phi(v_0^i(0,0,0), v_1^i(0,0,0)) \\ &+ \beta[1 - P_1^{\sim i}(0,0,0)]\phi(v_0^i(0,0,0), v_1^i(0,0,0)) \\ &= r^i(0,0) + \beta\phi(v_0^i(0,0,0), v_1^i(0,0,0)) \end{aligned} \tag{18}$$

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

$$P_1^{\sim i}(0,0,0) = \frac{\exp\{v_1^{\sim i}(0,0,0)/\eta\}}{\exp\{v_0^{\sim i}(0,0,0)/\eta\} + \exp\{v_1^{\sim i}(0,0,0)/\eta\}} \tag{19}$$

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

$$v_1^i(0,0,0) = v_0^i(0,0,0) - K(0), \tag{20}$$

which implies via equation (19) that

$$P_1^{\sim i}(0,0,0) = \frac{\exp\{-K(0)/\eta\}}{1 + \exp\{-K(0)/\eta\}}. \tag{21}$$

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

$$v_0^i(0,0,0) = r^i(0,0) + \beta\phi(v_0^i(0,0,0), v_0^i(0,0,0) - K(0)). \tag{22}$$

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 leap frog 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

$$\begin{aligned} v_0^1(c,0,0) &= r^1(c,0) + \beta\phi(v_0^1(c,0,0), v_1^1(c,0,0)) \\ v_1^1(c,0,0) &= r^1(c,0) - K(0) + \beta\phi(v_0^1(0,0,0), v_1^1(0,0,0)). \end{aligned} \tag{23}$$

Note that the solution for $v_1^1(c,0,0)$ in equation (23) 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 (22) and (20) above. Substituting the resulting solution for $v_1^1(c,0,0)$ into the first equation in (23) 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

$$P_1^2(c,0,0) = \frac{\exp\{-K(0)/\eta\}}{1+\exp\{-K(0)/\eta\}} \tag{24}$$

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 (21) 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,

$$v_1^1(0,c,0) = v_0^1(0,c,0) - K(0). \tag{25}$$

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

$$\begin{aligned} v_0^1(0,c,0) = r^1(0,c) \quad &+ \quad \beta P_1^2(0,c,0)\phi(v_0^1(0,0,0), v_0^1(0,0,0) - K(0)) \\ &+ \quad \beta[1 - P_1^2(0,c,0)]\phi(v_0^1(0,c,0), v_0^1(0,c,0) - K(0)). \end{aligned} \tag{26}$$

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

$$\begin{aligned} P_1^2(0,c,0) &= \frac{\exp\{v_1^2(0,c,0)/\eta\}}{\exp\{v_1^2(0,c,0)/\eta\}+\exp\{v_0^2(0,c,0)/\eta\}} \\ &= \frac{\exp\{v_1^1(c,0,0)/\eta\}}{\exp\{v_1^1(c,0,0)/\eta\}+\exp\{v_0^1(c,0,0)/\eta\}}, \end{aligned} \tag{27}$$

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 (23) above, these solutions can be substituted into equation (27) to obtain the probability that firm 2 invests, and then this probability can be substituted into equation (26) 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 (25).

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

$$\begin{aligned} v_0^1(c_1,c_2,0) = r^1(c_1,c_2) &+ \beta P_1^2(c_1,c_2,0)\phi(v_0^1(c_1,0,0),v_1^1(c_1,0,0)) \\ &+ \beta[1-P_1^2(c_1,c_2,0)]\phi(v_0^1(c_1,c_2,0),v_1^1(c_1,c_2,0)) \\ v_1^1(c_1,c_2,0) = r^1(c_1,c_2)-K(0) &+ \beta P_1^2(c_1,c_2,0)\phi(v_0^1(0,0,0),v_1^1(0,0,0)) \\ &+ \beta[1-P_1^2(c_1,c_2,0)]\phi(v_0^1(0,c_2,0),v_1^1(0,c_2,0)). \end{aligned} \tag{28}$$

Given the equation for $v_1^1(c_1,c_2,0)$ in equation (28) 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:

$$P_1^2(c_1,c_2,0) = \frac{\exp\{v_1^2(c_1,c_2,0)/\eta\}}{\exp\{v_0^2(c_1,c_2,0)/\eta\}+\exp\{v_1^2(c_1,c_2,0)/\eta\}}. \tag{29}$$

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 (28) for an arbitrary value of $P_1^2 \in [0,1]$. Substituting this into the equation for v_0^1, the top equation in (28), 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 (23) above). Using these values, we can write firm 1's probability of investing $P_1^1(c_1,c_2,0)$ as

$$P_1^1(c_1,c_2,0,P_1^2) = \frac{\exp\{v_1^1(c_1,c_2,0,P_1^2)/\eta\}}{\exp\{v_0^1(c_1,c_2,0,P_1^2)/\eta\} + \exp\{v_1^1(c_1,c_2,0,P_1^2)/\eta\}}. \tag{30}$$

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 (29) 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 probablity 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 (30) we obtain the following fixed point equation for firm 1's probability of investing

$$P_1^1 = \frac{\exp\{v_1^1(c_1,c_2,0,P_1^2(c_1,c_2,0,P_1^1))/\eta\}}{\exp\{v_0^1(c_1,c_2,0,P_1^2(c_1,c_2,0,P_1^1))/\eta\} + \exp\{v_1^1(c_1,c_2,0,P_1^2(c_1,c_2,0,P_1^1))/\eta\}}. \tag{31}$$

3.5 End Game Equilibrium Solutions

By Brouwer's fixed point theorem, at least one solution to the fixed point equation (31) exists. Further, when $\eta > 0$, the objects entering this equation (i.e. the value functions $v_0^1(c_1,c_2,0,P_1^2), v_1^1(c_1,c_2,0,P_1^2), v_0^2(c_1,c_2,0,P_1^1), v_1^2(c_1,$ and the logit choice probability function P_1^2 are all C^∞ functions of P_1^2 and P_1^1, and 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.

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

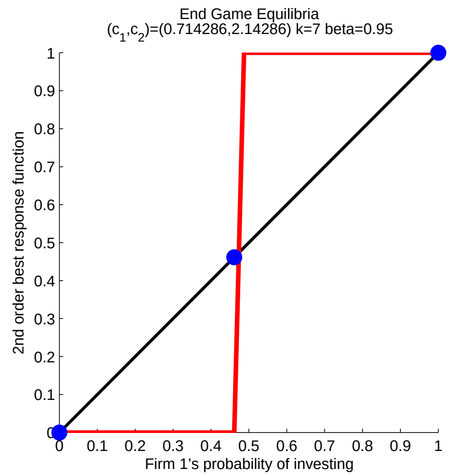

Figure 1 End Game Equilibria

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 leap frogging 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 (31) against the 45 degree line. We see that firm 1 is the low-cost leader with a substantially lower marginal cost of production than firm 2. In the mixed strategy equilibrium, firm 1 invests with probability 0.484, whereas the firm 2, the high cost follower, invests with probability 0.82. Thus, the high cost follower has a significantly higher chance of leap frogging 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	Firm 2: 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 0 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 leap frogged 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

$$V_1 = 0 + \beta V_1 \tag{32}$$

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

$$-K\pi_2 + (1-\pi_2)[\beta c_2/(1-\beta) - K] = 0, \tag{33}$$

or

$$\pi_2 = \frac{\beta c_2/(1-\beta) - K}{\beta c_2/(1-\beta)}. \tag{34}$$

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 emminently reasonable.

For firm 2 we have the following equation for V_2

$$V_2 = \pi_1(c_1 - c_2) + (1-\pi_1)(c_1 - c_2 + \beta V_2) \tag{35}$$

which implies that

$$V_2 = \frac{c_1 - c_2}{1 - \beta(1 - \pi_1)}. \tag{36}$$

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

$$V_2 = \pi_1(c_1 - c_2 - K) + (1 - \pi_1)(c_1 - c_2 + \beta c_1/(1 - \beta) - K). \tag{37}$$

Combining equations (36) and (37) 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. *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 leap frogging 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?

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 leap frogs its opponent, the game does not end, but rather the firms must anticipate additional leap frogging 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 leap frogging.

Specifically, we will focus on the following class of equilibria to the full game: *the cost follower invests whenever the state-of-the-art production cost c falls sufficiently below the marginal cost of the cost-leader to justify the investment cost $K(c)$, otherwise no investment occurs.* In order to "enforce" this equilibrium, we rely on a "credible threat" analogous to threats of a "price war" in the literature on tacit collusion in supergames. Specifically, 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. By simultaneously investing, the firms will move to the symmetric state $c_1 = c_2 = c$ where the equilibrium prescribes playing the "bad" mixed strategy equilibrium. This results in zero expected profits and this can be a sufficient "punishment" to deter the low cost leader from deviating from the implicit coordination that it should not invest when it is not its "turn."

In order to solve the full game, i.e. the pair of functional equations (16) and (17), it is helpful to rewrite them in the following way,

$$v_0^1(c_1,c_2,c) = r^1(c_1,c_2) + \beta\left[P_1^2(c_1,c_2,c)H^1(c_1,c,c) + (1-P_1^2(c_1,c_2,c))H^1(c_1,c_2,c)\right] \tag{38}$$

$$v_1^1(c_1,c_2,c) = r^1(c_1,c_2) - K(c) + \beta\left[P_1^2(c_1,c_2,c)H^1(c,c,c) + (1-P_1^2(c_1,c_2,c))H^1(c,c_2,c)\right] \tag{39}$$

where the function H^1 is given by

$$H^1(c_1,c_2,c) = p(c)\int_0^c \phi(v_0^1(c_1,c_2,c'),v_1^1(c_1,c_2,c'))f(c')dc' + (1-p(c))\phi(v_0^1(c_1,c_2,c),v_1^1(c_1,c_2,c)), \tag{40}$$

where $p(c)$ is the probability that a cost-reducing innovation will occur, and $f(c')$ is the density of the new (lower) cost of production under the current state-of-the-art conditional on an innovation having occurred.

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

$$v_0^2(c_1,c_2,c) = r^1(c_2,c_1) + \beta\left[P_1^1(c_1,c_2,c)H^2(c,c_2,c) + (1-P_1^1(c_1,c_2,c))H^2(c_1,c_2,c)\right] \tag{41}$$

$$v_1^2(c_1,c_2,c) = r^1(c_2,c_1) - K(c) + \beta\left[P_1^1(c_1,c_2,c)H^2(c,c,c) + (1-P_1^1(c_1,c,c))H^2(c_1,c,c)\right] \tag{42}$$

where the function H^2 is given by

$$H^2(c_1,c_2,c) = p(c)\int_0^c \phi(v_0^2(c_1,c_2,c'),v_1^2(c_1,c_2,c'))f(c')dc' + (1-p(c))\phi(v_0^2(c_1,c_2,c),v_1^2(c_1,c_2,c)), \tag{43}$$

If we set the arguments (c_1,c_2,c) to v_0 in equation (38) to (c,c,c), and similarly in equation (39) for v_1, we deduce that

$$v_1^1(c,c,c) = v_0^1(c,c,c) - K(c). \tag{44}$$

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

$$v_1^1(c,c_2,c) = v_0^1(c,c_2,c) - K(c). \tag{45}$$

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 (44) into equation (38) and use Newton's method to compute the unique fixed point $v_0^1(c,c,c)$. Similarly, we can solve for $v_0^1(0,c_2,0)$ by substituting equation (45) into equation (38) and solving. 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 (44) and (45) 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 (39). Then we substitute $v_1^1(c_1,c_2,c)$ into equation (38) and use Newton's method to compute $v_0^1(c_1,c_2,c)$.

Note that we assume that the integral term in equation (40) is "known". This is because the successive approximations solution algorithm is assumed to have computed $(v_0^1(c_1,c_2,c'), v_1^1(c_1,c_2,c'))$ for all $c' < c$ (although in actuality for a finite number of c' and other values of c' needed to numerically compute the integral in equation (40) are determined by interpolation).

Following the procedure we used to solve for equilibria in the end game, the set of equilibria for the investment "tage game" in each state (c_1,c_2,c) can be computed from the following fixed point equation

$$P_1^1 = \frac{\exp\{v_1^1(c_1,c_2,c,P_1^2(c_1,c_2,c,P_1^1))/\eta\}}{\exp\{v_0^1(c_1,c_2,c,P_1^2(c_1,c_2,c,P_1^1))/\eta\} + \exp\{v_1^1(c_1,c_2,c,P_1^2(c_1,c_2,c,P_1^1))/\eta\}}. \tag{46}$$

Depending on the rule we choose to select among the possible equilibria in each state (c_1,c_2,c) we can construct a 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 (38) and (39) 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 (41) and (42) 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

$$P_1^2(c_1,c_2,c,P_1^1) = \frac{\exp\{v_1^2(c_1,c_2,c,P_1^1)/\eta\}}{\exp\{v_0^2(c_1,c_2,c,P_1^1)/\eta\} + \exp\{v_1^2(c_1,c_2,c,P_1^1)/\eta\}}. \tag{47}$$

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

3. Using firm 2's best response probablity, P_1^2, calculate the value functions $(v_0^1(c_1,c_2,c,P_1^2),v_1^1(c_1,c_2,c,P_1^2))$ representing *firm 1's* values of not investing and investing in state (c_1,c_2,c), respectively, by solving the system (38) and (39).

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 (46).

A modified version of the algorithm involves solving only for the values off the end game (as done in section 3) and then solving for the values in the other states (c_1,c_2,c) for $c>0$ by the method of successive approximations, iterating on equations (38), (39), (41) and (42) until convergence is achieved. Any solution to this system constitutes a Markov-perfect Nash equilibrium to the dynamic investment/pricing game. Although successive approximations is not guaranteed to converge as is in the case when it is applied to solving Bellman equations in dynamic programming problems which are "single agent" problems (i.e. "games against nature"), when successive approximations does converge, it converges to an equilibrium of the game as we noted earlier.

We have found that simple successive approximations, using the modified approach where we do not force the investment actions for the two firms to constitute Nash equilibria at each state configuration (c_1,c_2,c) at *each iteration of the successive approximations algorithm*, will sometimes converge and sometimes not converge, depending on the initial conditions that we start out the algorithm and other details. When the successive approximations does converge, it converges to values and corresponding investment/pricing strategies that constitute mutual best responses at every (c_1,c_2,c) state point on the grid that we used to compute the problem, and these values satisfy each firm's Bellman equations (38), (39), (41) and (42). Thus the converged values implies equilibrium strategies for the full dynamic game, and these strategies are "perfect" in the sense that they are mutual best responses in every subgame and at all feasible states in the state space.

We have found that depending on how we initialize the successive approximations algorithm, when it does converge, it can converge to many different types of equilibria that have very different, interesting properties. These equilibria are generally of the pure strategy type, i.e. each firm has a unique best response to its opponent in each state (c_1,c_2,c). Some equilibria include the dynamic generalizations of "firm 1 invests" and "firm 2 invests" that we observed in the end game equilibria in section 3, so these equilibria

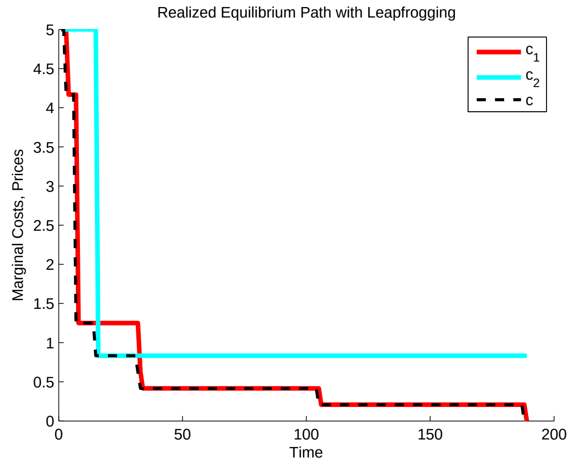

Figure 2 Equilibrium realization with leap frogging

lead to only one of the firms undertaking investments and setting a price equal to the initial price of their opponent. We refer to these as "monopoly equilibria" since one of the firms captures all of the benefits from the cost-reducing investments it undertakes and its opponent never challenges it by attempting to undertake 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. We conjecture that the investment paths in these equilibria are identical to the investment paths 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.

However there are also equilibria involving leap frogging behavior where the firms do compete dynamically by undertaking competing cost reducing investments. This causes 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.

However note from the figure 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 inital 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, leap frogging 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 leap frogged 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 leap frogs 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 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.[2] 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 leap frogs 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.

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 leap frogging and no investment, except for a single pre-

[2]Note that for these simulations we discretized the possible values that c could take on into 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 for the actual process we described in section 3 of the paper, the limiting value 0 would only be obtained asymptotically as $t \to \infty$. However as noted, when c_t becomes sufficiently small, the firms no longer have further incentive to invest. Thus, a more accurate simulation of the process (something we plan to do in future work) 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, the reader should interpret $c_t = 0$ as this small positive value of c_t at which further investment is no longer economic.

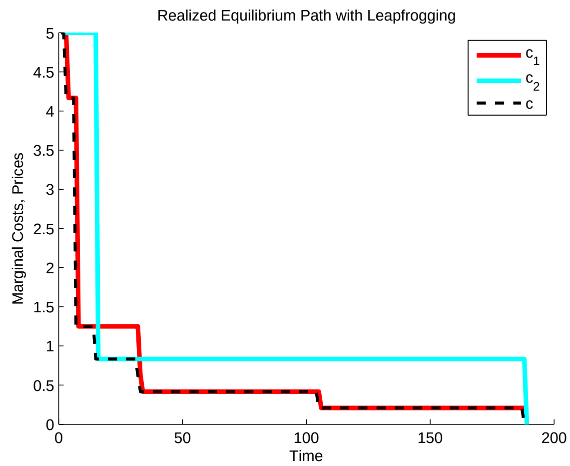

Figure 3 equilibrium realization with leap frogging

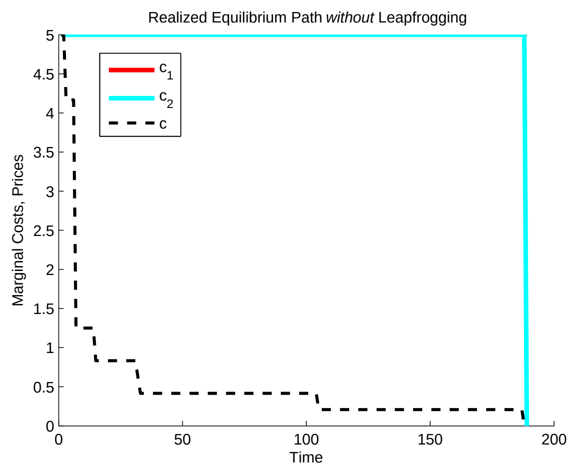

Figure 4 equilibrium realization *without* **leap frogging**

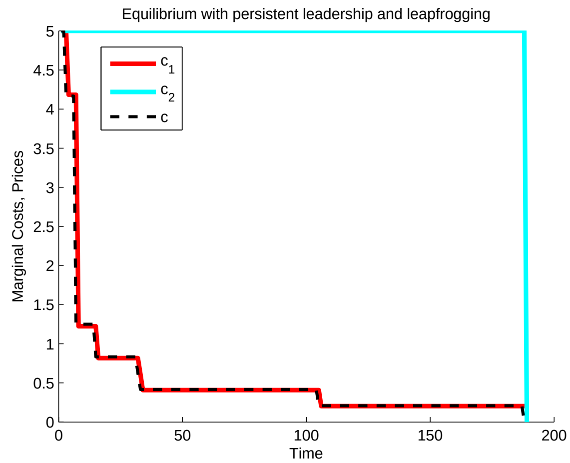

Figure 5 equilibrium realization with leap frogging and persistent leadership

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, leap-frogging firm 1 to attain a permanent position of 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

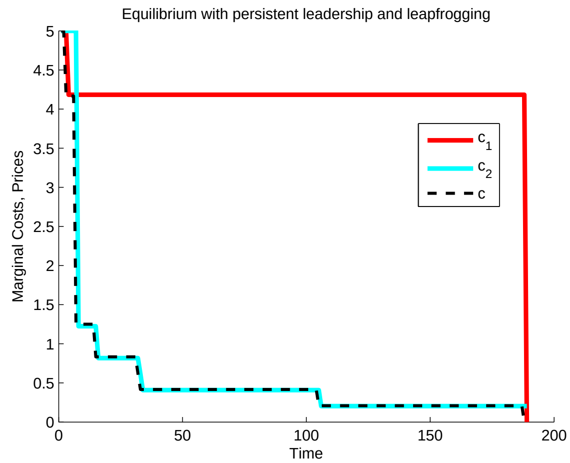

Figure 6 equilibrium realization with leap frogging and alternating leadership

firm 1 and not consumers. However unlike figure 4, when firm 2 finally invests and leap frogs 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 leap frogging and an alternating pattern of low cost leadership that results in more of the benefits of cost-reducing investments being passed on to 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 marinal 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 leap frog 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-reducinng 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 leap frogging 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

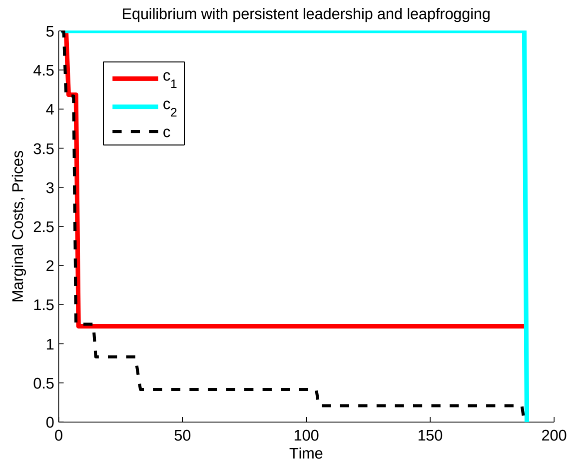

Figure 7 equilibrium realization with leap frogging and alternating leadership

undertake any further cost reducing investments after that, until in period $t = 190$ firm 2 invested and leap frogged 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 shoud 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 has a very simple, unique solution. Although we noted above that we have yet to systematically characterize the set of all equilibria to this model, and characterize the implied payoff sets (profits for firms 1 and 2 and prices to consumers), it is clear from figures 2 to 7 above, there is a very wide range of profits and prices that are consistent with equilibrium in 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 equilbria 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 leap frogging *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 leap frogs 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 leap frogs 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" frequently appears in the equilibrium simulations. By ths 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 leap frog 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 leap frog 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 conjecture in section 3 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: both firms earn zero profits and produce at a price equal to marginal cost.

However the static model begs the question of potential redundancy in production costs among the 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 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, and it provide the goods produced by this optimal plant 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

$$V(c_1,c) = \min\left[c_1 + \beta\int_0^c V(c_1,c')\pi(dc'|c), c_1 + K(c) + \beta\int_0^c V(c,c')\pi(dc'|c)\right]. \tag{48}$$

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 optimal threshold $\overline{c}(c_1)$ is the solution to the following equation

$$K(\overline{c}(c_1)) = \beta \int_0^{\overline{c}(c_1)} \left[V(c_1, c') - V(\overline{c}(c_1), c')\right] \pi(dc'|\overline{c}(c_1)). \tag{49}$$

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.

Comment: this section not yet complete. We intend to compare the overall efficiency of various duopoly equilibria to the social planning optimum and hope to show that various equilibria involve more efficient coordination in investment decisions between the two firms, and thus get closer (but not equal) to the social planning optimum. Duopoly will always involve some redundancy, and hence inefficiency, relative to the optimum that a social planner can achieve.

6 Conclusions

This draft is still preliminary and incomplete, so we hesitate to draw too many conclusions at this point. However several conclusions are possible from the work we have done so far. First, we have identified and resolved the *Bertrand invesment paradox* by showing that Bertrand duopolists do have incentive to undertake cost-reducing invesments. The cost-reducing investments can usually enable one of the firms to attain a temporary period of low cost leadership during which the discounted profits it can expect to earn are greater than the up-front fixed costs of undertaking the investment.

Our paper is not the first to establish the possibility of leap frogging 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 compeition with cost-reducing investments in a framework similar to our's. 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 leap-frogging equilibria in which investments occur in every period, but with the two firms alternating in their investments. The second main difference is the specification of demand, where Giovannetti assumed that the demand for goods is given by a constant elasticity of substitution demand curve rather than deriving aggregate demand from micro aggregation of individual discrete purchase decisions as we have done in our model. The stochastic nature of technological progress that is captured in our model leads to equilibria where there are long periods where there are no technological improvements and thus no investment by either firm, punctuated by technological break-throughs that sometimes induce one or both of the firms to invest in the state-of-art production machinery, thereby precipitating a price war.

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 econommic), thereby displacing its rival to attain a permanent low cost leadership position.

We also refer the reader to the very important paper by Goettler and Gordon (2009) that studies leap-frogging 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 comsumers 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. Interestingly, despite the substantial additional complexity of their modeling framework, Goettler and Gordon claim that there model has a *unique equilibrium.* It is of interest to us to better understand what features lead to the huge plethora of equilibria in our much simpler framework, and which features of the Goettler and Gordon framework lead to a unique equilibrium.[3]

We were surprised by how complex are the various tytpes of equilibrium behavior that can emerge from this simple model. Unlike the static Bertrand model of price competition where there is a single, simple, unqiue and fully efficient equilibrium outcome — the Walrasian competitive outcome — in the simple dynamic extension we have considered where firms compete over both price and investment strategies, there appear to be a vast multiplicity of different equilibria, virtually all of which are inefficient. Some of these equilibria can result in outcomes that are very bad for consumers even though the duopolists are never colluding and behaving as Bertrand price competitors in every period. We have more work to do to explore and characterize the set of equilibria in this model, and to better understand the dynamics of price and investment competition when the two firms are producing goods that are not perfect substitutes. We would also like to add capacity constraints to the model and understand whether the equilibria of this extended model would exhibit the result discovered by Kreps and Scheinkman (1983), namely, that capacity investment followed by Bertrand price competition yields an outcome identical to the Cournot-Nash equilibrium in a model where firms chooses quantities only.

A final contribution is that we provide a new interpretaion for price wars. In our model price wars occur when a high cost firm leap frogs 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 leap frogging 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 our results are simply theoretical curiosums, or whether this framework can be extended and the issues of multiple equilibria be addressed in a satisfactory way that this work might yield useful practical insights and new tools for applied Industrial Organization.

[3]Goettler and Gordon appealed to the *Unique Investment Choice* (UIC) admissibility criterion of Doraszelski and Satterthwaite (2010) to establish the uniqueness of equilibrium in their model.

7 Appendix: Proof of Lemma 1

Lemma 3.1 *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 contrudiction. We conclude then that $\pi_1>\pi_2$. □

References

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

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

[3] 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.

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

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

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

[7] Ericson, Richard and Ariel Pakes (1995) "Markov-perfect Industry Dynamics: A Framework for Empirical Work" *Review of Economic Studies* **62** 53–82.

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

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

[10] 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).

[11] 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.

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

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

[14] Riordan, Michael (2010) "Quality Competition and Multiple Equilibria" working paper, Department of Economics, Columbia University.

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

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