DISCUSSION PAPER SERIES

No. 4249

GO FOR BROKE OR PLAY IT SAFE? DYNAMIC COMPETITION WITH CHOICE OF VARIANCE

Axel Anderson and Luís M B Cabral

INDUSTRIAL ORGANIZATION

Centre for Economic Policy Research

www.cepr.org

Available online at: www.cepr.org/pubs/dps/DP4249.asp

ISSN 0265-8003

GO FOR BROKE OR PLAY IT SAFE? DYNAMIC COMPETITION WITH CHOICE OF VARIANCE

Axel Anderson, Georgetown University
Luís M B Cabral, New York University and CEPR

Discussion Paper No. 4249
February 2004

Centre for Economic Policy Research
90–98 Goswell Rd, London EC1V 7RR, UK
Tel: (44 20) 7878 2900, Fax: (44 20) 7878 2999
Email: cepr@cepr.org, Website: www.cepr.org

CEPR Discussion Paper No. 4249

February 2004

ABSTRACT

Go For Broke or Play it Safe? Dynamic Competition with Choice of Variance

We consider a differential game in which the joint choices of the two players influence the variance, but not the mean, of the one-dimensional state variable. We interpret this state variable as a summary of how far 'ahead' player 1 is in the game. At each moment in time, players receive a flow pay-off which is a continuous, monotonic and bounded function of the state variable. We show that a Markov Perfect Equilibrium exists and has the property that patient players chose to play it safe when sufficiently ahead and to take risks when sufficiently behind. We also provide a simple condition that implies both players choose risky strategies when neither one is too far ahead, a situation that ensures a dominant player emerges 'quickly'.

JEL Classification: C70 and L10
Keywords: differential games and R&D competition

Axel Anderson
Department of Economics
Georgetown University
580 ICC, 37th & O Streets, NW
Washington, DC 20057
USA
Email: aza@georgetown.edu

Luís M B Cabral
Stern School of Business
New York University
44 West 4th Street
New York, NY 10012-1126
USA
Tel: (1 212) 998 0858
Fax: (1 212) 995 4218
Email: lcabral@stern.nyu.edu

For further Discussion Papers by this author see:
www.cepr.org/pubs/new-dps/dplist.asp?authorid=160260

For further Discussion Papers by this author see:
www.cepr.org/pubs/new-dps/dplist.asp?authorid=114053

Submitted 15 December 2003

1 Introduction

Many industries, such as the pharmaceutical and the microprocessor industries, can be explained by a dynamic game with the following features: (a) in each period, each firm is characterized by the value of its product; (b) in each period, each firm's profit is a function of all firms' product values; (c) by investing resources into R&D, a firm stochastically improves the future quality of its product.

Several authors have looked at games with this basic structure; see Budd et al. (1993), Ericson and Pakes (1995), Fershtman and Pakes (2000), Hörner (1999). One feature that is common to all of these models is that firm strategies consist of choosing the level of R&D expenditures. Although this is unquestionably an important part of a firm's strategy, it is not the only one. In fact, there are cases when, arguably, the *type* of R&D expenditures is more important than its *level*. For example, the analysis of the microprocessor industry developed in Khanna and Iansiti (1997) suggests that firms are frequently financially constrained and work from a fixed budget. By contrast, they have significant flexibility as to what kind of R&D "strategy" to follow. It's not how much to spend, rather how to spend it. In particular, firms can choose R&D "strategies" with various degrees of risk.

In this paper, we study the dynamics of R&D competition when firms choose the variance of R&D outcomes. We consider a differential game with two players (firms). At each moment in time, each player's position is given by a real number q_i (e.g., the quality of its product). Player i receives a payoff flow given by $\pi(q_i - q_j)$, a strictly increasing function. The player's position, q_i, evolves according to a Wiener process with mean μ, which is exogenously given, and variance σ_i, which is chosen by firm i. Our goal is to characterize the Markov equilibrium of this game. In words, we want to understand when players choose safer or riskier R&D strategies as a function of their relative position in the race.

Strategic choice of risk (variance) plays an important role in sports. For example, in the fourth quarter of a (American) football game, the team that is behind calls more passing plays, while the team that is ahead runs the ball. Toward the end of a hockey game the team that is behind pulls their goalie, in favor of an additional offensive player, while the team that is ahead substitutes in more defensive players. In both of these situations the team that is behind is opting for a high variance strategy, while the team that is ahead is opting for a low variance strategy. As the saying goes, "if you're behind you have nothing to lose."

There is no reason to suspect *a priori* that such reasoning should carry over to infinite horizon games, as it seems to be the end game effect that drives the intuition. However, we think an infinite horizon is a better description of real-world oligopoly competition. So, we ask, do players still adopt a high risk strategy when behind and a low risk strategy when ahead in an infinite horizon game?

Consider first the case when players are very impatient. In this case, the value function is approximately equal to the flow profit function. In general, by Jensen's inequality, risk choices are determined by the curvature of the value function. Thus, risk choice for very impatient players follows from the local curvature of the profit function. Thus, for very impatient players the answer to our question is unsurprising: choice of variance is entirely dependent on the local curvature of the flow profit function.

Consider now the case of very patient players. We are able to get a very strong result: In every Markov Perfect Equilibrium, patient players chose to play it safe when ahead and to take risks when behind. Specifically, there exists a value x^* such that players choose low variance if $x > x^*$ and high variance if $x < x^*$, where x is the difference between the player's position and its rival's (the state of the game). Note that this result follows from the monotonicity and boundedness of the flow profit function, i.e. we need not make any assumptions about the local curvature of this profit function.

The main thrust of our results is that, when players are very patient, the second derivative of the value function is negatively related to the current payoff level. Specifically, a lagging player receives a low payoff and has a convex value function; whereas a leading player receives a high payoff and has a concave value function. Once this has been established, equilibrium strategies follow from Jensen's inequality.

As part of our equilibrium characterization, we provide a simple formula for x^*, the threshold dividing the low-risk and high-risk regions. When $x^* > 0$, both players choose risky strategies in states where x is close to zero. In other words, starting from a situation where players are more or less even, a dominant player will emerge "quickly." Previous research (Budd et al., 1993; Cabral, 2002; Cabral and Riordan, 1994) has characterized dynamic games featuring increasing dominance, the property whereby the gap between leader and follower tends to increase in expected value, resulting in an asymmetric outcome. In our model, player's choices cannot effect the drift of the state variable, so that the gap between leader and follower must remain constant in expected terms. Despite this restriction, our result shares the feature that asymmetry tends to emerge rapidly.

Technically we study a (very simple) stochastic differential game in which the agent's choices affect the variance of some state variable. The general equilibrium existence theory for such variance choice games is not well developed. Thus, it is not surprising that existence questions have been for the most part dodged in economic applications of stochastic differential games with endogenous variance. Luckily, our model is simple enough that we can apply a result from Harris (1993) to establish existence. For other examples of stochastic differential games in economics in which strategies affect the variance of the evolution of the state variable, see Bergemann and Valimaki (2000) and Bolton and Harris (1999).

In this paper, we characterize the equilibrium value functions and optimal strategies, and *then* investigate the behavior of these value functions and strategies as players become infinitely patient. An alternative approach taken recently by Bergemann and Valimkai (2002) and Bolton and Harris (2001) is to consider the MPE of the *undiscounted* gamedirectly. Dutta (1991) establishes that in the limit the equilibrium strategies and payoffs of the discounted game must converge to those of the undiscounted games when the strong long run average payoff is used. Thus, we could have considered the limiting equilibrium directly. However, given the simplicity of our model we feel that characterizing (and proving existence of) the MPE in the discounted game, and then taking the limit is the right approach. Note that in the two referenced papers the models are more complex, so that they *must* investigate the undiscounted game directly to make reasonable progress.

The paper is organized as follows. In Section 2, we present the basic elements of the model: a continuous time, differential game with a one-dimensional state space. In Section 3, we show that an equilibrium exists for this game. In Section 4, we derive a series of results regarding the value function. We use these results in Section 5, where we consider the limit case when players are very patient. This section includes the main results in the paper regarding the characterization of equilibrium strategies: the result that leaders tend to choose safe strategies and followers risky ones; and the result that, under some circumstances, the resulting dynamic system will quickly move away from symmetry. In Section 6, we consider a series of examples that illustrate the main results. Section 7 concludes the paper.

2 The model

In this section, we present the basic elements of the game we will study. Consider the following two player stochastic continuous time (differential) game.[1] At each instant in time, player $i \in \{1,2\}$ chooses $\sigma_i \in [\underline{\sigma}, \bar{\sigma}]$, the variance of its motion in a one-dimensional state space. The state of the game at time t is summarized by $x(t) \in \mathbb{R}$. Conditional on the joint choices of the two players, x evolves according to the following Ito process:[2]

$$dx(t) = \sqrt{2(\sigma_1 + \sigma_2)}dz \tag{1}$$

where dz is the increment of a Wiener process. Let $\pi(x)$ denote the flow profits that player 1 receives, while player 2 receives profit flows $\pi(-x)$. We assume that π is continuous, bounded and strictly increasing. Finally, we assume that players discount future profits at rate r.

We will be considering Markov equilibria. A *Markov strategy* for player i is a map $\sigma_i : (-\infty, +\infty) \mapsto [\underline{\sigma}, \bar{\sigma}]$. Given a strategy pair $\sigma = (\sigma_1, \sigma_2)$, the payoffs for the players are:

$$U_1(x,\sigma) \equiv E\left[\int_0^\infty e^{-rt}\pi(x(t))dt \mid x(0), (\sigma_1, \sigma_2)\right]$$
$$U_2(x,\sigma) \equiv E\left[\int_0^\infty e^{-rt}\pi(-x(t))dt \mid x(0), (\sigma_1, \sigma_2)\right]$$

In summary, we have a symmetric game on a one-dimensional space, $x(t) \in \mathbb{R}$. Player 1 benefits from a high value of x, whereas player 2 prefers a low value. The expected motion of x is zero, but its variance depends on the players' choices. Specif-

[1]See Harris (1993) for a very thorough treatment of one-dimensional stochastic differential games.
[2]A good (accessible) reference for basic stochastic control is Dixit and Pindyck (1994). For a (much) more technical reference see Øksendal (1998).

ically, at each point x of the state space, each player chooses either a high or a low value of variance, with the system variance equal to the sum of the players' choices.

3 Equilibrium existence

In this section, we show that an equilibrium exists to the game introduced in the previous section. Fix a Markov strategy σ_2 for player 2. Then player 1's Markov best response solves:

$$U_1^*(x;\sigma_2) = \sup_{\sigma_1(x)} U_1(x,\sigma_1,\sigma_2)$$

Assume that an optimal Markov best response $\sigma_1^*(\sigma_2)$ exists, then Theorem 11.2.3 in Øksendal (1998) establishes that player 1 can achieve as high a payoff using σ_1^* as he can using any (measurable) strategy. That is, a Markov strategy is a best response to a Markov strategy.

Define the Hamilton-Jacobi-Bellman equation (HJB) as (via Ito's Lemma):

$$rV_1(x;\sigma_2) = \max_{\sigma_1(x)} \left[\pi(x) + (\sigma_1(x)+\sigma_2(x))V_1''(x;\sigma_2)\right]$$

The following lemma allows us to focus our attention on the HJB equation.

Lemma 1 *If we restrict attention to C^2 solutions then $U_1^*(x;\sigma_2) = V_1(x;\sigma_2)$. More precisely,*

i. *If σ_1^* exists and $U_1^* \in C^2$ and is bounded then U_1^* solves HJB.*

ii. *If $V_1 \in C^2$ solves HJB, then $V_1 = U_1^*$ and any solution to HJB is a Markov best response.*

Proof: It is straightforward to verify that our assumptions satisfy Theorems 11.2.1 and 11.2.2 in Øksendal (1998). Applying Theorem 11.2.1 yields part **i**, while applying Theorem 11.2.2 yields part **ii**. ■

A Markov strategy σ_i is *simple* if any bounded interval $(a,b) \subset \mathbb{R}$ admits a partition $\{y_i, i=0,\ldots,n\}, a = y_0 < y_1 < \cdots y_n = b$ such that σ_i is constant on each subinterval (y_i, y_{i+1}). A MPE is simple if the equilibrium strategies are simple. We will focus our attention on simple MPE, a restriction we justify with the following proposition.

Proposition 1 *A simple MPE exists.*

Proof: The proof relies heavily on Harris (1993). First we show that any MPE in which the value function is C^2 must be simple. Then we show that a C^2 MPE exists.

Claim 1 *Assume a MPE with C^2 value functions, then the equilibrium is simple.*

Proof: We will analyze everything from the perspective of player 1, the arguments for player 2 are symmetric. Note that optimality requires: $\sigma_1^*(x) = \underline{\sigma}$ whenever $V_1''(x) < 0$ and $\sigma_1^*(x) = \overline{\sigma}$ whenever $V_1''(x) > 0$. Since V_1'' is continuous, the posited equilibrium will be simple, unless $V_1'' = 0$ on some interval. Now assume $V_1'' = 0$ on some interval. This implies V_1 is also constant on this interval. But given $V_1'' = 0$, equation (2) becomes $rV_1(x) = \pi(x)$. Since π is strictly increasing, this is a contradiction. ■

Claim 2 *A C^2 MPE exists.*

Proof: We wish to apply Theorem 11.7 from Harris (1993). To do so requires we analyze a static two player game in which the player's choose $\sigma_i \in [\underline{\sigma}, \overline{\sigma}]$ as before, but payoffs are:

$$\lambda_i'' + \frac{\pi_i(x) - r\lambda_i}{\sigma_1 + \sigma_2}$$

where $\lambda, \lambda'' \in \mathbb{R}^2$. Following Harris, let $\overline{ne}(x, \lambda, \lambda'')$ be the set of Nash equilibrium payoff vectors for this static game. Then by Theorem 6.6 in Harris, a MPE to the original dynamic game will exist if, $\overline{ne}(x, \lambda, \lambda'')$ is nonempty and convex for all λ, λ'', and x.

Note that for all (λ, x) such that $\pi_i(x) \neq r\lambda_i$ for all i, there is a unique equilibrium, so $\overline{ne}(x, \lambda, \lambda')$ is nonempty and trivially convex. If $\pi_i(x) = r\lambda_i$ for all i, any allowable σ is an equilibrium, and all equilibria have the same payoff vector. Finally, consider the case in which $\pi_1(x) > r\lambda_1$ and $\pi_2(x) = r\lambda_2$ (WLOG this is the only remaining case to consider). In this case, the set of Nash equilibria is $\sigma_1 = \underline{\sigma}, \sigma_2 \in [\underline{\sigma}, \overline{\sigma}]$. The payoff vector for player 2 is the same in every equilibrium. The set of payoff vectors for player 1 is the convex set:

$$\left[\frac{\pi_1(x) - r\lambda_1}{\underline{\sigma} + \overline{\sigma}}, \frac{\pi_1(x) - r\lambda_1}{\underline{\sigma} + \underline{\sigma}}\right]$$

and thus Theorem 6.6 in Harris (1993) applies. ■

This concludes the proof of the main result. ■

4 Curvature of the value function

In this section we investigate various characteristics of the value function, in particular its curvature. Some of the results will be used later when characterizing limit equilibrium behavior (as the discount rate goes to zero).

A simple strategy pair partitions $[-\infty, \infty]$ into open intervals such that σ is constant in each interval. Consider one of these intervals and abbreviate $V(x) \equiv V(x, \sigma) \forall x$ on this interval. Then the HJB equation implies that:

$$rV(x) = \pi(x) + (\sigma_1 + \sigma_2)V''(x). \tag{2}$$

The general solution to this differential equation is:[3]

$$V_G(x) = ae^{-\alpha x} + be^{\alpha x} + \tilde{\pi}(x; \alpha),$$

where

$$\tilde{\pi}(x; \alpha) \equiv \frac{1}{\gamma} \left[\int_{-\infty}^{x} e^{\alpha(s-x)} \pi(s) ds + \int_{x}^{\infty} e^{\alpha(x-s)} \pi(s) ds \right],$$

a and b are undetermined coefficients, $\gamma \equiv 2\sqrt{\sigma r} > 0$, $\alpha \equiv \sqrt{r/\sigma} > 0$, and $\sigma = \sigma_1 + \sigma_2$.

[3]To verify this solution, note that it must satisfy

$$V_G(x) = \frac{\pi(x) + \sigma V_G''(x)}{r} = \frac{\gamma - 2\alpha\sigma}{\gamma} \pi(x) + \frac{\alpha^2 \sigma}{r} \left[ae^{-\alpha x} + be^{\alpha x} + \tilde{\pi}(x) \right],$$

which is true iff $\gamma - 2\alpha\sigma = 0$ and $\alpha^2 \sigma / r = 1$. These two equations are satisfied for the given γ and α. Further, it must be the case that $\tilde{\pi}$ is bounded. To see this, take the first term in brackets and simplify:

$$\begin{aligned} \int_{-\infty}^{x} e^{\alpha(s-x)} \pi(s) ds &= e^{-\alpha x} \int_{-\infty}^{x} e^{s} \pi(s) ds \\ &\leq e^{-\alpha x} \int_{-\infty}^{x} e^{s} \bar{\pi} ds \\ &= \frac{\bar{\pi}}{\alpha}, \end{aligned}$$

where the inequality follows from the assumption that π is bounded ($\pi(s) < \bar{\pi}$).

Direct computation yields $\tilde{\pi}(x;\alpha)$ equal to the total expected discounted value of profits starting in state x if no one switched projects. In fact, given that $x(t)$ is an Ito process, the distribution over future values x at time t starting from x at time $t=0$ is normal with mean x and variance $2\sigma t$. Thus, we can directly compute

$$
\begin{aligned}
&E\left[\int_0^\infty e^{-rt}\pi(x_t)dt|x(0)=x,\sigma\right]= \\
&=\int_{-\infty}^\infty\int_0^\infty e^{-rt}\pi(s)(4\pi\sigma t)^{-\frac{1}{2}}e^{-\frac{(x-s)^2}{4\sigma t}}dtds \\
&=\int_{-\infty}^x\left[\int_0^\infty(4\pi\sigma t)^{-\frac{1}{2}}e^{-rt-\frac{(x-s)^2}{4\sigma t}}dt\right]\pi(s)ds+\int_x^\infty\left[\int_0^\infty(4\pi\sigma t)^{-\frac{1}{2}}e^{-rt-\frac{(x-s)^2}{4\sigma t}}dt\right]\pi(s)ds.
\end{aligned}
$$

Evaluating the bracketed expressions yields the desired result.

For an intuition of why this must be so, note that the value function is bounded, and must always satisfy the general form of V_G. If no one switched projects, the value function has the same form for all values of x, and yet $V_G(x)$ is unbounded unless $a=0$ and $b=0$. Thus, if no one switched projects, $V(x)=\tilde{\pi}(x\alpha)$. For this reason, we shall refer to $\tilde{\pi}$ as the *fundamental value*.

If $\tilde{\pi}$ is the value when no one switches projects, then the other two terms must be the sum of a player's option to switch projects, and the effect of that player's rival switching projects (we shall refer to these terms as the *induced values*).

Rearranging equation (2) yields the following convexity/concavity lemma.

Lemma 2 *The value function is strictly convex (concave) iff $rV(x)>\pi(x)$ ($rV(x)<\pi(x)$).*

It will often prove convenient to work with $r\tilde{\pi}$. Multiplying and dividing $r\tilde{\pi}$ by α yields:

$$r\tilde{\pi}(x;\alpha)=\frac{1}{2}\left[\int_{-\infty}^x\alpha e^{\alpha(s-x)}\pi(s)ds+\int_x^\infty\alpha e^{\alpha(x-s)}\pi(s)ds\right] \tag{3}$$

Since, $\gamma\alpha=2r$, so that $r/(\gamma\alpha)=1/2$.

Lemma 3 *If π is C^k, then V_G is C^{k+2}. Thus, since π is continuous, V_G is C^2.*

Proof: The option and induced value terms are C^∞, thus we need only be concerned with $\tilde{\pi}$. We first show that $\tilde{\pi}$ is C^2. Since, $\tilde{\pi}$ is the integral of a continuous function, it is continuous and differentiable. Taking the first derivative of (3) yields:

$$r\tilde{\pi}'(x;\alpha) = \frac{1}{2}\left[\int_{-\infty}^{x} -\alpha^2 e^{\alpha(s-x)}\pi(s)ds + \int_{x}^{\infty} \alpha^2 e^{\alpha(x-s)}\pi(s)ds\right], \tag{4}$$

which is again continuous and differentiable by the continuity of π.

Differentiating a second time yields:

$$\begin{aligned} r\tilde{\pi}''(x;\alpha) &= \frac{1}{2}\left[\int_{-\infty}^{x} \alpha^3 e^{\alpha(s-x)}\pi(s)ds + \int_{x}^{\infty} \alpha^3 e^{\alpha(x-s)}\pi(s)ds - 2\alpha^2\pi(x)\right] \\ &= \alpha^2(r\tilde{\pi}(x) - \pi(x)), \end{aligned} \tag{5}$$

which is again continuous, so $\tilde{\pi}$ is C^2. Finally, the result follows from equation (5) by induction. ■

The following lemma provides a simple characterization for the second derivative of the fundamental value.

Lemma 4 *The fundamental value is strictly convex (concave)* iff $r\tilde{\pi}(x;\alpha) > \pi(x)$ *(*$\tilde{\pi}(x\alpha) < \pi(x)$*).*

Proof: The result follows immediately from equation (5). ■

5 Result with high patience

The player's attitudes toward risk will be influenced by the curvature of the flow profit function π. If the players are very impatient (high r), then the local curvature of π will weigh heavily in their decision making. In fact, we could show that if π is convex (concave) in a neighborhood of x, then player 1 will choose the risky (safe) process at x if r is above a certain threshold. We do not find this surprising or particularly interesting. Instead we are interested in what happens for low r. This is not *a priori* obvious given our minimal assumptions on π. First we need to understand the behavior of the fundamental value $\tilde{\pi}$ for low r.

5.1 The value function in the limit

Let

$$\overline{\pi} = \lim_{s\to\infty} \pi(s), \qquad \text{and} \qquad \underline{\pi} = \lim_{s\to-\infty} \pi(s).$$

The following lemma establishes the limit behavior of the fundamental value, given very weak assumptions on profits.

Lemma 5 *(a)* $\lim_{r\to 0} r\tilde{\pi}(x) = (\overline{\pi} + \underline{\pi})/2$
(b) $\lim_{r\to 0} \sqrt{r\sigma}\tilde{\pi}'(x;\alpha) = (\overline{\pi} - \underline{\pi})/2$
(c) $\lim_{r\to 0} \sigma\tilde{\pi}''(x;\alpha) = (\overline{\pi} + \underline{\pi})/2 - \pi(x)$

Proof: This result follows easily from the following facts:

$$\lim_{r\to 0} \int_{-\infty}^{x} \alpha e^{\alpha(s-x)}\pi(s)ds = \underline{\pi} \quad \text{and} \quad \lim_{r\to 0} \int_{x}^{\infty} \alpha e^{\alpha(x-s)}\pi(s)ds = \overline{\pi}$$

We will only establish the first limit; the second is established similarly. Integrating by parts (π monotonic implies $\pi'(x)$ exists a.e.) we find:

$$\begin{aligned}
\int_{-\infty}^{x} \alpha e^{\alpha(s-x)}\pi(s)ds &= \left[e^{\alpha(s-x)}\pi(s)\right]_{-\infty}^{x} - \int_{-\infty}^{x} e^{\alpha(s-x)}\pi'(s)ds \\
&= \pi(x) - \lim_{s\to-\infty} e^{\alpha(s-x)}\pi(s) - \int_{-\infty}^{x} e^{\alpha(s-x)}\pi'(s)ds \\
&= \pi(x) - \lim_{s\to-\infty} e^{\alpha(s-x)}\underline{\pi} - \int_{-\infty}^{x} e^{\alpha(s-x)}\pi'(s)ds \\
&= \pi(x) - \int_{-\infty}^{x} e^{\alpha(s-x)}\pi'(s)ds.
\end{aligned}$$

Note that we are done if we can establish that $\lim_{r\to 0} \int_{-\infty}^{x} e^{\alpha(s-x)}\pi'(s)ds = \pi(x) - \underline{\pi}$. Since, $\alpha = \sqrt{r/\sigma}$, this is equivalent to establishing the limit as $\alpha \to 0$. Taking this limit we find:

$$\begin{aligned}
\lim_{\alpha\to 0} \int_{-\infty}^{x} e^{\alpha(s-x)}\pi'(s)ds &= \int_{-\infty}^{x} \lim_{\alpha\to 0} e^{\alpha(s-x)}\pi'(s)ds \\
&= \int_{-\infty}^{x} \pi'(s)ds \\
&= \pi(x) - \lim_{s\to-\infty} \pi(s) \\
&= \pi(x) - \underline{\pi}.
\end{aligned}$$

Given these established limits, the results follow directly from (3), (4), and (5), respectively. ■

Lemma 6 *$rV''(x) \to r\tilde{\pi}''(x;\alpha)$ a.e. as $r \to 0$.*

Proof: Recall,

$$rV''(x) = r\tilde{\pi}''\alpha^2(b(r)e^{\alpha x} + a(r)e^{-\alpha x})$$

for some $a(r)$, $b(r)$. Since, $\alpha^2 = r/\sigma$, if $rV''(x)$ does not converge to $r\tilde{\pi}''(x;\alpha)$ then $(b(r)e^{\alpha x} + a(r)e^{-\alpha x})$ is unbounded, but this implies that $rV(x)$ is unbounded, a violation. ■

5.2 Risk choice in the limit

Given the results in the last section we can conclude almost immediately that players will pursue the safe process when ahead and the risky process when behind.

Proposition 2 *Let x^* solve $\pi(x^*) = (\overline{\pi}+\underline{\pi})/2$. There exists an $\varepsilon(r)$, with $\lim_{r\to 0}\varepsilon(r) = 0$, such that $\sigma_1 = \overline{\sigma}$ for all $x < x^* - \varepsilon(r)$ and $\sigma_1 = \underline{\sigma}$ for all $x > x^* + \varepsilon(r)$.*

Proof: By the first part of Lemma 5, if $x < x^* - \varepsilon(r)$ then $\pi(x) < (\overline{\pi} + \underline{\pi})/2 = \lim\limits_{r\to 0} r\tilde{\pi}(x)$. By Lemma 4, $\pi(x) < r\tilde{\pi}(x)$ implies that $\tilde{\pi}$ is convex. Finally, by Lemma 6 if $\tilde{\pi}$ is convex then V is convex. ■

In words, Proposition 2 states that, in the limit, there will be a threshold dividing the low risk and the high-risk regions of the strategy space. For low x (follower), high risk is the equilibrium strategy, whereas for high x (leader) low risk is better. Figure 2 illustrates this. As $r \to 0$, V flattens out. Since π is increasing, there exits an x^* such that $\pi > rV$ iff $x > x^*$. In the region where $\pi > rV$, that is, for $x > x^*$, the optimal strategy is $\sigma = \underline{\sigma}$. For $x < x^*$ the opposite is true.

One way to understand the result is to recall that the profit function is increasing and bounded. This implies that, for very high values of x, the value function must be concave, whereas for very low values it must be convex. What the $r \to 0$ limit does

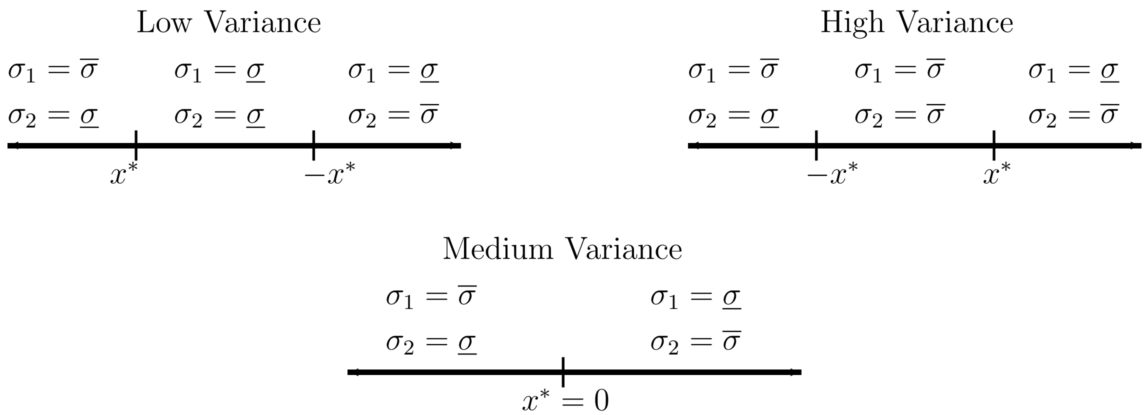

Figure 1: The Three MPE for Low r. If the critical threshold x^* is less than zero, then in the interval $(x^*, -x^*)$ both players choose $\sigma_i = \underline{\sigma}$ (Low Variance case); if $x^* > 0$, then both players choose $\sigma_i = \overline{\sigma}$ in $(x^*, -x^*)$ (High Variance case); finally, if $x^* = 0$, then total variance is constant in the entire space and equal to $\underline{\sigma} + \overline{\sigma}$ (Medium Variance case).

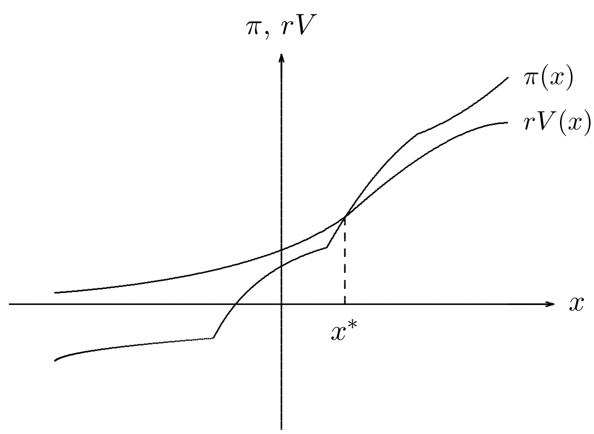

Figure 2: For low values of r, $V(x)$ is concave if and only if $\pi(x) > rV(x)$, which happens to the right of x^*.

is to make the intermediate values of x irrelevant, so that only the convex / concave regions remain.

Intuitively, Proposition 2 can also be understood with reference to Lemma 2. The value function is convex if and only if $\pi(x) < rV(x)$. In words, if current profit is less than average discounted payoff, then "things can only get better." If things are going to get better it's because the discounted payoff in neighboring states is better than in the current state; and so a high risk strategy is optimal, insofar as it will move us away from the current state. If we show that $\pi(x) < rV(x)$ for a laggard then we're home: a laggard wants to choose a high risk strategy. So, instead of the sports intuition that a laggard has "nothing to lose," we show that a laggard has only to gain from moving away from the current state, and does so by choosing a high risk strategy.

Note that in the limit the unique equilibria can only be one of the three types, pictured in Figure 1. The knife edged case of $x^* = 0$ is straightforward. Note that in this case $\sigma = \underline{\sigma} + \overline{\sigma}$, which we call the *Medium Variance* case. When $x^* \neq 0$, the state space is divided into three intervals. When $x^* > 0$, each player chooses high variance ($\sigma_i = \overline{\sigma}$) around $x = 0$, so we call this the *High Variance* case. When $x^* < 0$, we again have medium variance at the extremes, but low variance in a neighborhood of $x = 0$, so we call this the *Low Variance* case. This definitions allow us to state the following simple corollary to Proposition 2.

Corollary 1 *The High, Medium and Low Variance cases obtain as* $\pi(0)$ *is lower than, equal to, or greter than* $(\overline{\pi} + \underline{\pi})/2$*, respectively.*

To illustrate these results we graphed $rV(x) = \tilde{\pi}(x;\alpha)$ (Figure 3) for differing values of α for the following constant sum case:

$$\pi(x) = \begin{cases} -1 & \text{if} \quad x \leq -2 \\ 1 + x & \text{if} \quad -2 < x < -1 \\ 0 & \text{if} \quad -1 \leq x \leq 1 \\ x - 1 & \text{if} \quad 1 < x < 3 \\ 2 & \text{if} \quad x \geq 3 \end{cases}$$

Since $\underline{\pi} = -1, \overline{\pi} = 2$, we have $\pi(1) = (\overline{\pi} + \underline{\pi})/2$. It follows that $x^* = 1$ (High Variance case); that is, a player chooses low variance if and only if he is ahead by at least one

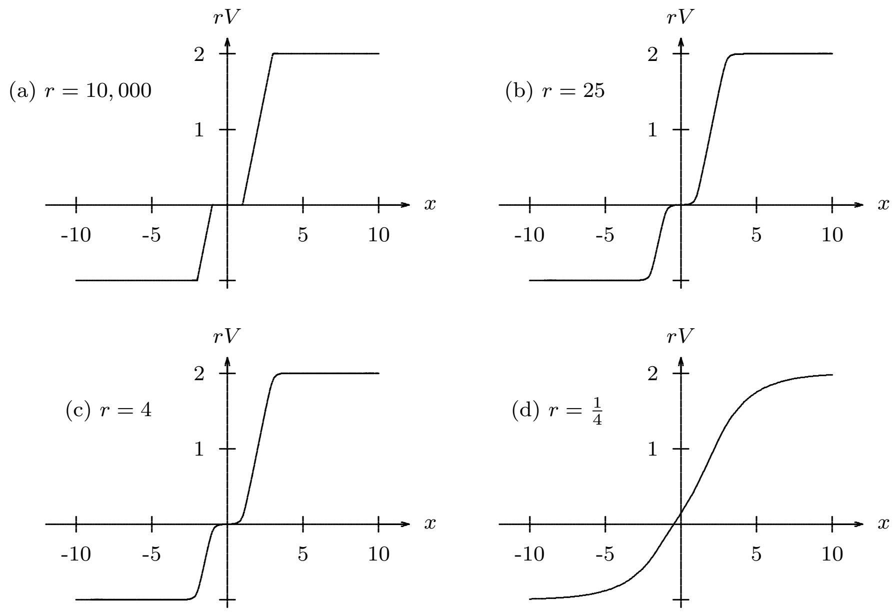

Figure 3: How $rV(x)$ Changes in r. From top left to bottom right $\sigma = 1$ and $r = \{10,000, 25, 4, \frac{1}{4}\}$. Note that the convexity / concavity result obtains long before V converges to a constant.

unit. In fact, as Figure 3 shows, even for values of r away from zero (that is, long before V converges to a constatn), the convexit/concavity result obtains.

5.3 How long until one player dominates?

One question that has received a lot of empirical and theoretical attention, is whether R and D competition leads to increasing dominance. That is, is it the case that firms that are ahead tend to pull farther ahead, or do firms that are behind tend to catch up to the market leaders? This question concerns the expected drift in x, but since we have ruled out expected drift *a priori* we cannot opine on this question as it is usually posed in the literature.

We can, however, ask a similar question: if two firms were located 'close together' at time 0, how long do we expect them to stay close together? Intuition suggests that the higher the variance in $\{x_t\}$ the faster (on average) the two firms should separate. To see this, think of the extreme case of zero variance; in that case the two firms would never separate. This intuition turns out to be correct. Specifically, if we let τ_x be the first exit time from the interval $(-x, x)$, given $x_0 \in (-x, x)$ for some $x > 0$, and $E^{x_0}[\tau_x]$

be the expected τ_x given x_0 then we have the following proposition.

Proposition 3 *$E^{x_0}[\tau_x]$ is highest in the Low Variance case, and lowest in the High Variance case.*

Proof: If $\{x_t\}$ has constant variance over this interval, then we can modify example 7.4.2 in Øksendal (1998) to determine the expected value of τ_x

$$E^{x_0}[\tau_x] = \frac{1}{2\sigma}(x^2 - x_0^2)$$

Notice that this expected value falls in σ. Now note that in each case $\{x_t\}$ has the same variance save for the interval $(-|x^*|, |x^*|)$. ■

Figure 4 illustrates Proposition 3. Instead of working with a primitive π function, we simply assume that x^* and apply Proposition 2: if r is sufficiently small, which we assume, then players choose $\sigma = \underline{\sigma}$ if $x > x^*$ and $\sigma = \overline{\sigma}$ if $x < x^*$. It follows that $\sigma_1 + \sigma_2 = 2\underline{\sigma}$ for $-x^* < x < x^*$ $\sigma_1 + \sigma_2 = \underline{\sigma} + \overline{\sigma}$ for $x < -x^*$ or $x > x^*$. Figure 4 plots a series of equilibrium paths $\{x_t\}$ for particular values of $x^*, \underline{\sigma}, \overline{\sigma}$. Even though the expected motion of x is zero, starting from $x = 0$ the system moves away from the "symmetry region" $[-x^*, x^*]$ relatively quickly.

Budd et al. (1993) and Cabral and Riordan (1994) provide conditions such that a dynamic competitive system will move away from symmetry in expected value (increasing dominance). In both papers, the fundamental condition is the "joint profit" or "efficiency" effect, namely that joint profits, $\pi(x) + \pi(-x)$, be increasing in $|x|$.[4] Cabral (2002) shows that increasing dominance may also result when firms choose the correlation of their motion with respect to their rival's, even if $\pi(x) + \pi(-x)$ is constant (no efficiency effect). Our result, by contrast, requires no particular assumption regarding $\pi(x) + \pi(-x)$. It does not directly pertain to increasing dominance. In fact, we *assume* that, in expected terms, the system will remain at the current state x. However, Corollary 1 and Proposition 3 have a flavor similar to increasing dominance, in the sense that, if $\pi(0) < (\overline{\pi} + \underline{\pi})/2$, then the system will have a tendency to move away from symmetry ($x = 0$).

[4]Cabral and Riordan (1994) consider, as we do, the limit case of very small discounting; Budd et al. (1993), by contrast, consider the case of high discounting.

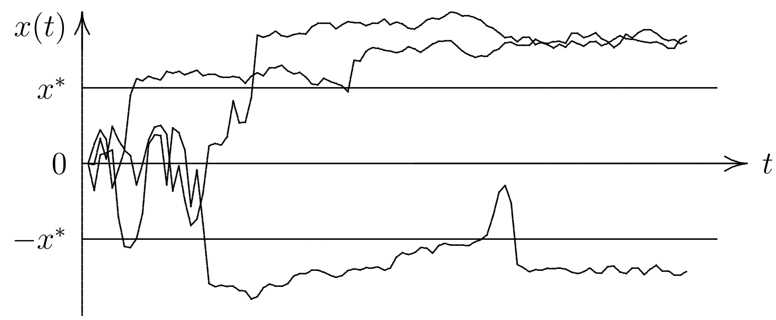

Figure 4: Sample paths of system dynamics when: $\underline{\sigma} = .5, \overline{\sigma} = 5, x^* = 10$.

6 Examples

In this section, we present three examples that illustrate our results. In all three examples, we consider a two stage game in each period, the first stage corresponding to R&D investment. Specifically, players first simultaneously choose σ_i, the variance of the increase in their product's quality level, q_i. The state of the game is given by $x = q_i - q_j$, the difference in quality levels.

The three examples differ with respect to the second stage of the period game. This leads to different profit functions $\pi(x)$ and different system dynamics.

6.1 Bertrand competition with differentiated products.

We begin by considering the case when, after R&D investments are made, price competition takes place. Specifically, suppose that each consumer receives utility $u = \max\{z_1 q_1, z_2 q_2\} + z_0$, where z_i is the quantity of good i, q_i is the quality of good i, and z_0 denotes other goods. Suppose that each consumer buys at most one unit from each firm ($z_i \in \{0, 1\}$) and is subject to a budget constraint such that he can only spend y. Finally, assume that marginal cost is constant and equal across firms (with no further loss of generality, assume marginal cost is zero). Firms simultaneously set prices and consumers then choose z_0, z_1, z_2. In equilibrium consumers buy from the firm with highest quality (say, firm i) at a price given by $\min\{q_i - q_j, y\}$. The profit function is therefore given by

$$\pi(x) = \begin{cases} 0 & \text{if} \quad x < 0 \\ x & \text{if} \quad 0 \leq x \leq y \\ y & \text{if} \quad x > y \end{cases}$$

In this example, $\pi(0) = 0, \overline{\pi} = y$, and $\underline{\pi} = 0$. The condition $\pi(0) < (\overline{\pi} + \underline{\pi})/2$ is therefore satisfied. Corollary 1 applies: if very patient players start play from $x = 0$, then the system will move away from symmetry very rapidly.

6.2 Price competition with brand loyalty

Consider a market where consumers are divided into four segments. $(1-\mu)/2$ consumers are highly loyal to firm 1's brand, and an equal fraction to firm 2's. Highly loyal consumers are willing to pay $\overline{p}$ for their favorite firm's product, zero for the rival's. The remaining consumers have lower levels of brand loyalty. A fraction $\mu/2$ is willing to pay $\underline{p} + q_1$ for product 1 and q_2 for product 2; an equal fraction is willing to pay q_1 for product 1 and $\underline{p} + q_2$ for product 2.

Suppose that μ is small and that the initial product quality levels are such that $q_i > \overline{p}$, for all i. Then the unique equilibrium of the pricing game is for firms to set $p_i = \overline{p}$, the highly loyal consumers willingness to pay. If $|q_i - q_j| \leq \underline{p}$, then mildly loyal consumers choose their favorite brand. If however $q_i - q_j > \underline{p}$, then all mildly loyal consumers choose firm i.

This situation leads to the following profit function:

$$\pi(x) = \begin{cases} \frac{1}{2}(1-\mu)\overline{p} & \text{if} \quad x \leq -\underline{p} \\ \frac{1}{2}\overline{p} & \text{if} \quad -\underline{p} < x < \underline{p} \\ \frac{1}{2}(1+\mu)\overline{p} & \text{if} \quad \underline{p} \leq x \end{cases}$$

Note that this case violates our assumption that π is strictly increasing in x. Nevertheless, solving for the equilibrium is straightforward. In fact, for cases when total payoffs add to a constant sum, we can prove a strong characterization result *for any* r:

Proposition 4 *If $\pi(x) + \pi(-x) = c$ for some constant c, then in equilibrium $\sigma(x) = \underline{\sigma} + \bar{\sigma}$ for all x. Thus, $rV(x) = r\tilde{\pi}(x;\alpha)$ with $\alpha = \sqrt{r/(\underline{\sigma} + \overline{\sigma})}$.*

Proof: The result obtains if $sign(V_1''(x)) = -sign(V_2''(x))$ for all x. Assume $V(x) = r\tilde{\pi}(x;\alpha)$ for the given α. Then

$$V_1(x) + V_2(x) = r\tilde{\pi}(x;\alpha) + r\tilde{\pi}(x;\alpha) = c,$$

which implies $V_1''(x) = -V_2''(x)$. Thus, $\sigma(x) = \underline{\sigma} + \bar{\sigma}$ and $rV(x)r = \tilde{\pi}(x;\alpha)$ as assumed. ■

We thus have $rV(x) = r\tilde{\pi}(x;\alpha)$ with $\alpha = \sqrt{r/(\underline{\sigma}+\bar{\sigma})}$. Integrating we find:

$$r\tilde{\pi}(x;\alpha) = \begin{cases} \bar{p}(1-\mu+\mu\cosh(\alpha\bar{p})e^{\alpha x}) & \text{if} \quad x \leq -\underline{p} \\ \bar{p}(1+\mu\sinh(\alpha x)e^{-\alpha\bar{p}}) & \text{if} \quad -\underline{p} < x < \underline{p} \\ \bar{p}(1+\mu-\mu\cosh(\alpha\bar{p})e^{\alpha x}) & \text{if} \quad \underline{p} \leq x \end{cases}$$

where $\cosh(z) = (e^z + e^{-z})/2$ and $\sinh(z) = (e^z - e^{-z})/2$. We can then twice differentiate to find:

$$r\tilde{\pi}''(x;\alpha) = \begin{cases} \bar{p}\alpha^2\mu\cosh(\alpha\bar{p})e^{\alpha x} > 0 & \text{if} \quad x \leq -\underline{p} \\ \bar{p}\alpha^2\mu\sinh(\alpha x)e^{-\alpha\bar{p}} & \text{if} \quad -\underline{p} < x < \underline{p} \\ -\bar{p}\alpha^2\mu\cosh(\alpha\bar{p})e^{-\alpha x} < 0 & \text{if} \quad \underline{p} \leq x \end{cases}$$

Thus, firms choose the risky strategy when behind by more than $\bar{p}$, and the safe strategy when ahead by more than $\bar{p}$. However, since $\sinh(z)$ is negative for $z < 0$ and positive for $z > 0$, firm 1 chooses the safe strategy when $x \in (-\bar{p}, 0)$, the risky strategy when $x \in (0, \bar{p})$, *regardless of the value of* r.

This example illustrates the importance of the strict monotonicity assumption on π in Proposition 2. If we drop the strict monotonicity requirement, we get a weaker result. It is still the case that there exist a low and a high threshold, such that the firm chooses the risky strategy below the lower threshold and the safe strategy above this threshold, but the distance between the two thresholds no longer shrinks to zero as r goes to zero.

6.3 Competitive balance in sports

In sports leagues, a team's value is a function of its competitive success as well as the overall success of its league; and the league's success is a function of competitive balance. For simplicity, consider a league with two "important" teams. Let x be the difference in quality between the teams (e.g., the average skill of its roster). Suppose that each period corresponds to a season and that at the beginning of the season each team gets to choose the variance of its quality variation. Let $\rho(x)$ be the probability of winning the league and $\nu(x)$ the value of the league. We assume that $\rho(x)$ is increasing and that $\nu(x)$ is decreasing in $|x|$, a measure of competitive imbalance.

Specifically, suppose that $\nu(x)$ declines exponentially with competitive imbalance:

$$\nu(x) = \begin{cases} \gamma + (1-\mu)e^{-|x|} & \text{if } \ |x| \leq \ln 2 \\ \frac{1}{2}(1+\mu) & \text{if } \ |x| > \ln 2 \end{cases},$$

where $\mu \in (\frac{2}{3}, 1)$. Suppose moreover that the likelihood that team i wins each league is exponentially increasing in its quality lead: $\rho(x) = \frac{1}{2}e^x$ (for values of x in $[0, \ln 2]$). Pulling all of these elements together, we have a profit function

$$\pi(x) = \rho(x)\nu(x) = \begin{cases} 0 & \text{if } \quad x < -\ln 2 \\ (\mu + (1-\mu)e^x)\left(1 - \frac{1}{2}e^{-x}\right) & \text{if } \ -\ln 2 \leq x \leq 0 \\ \frac{1}{2}\left(1 - \mu + \mu e^x\right) & \text{if } \quad 0 \leq x \leq \ln 2 \\ \frac{1}{2}(1+\mu) & \text{if } \quad x > \ln 2 \end{cases}$$

Note that since $\mu \in (\frac{2}{3}, 1)$, $\pi(x)$ is monotonically increasing, as our assumptions require.

Consider now the equilibrium strategies, beginning with the case when r is very high (high discounting). Straightforward computation shows that

$$\pi''(0^-) = 1 - \frac{3}{2}\mu$$
$$\pi''(0^+) = \frac{1}{2}\mu.$$

Since $\mu \in (\frac{2}{3}, 1)$, it follows by continuity that $\pi''(0^-) < 0$ whereas $\pi''(0^+) > 0$. That is, for x close to zero, $\pi(x)$ is concave for the laggard and convex for the leader. This implies that, for low discounting (high r) and when x is close to zero the leader chooses high variance whereas the laggard chooses low variance.

Consider now the case when r is low (low discounting). Notice that $\underline{\pi} = 0, \overline{\pi} = \frac{1}{2}(1+\mu)$, and $\pi(0) = \frac{1}{2}$. Therefore, since $\mu < 1$, $\pi(0) > (\underline{\pi} + \overline{\pi})/2$, which implies that the Low Variance case obtains by Corollary 1. In fact, it case be shown that players choose low variance if and only if $x > x^*$, where $x^* \approx -.693$. Thus, for low r, both firms choose low variance near $x = 0$.

This example illustrates that discounting may be quite important in determining leader and laggard's variance choices. In fact, for $-2 < x < -.693$ and for $0 < x < \ln 2$, the equilibrium strategies change from one extreme to the opposite as the discount rate goes from a high value to a low value.

7 Conclusion

Conventional wisdom from sports indicates that, close to the end of a game or race, the laggard should choose a high-variance strategy and the leader a low-variance strategy. In fact, the laggard has "nothing to lose:" his payoff does no decrease if he falls farther behind but his value may increase substantially if he moves ahead; in other words, his value function is convex. In this paper, we consider the situation of an infinite race. We show that, if players are sufficiently patient, then a laggard, if sufficiently behind, will choose a high-variance strategy, and the leader a low-variance strategy.

The summary intuition for our result is derived from the HJB equation, which in our game becomes

$$rV(x) = \pi(x) + (\sigma_1 + \sigma_2)V''(x).$$

This implies that the second derivative of the value function is negatively related to the current payoff level. Specifically, a lagging player receives a low payoff and has a convex value function; whereas a leading player receives a high payoff and has a concave value function. Finally, Jensen's inequality implies that a lagging player chooses high variance whereas a leading player choose low variance.

There are at least two possible extensions of our game that would be of interest. First, we could combine our model with previous models of R&D races and consider the case when players choose variance *and* effort. Second, we could also consider a model when the players' strategy choice determine the correlation between their motions.[5]

[5] Cabral (2002) does so for the particular case when total payoff is constant, i.e., $\pi(x) + \pi(-x) = c$.

References

Bergemann, Dirk, and Juuso Välimäki (2002), "Entry and Vertical Differentiation," *Journal of Economic Theory* **106**, 91–125.

Bolton, Patrick, and Christopher Harris (1999), "Strategic Experimentation," *Econometrica* **67**, 349–374.

Bolton, Patrick, and Christopher Harris (2003), "Strategic Experimentation: The Undiscounted Case," in P. Hammond and G. Myles (Eds.), *Incentives and Organization: Essays in Honour of James Mirrlees*.

Budd, Christopher, Christopher Harris and John Vickers (1993), "A Model of the Evolution of Duopoly: Does the Asymmetry Between Firms Tend to Increase or Decrease?," *Review of Economic Studies* **60**, 543–573.

Cabral, Luís M. B. (2002), "Increasing Dominance With no Efficiency Effect," *Journal of Economic Theory* **102**, 471–479.

Cabral, Luís M. B., and Michael H. Riordan (1994), "The learning curve, Market Dominance and Predatory Pricing," *Econometrica* **62**, 1115–1140.

Dixit, Avinash, and Robert Pindyck (1994), *Investment Under Uncertainty*, Princeton, NJ: Princeton University Press .

Dutta, Prajit (1991), "What do Discounted Optima Converge to?," *Journal of Economic Theory* **55**, 64–94.

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

Fershtman, Chaim, and Ariel Pakes (2000), "A Dynamic Oligopoly With Collusion and Price Wars," *Rand Journal of Economics* **31**, 207–36.

Harris, Christopher (1993), "Generalized Solutions of Stochastic Differential Games in One Dimension," Nuffield College.

Hörner, Johannes (1999), "A Perpetual Race to Stay Ahead," forthcoming in *Review of Economic Studies*.

Khanna, Tarun, and Marco Iansiti (1997), "Firm Asymmetries and Sequential R&D: Theory and Evidence from the Mainframe Computer Industry," *Management Science* **43**, 405–421.

Øksendal, Bernt (1998), *Stochastic Differential Equations*, Springer-Verlag.