RAND Journal of Economics
Vol. 38, No. 3, 2007
pp. 593–609

Go for broke or play it safe? Dynamic competition with choice of variance

Axel Anderson*

and

Luís M. B. Cabral**

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 show that a pure strategy perfect equilibrium in stationary Markov strategies (ME) exists and has the property that patient players choose 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."

1. Introduction

■ Characterizing observed firm behavior in terms of R&D budgeting, Cyert and March (1963) argue that "most organizations are aware of and probably use such simple rules as per cent of revenue as a guide to research and development allocations" (p. 274). In their study of the microprocessor industry, Khanna and Iansiti (1997) report that "interfirm researcher mobility is remarkably low" (p. 406). Moreover, evidence from the microprocessor and other industries suggests that there are frequently different paths to achieve the same goal. For example, a given level of microprocessor speed can be attained through different computer architectures.

Together, the above observations suggest that, from a manager's point of view, the decision is not just how much to spend on R&D but also how to spend it. In fact, in some cases, the main decision may be to choose among R&D strategies with different degrees of risk. In this article, we focus on this dimension of R&D policy. Specifically, 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. Each player's position may be interpreted as its current quality level. In the example above, q_i might denote the speed of firm i's current microprocessor. Player i receives a payoff flow given by $\pi(q_i - q_j)$. The player's position, q_i, evolves according to a Wiener process with mean μ, which is exogenously given, and variance

* Georgetown University; aza@georgetown.edu.
** New York University; lcabral@stern.nyu.edu.
We are grateful to the editor, two anonymous referees, Dirk Bergemann, and various seminar audiences for useful comments and suggestions. The usual disclaimer applies.

σ_i, which is chosen by firm i. Our goal is to characterize the pure strategy perfect equilibrium in stationary Markov strategies (Markov equilibrium; ME) of this game. In other words, we want to understand when players choose safer or riskier R&D strategies as a function of their relative position.

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, whereas 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, whereas 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, whereas 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."

What is common to the sports examples is that (i) we are close to the end of the game and (ii) the final payoff function is locally convex for the laggard and locally concave for the leader. For example, suppose a hockey team trails by one goal one minute from the end. In terms of final outcome, the payoff is the same if the team allows an additional goal, but higher if it scores an additional goal. So, the final payoff function is convex at -1. The fact that we are close to the end of the game makes it easy (at least conceptually) to compute the value functions. In fact, if we are close to the end of the game, then convexity of the final payoff implies convexity of the value function. Finally, by Jensen's inequality, it follows that the trailing team benefits from a mean-preserving spread in the goal-scoring function.

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.[1] 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, 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. Let x be the relative difference between the players in the game. If flow profits as a function of this state variable are bounded, and admit limits as the state variable tends to the extremes ($+/-\infty$), and satisfy a single-crossing property[2] at some value x^*, then in Markov equilibrium, patient players choose to play it safe when ahead and to take risks when behind. Specifically, if players are patient enough, they will choose low variance if $x > x^*$ and high variance if $x < x^*$. Note that we need not make any assumptions about the local curvature of the 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; 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. 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.

When $x^* > 0$, both players choose risky strategies in states where x is close to zero. It follows that, starting from a situation where players are more or less even, a dominant player will emerge

[1] However, some infinite horizon games may share some of the features of finite games as in the previous examples. For example, suppose that if one of the players falls sufficiently far behind, then it must exit the game, receiving a payoff of zero. See Section 4 for a related example.

[2] The single-crossing property we require is weaker than monotonicity.

"quickly." Previous research (Athey and Schmutzler, 2001; Budd, Harris, and Vickers, 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, players' choices do not influence 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.

We compare the ME outcome with the policy that maximizes the sum of the expected discounted profits of the two players (the Planner's solution). We show that the Planner will choose either the highest or the lowest variance possible for any discount factor. This immediately yields that, with enough patience, the ME outcome is inefficient outside of the interval $[-|x^*|, |x^*|]$. However, we also show that inside this interval the equilibrium is efficient; that is, when the players are "close enough" together, the Planner's choice corresponds to the ME outcome.

□ **Related literature.** Bhattacharya and Mookherjee (1986) and Klette and de Meza (1986) consider patent race models where players choose variance. Although they explicitly consider time, their models are static in the sense that firms make a once-and-for-all choice. They show that, in equilibrium, firms choose too much risk from a social welfare point of view. The intuition is that there is an externality in patent races: a firm's gain from anticipating its rival is less than the social benefit from earlier adoption.

Judd (2003) develops an explicitly dynamic patent race in continuous time. He assumes that, at each moment, each player may choose between a *partial jump* and a *leap* motion technology. Because the latter implies a bigger variation in motion (zero motion or winning the race), placing more resources into the *leap* technology effectively corresponds to a higher-risk strategy. Judd's Theorem 8 states that, if the race prize is close to zero, then social welfare would be increased if resources were shifted from the risky R&D projects to the less risky projects.

All three papers concur that there is too much variance in equilibrium. In broad strokes, the intuition is that there is an externality in patent races: the marginal private benefit from winning the race is lower than the social benefit as part of the increase in the probability of winning is associated with a lower probability that others win (which would be equally good, from a welfare point of view); and the higher the degree of risk, the greater the probability of an immediate end to the race, and the greater the above externality. Our model, in turn, shows that the equilibrium level of variance may be greater, smaller, or equal to the socially optimal level. The idea is that, given the linearity of the stochastic process we consider, the social optimum is either the highest or the lowest level of variance; but if players are sufficiently apart, then the curvature of their value functions must have opposite signs, and so one of them (exactly one of them) will choose the opposite of the social optimum.

More closely related to our model, Cabral (2003) considers a discrete-time, discrete-space R&D game where firms choose variance. He presents a series of examples from economics and management. However, his formal analysis is rather limited, as it does not include an existence result or a complete characterization of equilibrium strategies, both of which we provide in this article.

Several authors have looked at dynamic games with the properties that (i) in each period, each firm is characterized by the value of its product; (ii) in each period, each firm's profit is a function of all firms' product values; and (iii) by investing resources into R&D, a firm stochastically improves the future quality of its product. The list includes Budd, Harris, and Vickers (1993), Ericson and Pakes (1995), Fershtman and Pakes (2000), and Hörner (2004). One feature that is common to all of these models is that firm strategies consist of choosing the *level* of R&D expenditures. Our analysis complements theirs: we fix the level of R&D expenditures and consider the strategic choice of risk.

Technically, we study a (very simple) one-dimensional stochastic differential game in which the agents' choices affect the variance of some state variable.[3] The 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. We are aware of only two other papers that prove existence for particular variance choice games: Bergemann and Välimäki (2002) and Bolton and Harris (2001). These papers consider the ME of the *undiscounted* game directly. 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. We could also have considered the limiting equilibrium directly. Instead, we characterize the equilibrium value functions and optimal strategies, and *then* investigate their behavior as players become infinitely patient. Given the simplicity of our model, we feel that this is the right approach. Note that in Bergemann and Välimäki (2002) and Bolton and Harris (2000) the models are more complex, and the restriction to the undiscounted game is necessary in order to make reasonable progress.

The article is organized as follows. In Section 3, we present the dynamic game and show that an ME exists. In Section 4, we characterize the equilibrium in the cases when players are very patient. We also present results for the particular case of constant sum games, and derive implications for industry dynamics. In Section 5, we solve the Planner's problem and investigate the efficiency of the ME. In Section 6, we discuss some natural extensions of the basic model. Section 7 concludes the article.

3. The model and existence

■ Consider the following two-player stochastic continuous time (differential) game.[4] At each instant in time, player $i \in \{1, 2\}$ chooses $\sigma_i \in [\underline{\sigma}, \bar{\sigma}]$ ($\underline{\sigma} > 0$), 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[5]:

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

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 $\pi(-x)$. Let π have limits $\lim_{x\to-\infty} \pi(x) = \underline{\pi} > -\infty$ and $\lim_{x\to\infty} \pi(x) = \bar{\pi} < \infty$. Our proof critically depends on these limits boundedly existing. One situation in which this assumption would be automatically satisfied is if there is a threshold value of $\hat{x}$ such that the laggard is eliminated from the market, so that $\pi(x)$ would be the flow profits from some outside option for $x < \hat{x}$, and $\pi(x)$ would be the flow monopoly profits for $x > \hat{x}$.

We assume π satisfies the following single-crossing property: there exists an x^* such that

$$\pi(x) < (>) \frac{1}{2}(\underline{\pi} + \bar{\pi}) \quad \text{if and only if} \quad x < (>) x^*.$$

[3] Several authors have considered one-dimensional games as models of duopoly competition: see Harris and Vickers (1987), Budd, Harris, and Vickers (1993), and Athey and Schmutzler (2001). Budd, Harris, and Vickers (1993) present some examples of oligopoly games that satisfy the one-dimensionality restriction. Additional examples are presented in Section 4. These examples notwithstanding, we must acknowledge that the assumption of a one-dimensional state space is fairly restrictive, and is violated by a number of standard oligopoly models, such as logit demand with an outside good. Referring to models of effort choice, Budd, Harris, and Vickers (1993) claim that "the effects found in the one-dimensional model were found to be at work also in a two-dimensional model" (footnote 2). In fact, Cabral and Riordan (1994) consider a two-dimensional game and derive results similar to those of Budd, Harris, and Vickers (1993). However, it is unclear whether such extension would work in the context of variance choice.

[4] See Harris (1993) for a very thorough treatment of one-dimensional stochastic differential games.

[5] A good (accessible) reference for basic stochastic control is Dixit and Pindyck (1994). For a more technical reference, see Øksendal (1998).

As mentioned in the introduction, this single-crossing property is satisfied by all strictly monotonic flow profit functions. However, strict monotonicity is not required to satisfy this assumption. Weakly monotonic flow payoff functions are fine as long as they are not flat around $(\underline{\pi} + \bar{\pi})/2$. Finally, we assume that players discount future profits at rate r.

We will be considering pure strategy equilibria in stationary Markov strategies, which we will henceforth abbreviate as Markovian equilibria. A *Markov strategy* for player i is a measurable map $\sigma_i : (-\infty, +\infty) \mapsto [\underline{\sigma}, \bar{\sigma}]$.[6] Given a strategy pair $\sigma = \sigma_1 + \sigma_2$, the payoffs for the players are

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

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

In summary, we have a symmetric game on a one-dimensional space, $x(t) \in \mathbb{R}$. The expected motion of x is zero, but its variance depends on the players' choices. Specifically, at each point x of the state space, each player chooses variance within the interval $[\underline{\sigma}, \bar{\sigma}]$, with the system variance equal to the sum of the players' choices.

We now show that an equilibrium exists for this game. Fix a Markov strategy σ_2 for player 2. Then player 1's Markov best response solves:

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

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.

The Hamilton-Jacobi-Bellman equation (HJB) associated with this maximization problem is (via Ito's Lemma):

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

Proposition 1. A Markov equilibrium exists, and $V_i = U_i^*$ is continuous for $i \in \{1, 2\}$ in ME.

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 players choose scalars $\sigma_i \in [\underline{\sigma}, \bar{\sigma}]$ and the payoff for player i is[7]

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

where $(\lambda_i, \lambda_i'') \in \mathbb{R}^2$. Following Harris, let $\overline{ne}(x, \lambda, \lambda'')$ be the set of Nash equilibrium payoff vectors for this static game, where $\lambda = (\lambda_1, \lambda_2)$ and $\lambda'' = (\lambda_1'', \lambda_2'')$. Then by Theorem 6.6 in Harris, an ME 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,

[6] We focus on Markovian equilibria, rather than the more general feedback Nash equilibria for two reasons. (i) We are specifically interested in how variance choice depends on the state variable rather than on calendar time; and (ii) by focusing on stationary Markov strategies, we dodge technical difficulties in defining time-dependent strategies in continuous time. For a discussion of feedback Nash equilibria, see Basar and Olsder (1998). For a discussion of issues related to defining time-contingent strategies in continuous time, see Simon and Stinchcombe (1989) and Bergin and MacLeod (1993).

[7] For some intuition, substitute (λ_i, λ_i'') for $(V_i(x), V_i''(x))$ and rearrange the Bellman equation payoffs.

the set of Nash equilibria is $\sigma_1 = \underline{\sigma}, \sigma_2 \in [\underline{\sigma}, \bar{\sigma}]$. The payoff vector for player 2 is the same in every equilibrium. The set of payoff vectors for player 1 is the convex set

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

and thus Theorem 6.6 in Harris (1993) applies. *Q.E.D.*

4. Results with high patience

■ Players' attitudes toward risk will be influenced by the curvature of the flow profit function π. If players are very impatient (high r), then the local curvature of π will weigh heavily in their decision making. In fact, if π is convex (concave) in a neighborhood of x, then player 1 chooses the risky (safe) process at x if r is above a certain threshold. Because we can provide examples of functions that alternate between convex and concave throughout the range of x, we cannot hope for low-patience analogs of our high-patience results. One class of examples is

$$\pi(x) = \begin{cases} \frac{a-x}{(a-x)^2 - b \sin x} & x < 0 \\ \frac{2}{a} - \frac{a+x}{(a+x)^2 + b \sin x} & x \geq 0 \end{cases} \tag{1}$$

with $b > 2$ (Figure 1).

We do not find these insights for the high r case surprising or particularly interesting. Instead we focus on what happens for low r. Given our minimal assumptions on π, it is not obvious *a priori* what the nature of the equilibrium strategy is.

□ **Risk choice in the limit.** Our main result is that x^* divides the state space so that with enough patience, high variance is chosen by player 1 when $x < x^*$ and low variance is chosen by player 1 when $x > x^*$. The structure of the argument that establishes this result is straightforward and proceeds in the following three steps:

Step 1. The Bellman equation implies that *sign*$(rV_1(x) - \pi(x)) =$ *sign*$(V_1''(x))$. Also, the Bellman equation is linear in variance choice with coefficient $V_1''(x)$, so $rV_1(x) - \pi(x) > 0$ implies $\sigma_1(x) = \bar{\sigma}$ and $rV_1(x) - \pi(x) < 0$ implies $\sigma_1(x) = \underline{\sigma}$.

Step 2. In the long run, x spends almost all of its time arbitrarily far from 0. With no drift, x is equally likely to be arbitrarily close to ∞ and $-\infty$. Given $\lim_{x\to-\infty}\pi(x) = \underline{\pi}$ and $\lim_{x\to\infty}\pi(x) = \bar{\pi}$, we have $\lim_{r\to 0} rV_1(x) = (\underline{\pi} + \bar{\pi})/2$ (Lemma 1).

Step 3. The single-crossing property combined with Step 2 implies that for any $x < x^*$, r low enough yields $rV_1(x) > \pi(x)$ and thus $\sigma_1(x) = \bar{\sigma}$ by Step 1.

FIGURE 1

PLOT OF FUNCTION GIVEN IN (1) FOR $a = 3$, $b = 5$

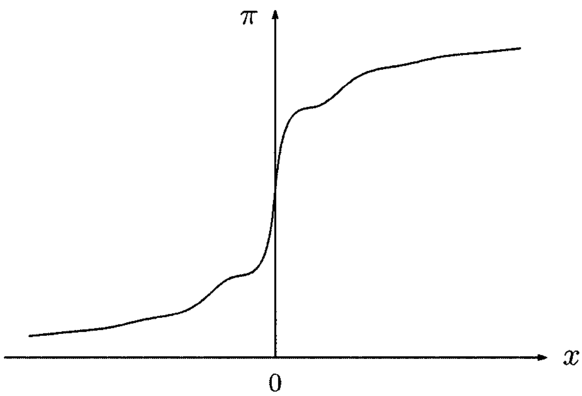

First we formally establish Step 2:

Lemma 1. $\lim_{r\to 0} r V_i(x) = (\underline{\pi} + \bar{\pi})/2$ for $i \in \{1, 2\}$.

The proof is in the Appendix.

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. For all $x < x^*$, $\exists r^*(x)$ such that $\sigma_1(x) = \bar{\sigma}, \forall r < r^*(x)$. Conversely, $\forall x > x^*$, $\exists r^*(x)$ such that $\sigma_1 = \underline{\sigma}, \forall r < r^*(x)$.

Proof. We have $\pi(x) > (\underline{\pi} + \bar{\pi})/2$ for all $x > x^*$ and $\pi(x) < (\underline{\pi} + \bar{\pi})/2$ for all $x < x^*$. By Lemma 1, $\lim_{r\to 0} r V_1(x) = (\underline{\pi} + \bar{\pi})/2$. Finally, by the HJB equation for player 1, $r V_1(x) > \pi(x)$ implies that $V_1''(x) > 0$, which in turn implies $\sigma_1(x) = \bar{\sigma}$, while $r V_1(x) < \pi(x)$ implies that $V_1''(x) < 0$, which in turn implies $\sigma_1(x) = \underline{\sigma}$. *Q.E.D.*

To illustrate this results, we graphed $rV_1(x)$ (Figure 2) for differing values of r for the following constant sum case:

$$\pi(x) = \begin{cases} -2 & \text{if } x \le -2 \\ 2 + 2x & \text{if } -2 < x < -\frac{3}{2} \\ \frac{2}{3}x & \text{if } -\frac{3}{2} \le x \le \frac{3}{2} \\ 2x - 2 & \text{if } \frac{3}{2} < x < 2 \\ 2 & \text{if } x \ge 2 \end{cases} .$$

Because $\underline{\pi} = -1, \bar{\pi} = 2$, we have $\pi(0) = (\underline{\pi} + \bar{\pi})/2$. It follows that $x^* = 0$; that is, a patient player chooses low variance if and only if he is ahead by at least one unit. In fact, as Figure 2 shows, even for values of r away from zero (that is, long before rV converges to a constant), the value function is concave below $x^* = 0$ and concave above $x^* = 0$, for example when $r = 1/4$.

Proposition 2 states that x^* divides the state space, so that for all $x < x^*$ (laggard), high risk is the equilibrium strategy given sufficient patience, whereas for $x > x^*$ (leader), low risk is better given sufficient patience. Figure 3 illustrates this. Notice that the figure also illustrates that the threshold value of r depends on the particular x considered. In this example, the closer x is to zero the lower the threshold $r^*(x)$.

Intuitively, Proposition 2 can be understood with reference to each player's HJB. Clearly, the value function is convex if and only if $\pi(x) < r V(x)$. In other words, if current profit is less than average discounted payoff, then "things can only get better." If things are going to get better it is 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 $\pi(x) < r V(x)$ for a laggard then we are done: 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 equilibrium can only be one of the three types pictured in Figure 4. The knife-edged case of $x^* = 0$ is straightforward. Note that in this case $\sigma = \underline{\sigma} + \bar{\sigma}$, 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 ($\sigma_i = \bar{\sigma}$) 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. These 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 greater than $(\underline{\pi} + \bar{\pi})/2$, respectively.

■ *Example: Bertrand competition.* Consider an industry with two firms in which price competition takes place after R&D investments are made. Specifically, suppose that each consumer

FIGURE 2

HOW $rV(x)$ CHANGES IN r

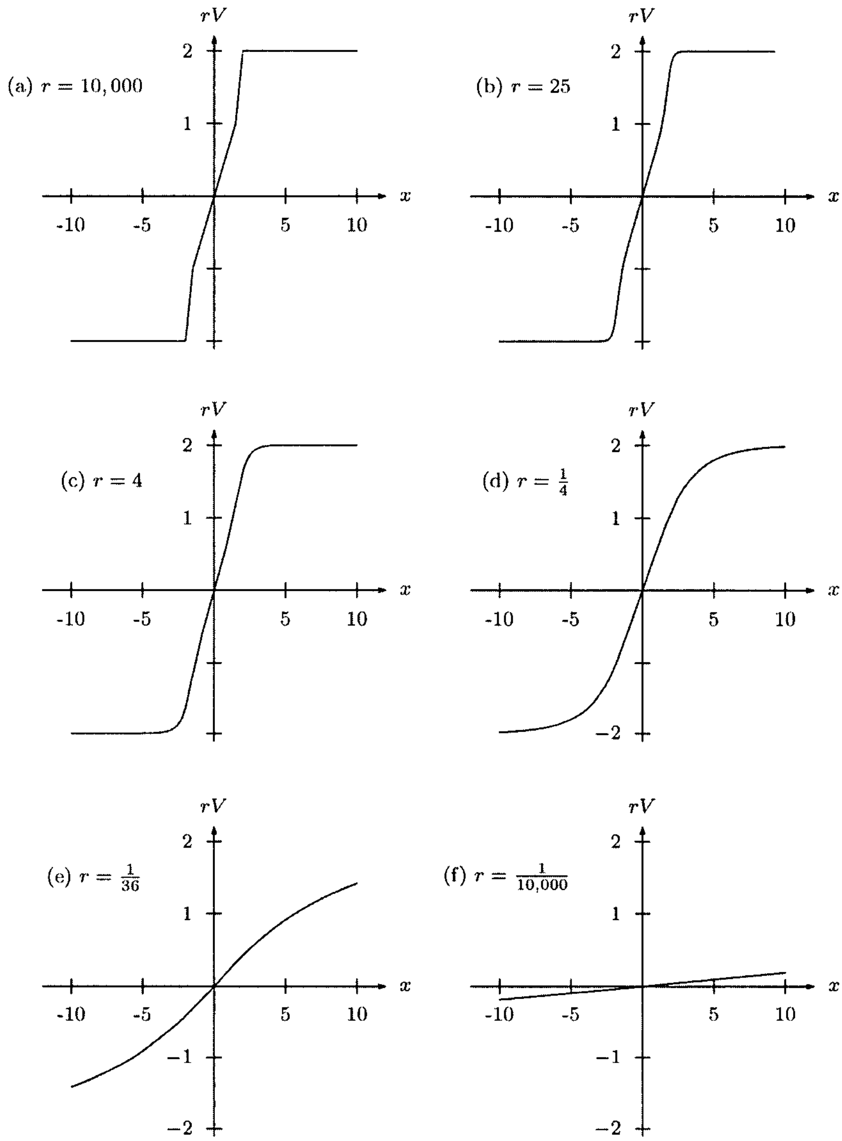

receives utility $u = \max\{z_1 q_1, z_2 q_2\} + z_0$, where z_i is the quantity of good i, q_i is the quality of good i, and z_0 denotes other goods. Suppose that each consumer buys at most one unit from each firm ($z_i \in \{0, 1\}$) and is subject to a budget constraint such that he can only spend 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 the highest quality (say, firm i) at a price given by min $\{q_i - q_j, y\}$. The profit function is therefore given by

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

FIGURE 3

THE CONCAVITY OF $V(x)$

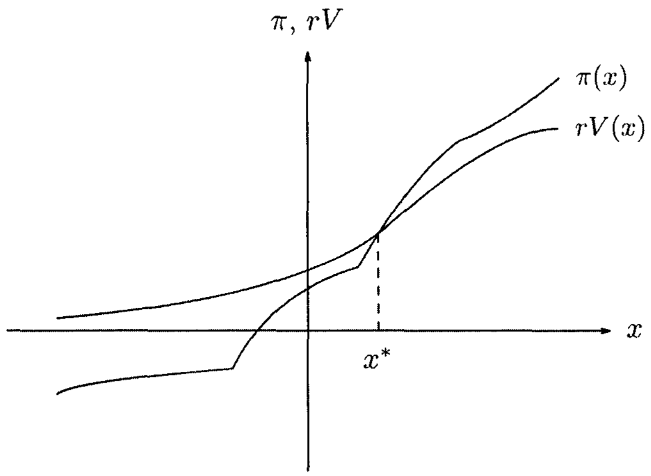

FIGURE 4

THE THREE ME FOR LOW r

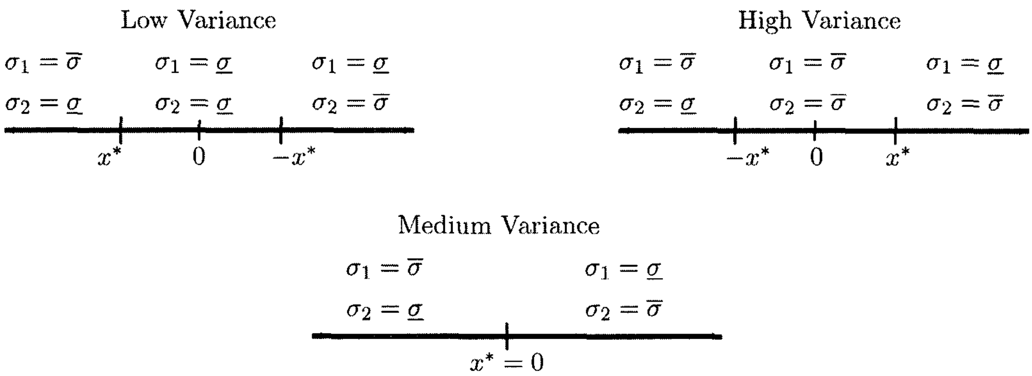

In this example, $(\underline{\pi} + \overline{\pi})/2 = y/2$ and thus $x^* = 1/2 > 0$; thus, this is the high-variance case. Corollary 1 applies: near $x = 0$, both firms choose high variance.

■ *Example: 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 instant corresponds to a season and that at the beginning of the season each team gets to choose the variance of its quality change. 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:

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

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

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

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

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

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

As $\mu\in(\frac{2}{3},1)$, 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 enough discounting (high r) and when x is close to zero, the leader chooses high variance whereas the laggard chooses low variance, a reversal of what must occur with patient players. This example illustrates that discounting may be quite important in determining the variance choices of leaders and laggards.

Finally, consider low r. Notice that $\underline{\pi}=0$, $\bar{\pi}=\frac{1}{2}(1+\mu)$, and $\pi(0)=\frac{1}{2}$. As, $\mu<1$, $\pi(0)>(\underline{\pi}+\bar{\pi})/2$. In addition, $\mu\in(\frac{2}{3},1)$ implies π is monotonically increasing. All together, this implies that $x^*<0$, and the low-variance case obtains by Corollary 1. Thus, for low r, both firms choose low variance near $x=0$.

□ **Constant Sum Games.** Notice that if $\pi(x)+\pi(-x)=c$ for some constant c (i.e., we have a constant sum game), then $\bar{\pi}+\underline{\pi}=2c$, so $x^*=0$ and we are in the medium-variance case. Thus, in any constant sum game, patient players will choose high variance when behind and choose low variance when ahead. In fact, we can prove a stronger result in the constant sum case.

Proposition 3. If $\pi(x)+\pi(-x)=c$ for some constant c, then in equilibrium $\sigma(x)=\underline{\sigma}+\bar{\sigma}$ for all x.

Proof. By definition:

$$rV_1(x)+rV_2(x)=E\left[\int_0^\infty re^{-rt}(\pi(x(t))+\pi(-x(t))dt \mid x(0),(\sigma_1,\sigma_2)\right]$$

$$=c\int_0^\infty re^{-rt}dt=c.$$

Thus, $rV_1(x)+rV_2(x)=c=\pi(x)+\pi(-x)$. So,

$$rV_1-\pi(x)=-(rV_2(x)-\pi(-x)),$$

and thus,

$$sign(rV_1(x)-\pi(x))=-sign(rV_2(x)-\pi(-x)).$$

Finally, note that the variance choice of player i is determined by $sign(rV_i(x)-\pi_i(x))$. *Q.E.D.*

This result obtains for any r.

■ *Example: 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 is highly loyal to firm 2's. Highly loyal consumers are willing to pay $\bar{p}$ for their favorite firm's product and 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 $q_i > \bar{p}$ for all i. Then the unique equilibrium of the pricing game is for firms to set $p_i = \bar{p}$, 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 instead.

This situation leads to the following profit function:

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

This is a constant sum example, so solving for the equilibrium is straightforward. We have

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

where $\gamma = 2\sqrt{r(\underline{\sigma} + \bar{\sigma})}$ and $\alpha = \sqrt{r/(\underline{\sigma} + \bar{\sigma})}$, by Proposition 3. Integrating we find:

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

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

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

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, because $\sinh(z)$ 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. Thus, x^* does not behave as a cutoff as in Proposition 2.

What fails? Notice that $(\underline{\pi} + \bar{\pi})/2 = \bar{p}/2$, and that $\pi(x) = \bar{p}/2$ for a range of x values. That is, the single-crossing property is not satisfied in this example. Thus, despite the fact that $rV_i(x)$ is converging to this constant, we cannot sign $rV_1(x) - \pi(x)$ regardless of the value of r on this range.

□ **How long until one player dominates?** One question that has received a lot of empirical and theoretical attention is whether R&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 as 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 Ito process $x(t)$, 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 4. $E^{x_0}[\tau_x]$ is highest in the low-variance case, and lowest in the high-variance case.

Proof. Recall that $f(x' \mid t, x(0), \sigma)$ is the probability that $x(t) = x'$ at time t, which is a normal density with mean $x(0)$ and some variance. Note that in the high-variance case the variance

FIGURE 5

SAMPLE PATHS OF SYSTEM DYNAMICS WHEN $\underline{\sigma} = 0.5, \bar{\sigma} = 5, x^* = 10$

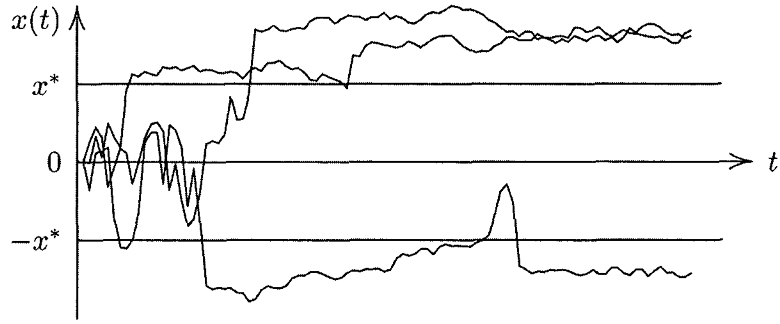

$\sigma_1(x) + \sigma_2(x)$ is higher for all x, thus the overall variance is higher as well. This means that starting from the same $x(0)$, the probability of being outside of any interval $(-x, x)$ at any time t is higher in the high-variance case. Thus, the expected exit time is lower. *Q.E.D.*

Figure 5 illustrates Proposition 4. Instead of working with a primitive π function, we simply assume that $x^* > 0$ and apply Proposition 2: if r is sufficiently small, which we assume, then players choose $\sigma = \underline{\sigma}$ if $x > x^*$ and $\sigma = \bar{\sigma}$ if $x < x^*$. It follows that $\sigma_1 + \sigma_2 = 2\underline{\sigma}$ for $-x^* < x < x^*$ and $\sigma_1 + \sigma_2 = \underline{\sigma} + \bar{\sigma}$ for $x < -x^*$ or $x > x^*$. Figure 5 plots a series of equilibrium paths $\{x(t)\}$ for particular values of $x^*, \underline{\sigma}, \bar{\sigma}$. 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, Harris, and Vickers (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 $|x|$.[8] Cabral (2002) shows that increasing dominance may also result when firms choose the correlation of their motion with respect to their rivals', 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 4 have a flavor similar to increasing dominance, in the sense that, if $\pi(0) < (\underline{\pi} + \bar{\pi})/2$, then the system will have a tendency to move away from symmetry ($x = 0$).

5. Equilibrium and efficiency

■ One question that has received some attention in the R&D literature is the relationship between equilibrium and efficient choices of risk. Bhattacharya and Mookherjee (1986) and Klette and de Meza (1986) show that, in a static patent race model, firms choose risk levels that are inefficiently high.[9] In this section, we solve for the efficient solution (i.e., the solution that maximizes joint payoffs), and compare this to the equilibrium solution. As we will see, the result from the static patent race models does not extend to our model.

We consider an extension of the basic model as follows. Instead of two players, we now consider a single player—the Planner—who receives a flow payoff given by $\pi_P(x) \equiv \pi(x) + \pi(-x)$. The state of the game, x, evolves according to a Wiener process with zero drift and variance $\sigma \in [2\underline{\sigma}, 2\bar{\sigma}]$, where σ is the Planner's choice. Specifically, a *Markov control* for the

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

[9] Although there is time in their models, we refer to them as static in the sense that players make a one-time decision regarding risk level.

Planner is a measurable map $\sigma : (-\infty, +\infty) \mapsto [2\underline{\sigma}, 2\bar{\sigma}]$. Fix a Markov control σ and define the expected discounted value of joint profits starting from $x(0)$ as:

$$U_P(x, \sigma) \equiv E\left[\int_0^\infty e^{-rt}\pi_P(x(t))\,dt \mid x, \sigma\right].$$

The social Planner then solves:

$$U_P^*(x) = \sup_\sigma U_P(x, \sigma).$$

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

$$r V_P(x) = \max_{\sigma(x)}\left[\pi_P(x) + \sigma(x)V_P''(x)\right].$$

Proposition 5. A solution to the Planner's problem exists, joint profits are maximized by a Markov control, and $rV_P(x) = rU_P^*(x)$, where V_P is continuous.

Proof. This is a standard stochastic control problem. Theorem 11.2.1 in Øksendal (1998) establishes the necessity of the HJB, while Theorem 11.2.2 establishes sufficiency. Finally, 11.2.3 yields that the maximum is obtained by a Markov control. *Q.E.D.*

The Planner's solution is *bang-bang* if an optimal control σ^* can be chosen such that $\sigma^* \in \{2\underline{\sigma}, 2\bar{\sigma}\}$. A Markov control σ 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 is constant on each subinterval (y_i, y_{i+1}). The Planner's solution is simple if there exists an optimal Markov control that is simple.

Lemma 2. The Planner's solution is simple and bang-bang.

Proof. Let σ^* be an optimal Markov control. Assume $rV_P(x) > \pi_P(x)$. The continuity of V_P and π then implies that there exists an ε such that $rV_P(y) > \pi_P(y)$ for all $y \in (x - \varepsilon, x + \varepsilon)$, and thus $r V_P''(y) < 0$ and $\sigma^*(y) = 2\underline{\sigma}$ on this interval. Like reasoning establishes that σ^* is equal to $2\bar{\sigma}$ on an open interval whenever $rV_P(x) < \pi_P(x)$. Finally, for any x such that $rV_P(x) = \pi_P(x)$ the Planner is indifferent across all σ, and thus we may choose an optimal Markov control $\hat{\sigma}$ such that $\hat{\sigma}(x) = \sigma^*(x)$ for all x such that $rV_P(x) \neq \pi_P(x)$ and $\hat{\sigma}(x) = 2\bar{\sigma}$ otherwise. *Q.E.D.*

□ **Planner's value function characterization.** Now that we know that an optimal control can be chosen such that σ is constant on open intervals, we can *explicitly* solve for the form of the value function. Consider any interval on which σ is constant. The HJB equation implies that:

$$r V_P(x) = \pi_P(x) + \sigma V_P''(x). \tag{2}$$

The general solution to this differential equation is:

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

where

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

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

(To verify this solution, note that it must satisfy

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

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 $\tilde{\pi}_P$ is bounded. To see this, take the first term in brackets and simplify:

$$\int_{-\infty}^{x} e^{\alpha(s-x)}\pi_P(s)\,ds = e^{-\alpha x}\int_{-\infty}^{x} e^{s}\pi_P(s)\,ds$$
$$\leq e^{-\alpha x}\int_{-\infty}^{x} e^{s}2\bar{\pi}\,ds$$
$$= \frac{2\bar{\pi}}{\alpha},$$

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 σ remained unchanged. Because $x(t)$ 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 $2\sigma t$ if σ does not change. Thus, we compute

$$E\left[\int_0^{\infty} e^{-rt}\pi_P(x(t))\,dt \mid x(0) = x, \sigma\right]$$
$$= \int_{-\infty}^{\infty}\int_0^{\infty} e^{-rt}\pi_P(s)(4\pi\sigma t)^{-\frac{1}{2}}e^{-\frac{(x-s)^2}{4\sigma t}}\,dt\,ds$$
$$= \int_{-\infty}^{x}\left[\int_0^{\infty}(4\pi\sigma t)^{-\frac{1}{2}}e^{-rt-\frac{(x-s)^2}{4\sigma t}}dt\right]\pi_P(s)\,ds + \int_{x}^{\infty}\left[\int_0^{\infty}(4\pi\sigma t)^{-\frac{1}{2}}e^{-rt-\frac{(x-s)^2}{4\sigma t}}dt\right]\pi_P(s)\,ds.$$

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_P. If the Planner chooses $\sigma(x)$ equal to a constant, the value function has the same form for all values of x, yet $V_P(x)$ is unbounded unless $a = 0$ and $b = 0$. Thus, if no one switched projects, $V_P(x) = \tilde{\pi}_P(x;\alpha)$.

As $\tilde{\pi}_P$ is the value when no one switches projects, then the other two terms must be the value to the Planner of the option to switch projects, which implies $a, b \geq 0$.

□ **The patient Planner case.** We know that the Planner will either choose the highest or lowest possible variance, as the solution is bang-bang. It turns out that there is a simple condition that determines which extreme the Planner will choose near $x = 0$ as long as the Planner is patient. As in ME, the Planner's choice will be determined by the local curvature of the profit function for high enough r. Thus, we focus on what happens for low r. Substituting π_P for π in the proof of Lemma 1 yields a similar result for the Planner's value.

Lemma 3. $\lim_{r\to 0} rV_P(x) = \underline{\pi} + \bar{\pi}$.

□ **Efficient variance choice when firms are "close".** Given the results in the last section, we can offer a simple condition that determines the Planner's choice of variance near $x = 0$ given enough patience.

Proposition 6. If $x^* > 0$, then $\forall x \in (-x^*, x^*)$, $\exists r^* > 0$ such that the Planner sets $\sigma(x) = 2\bar{\sigma}$ for all $r < r^*$. Conversely, if $x^* < 0$, then $\forall x \in (-|x^*|, |x^*|)$, $\exists r^* > 0$ such that the Planner sets $\sigma(x) = 2\underline{\sigma}$ for all $r < r^*$.

Proof. If $x^* > 0$, then $\pi(x) < (\underline{\pi} + \bar{\pi})/2$ for all $x < x^*$, while $\pi(-x) < (\underline{\pi} + \bar{\pi})/2$ for all $x > -x^*$ so that $\forall x \in (-x^*, x^*)$, $\pi_P(x) \equiv \pi(x) + \pi(-x) < \underline{\pi} + \bar{\pi} = \lim_{r\to 0} rV_P(x)$ (by Lemma 3). Finally, by the HJB equation for the Planner, $\pi_P(x) < rV_P(x) \Rightarrow V_P''(x) > 0 \Rightarrow \sigma(x) = 2\bar{\sigma}$. *Q.E.D.*

We are able to characterize the patient Planner's variance choice on the interval $(-|x^*|, |x^*|)$ but not outside of this interval. To do so, we would need to assume that π_P satisfies the single-crossing property. Notice that *none* of the examples we have presented satisfy this property.

□ **Equilibrium and efficiency.** We are now ready to compare the equilibrium outcome with the Planner's solution. Notice that the Planner will always choose the highest or lowest variance possible. Thus, whenever the players choose different variances in ME, the ME is inefficient. By Proposition 2, the players choose different variances outside of the interval $(-|x^*|, |x^*|)$ for r low enough, whereas by Propositions 2 and 6, the Planner's choice corresponds to the ME on this interval. Thus we have:

Corollary 2. For r low enough, total variance in ME is socially efficient in $[-|x^*|, |x^*|]$ and inefficient outside of this interval.

6. Extensions

■ There are a number of ways in which the simple model presented here could be extended. In this section we will consider three of them: making the variance term a more general function of player choices; adding exogenous drift; and adding in cost of variance to the flow payoffs.

Instead of the linear specification considered here, we could instead have the instantaneous variance be some more general function $\Sigma(\sigma_1, \sigma_2)$. As long as Σ is bounded away from 0 and ∞ and monotonic in both σ_1 and σ_2 individually, all of our results extend trivially. Thus, our results are not driven by our linear specification.

We assumed no drift in $x(t)$. The first step to relaxing this assumption would be to assume some exogenous drift, $\mu(x)$. Our existence results extend immediately with this change. The low r characterization results are a bit more delicate. With drift, the Bellman equation becomes

$$rV_1(x) = \max_{\sigma_1(x)}[\pi(x) + \mu(x)V_1'(x) + (\sigma_1(x) + \sigma_2(x))V_1''(x)].$$

There are two issues: first, the limit of $V_1'(x)$ must be characterized. Intuitively, this should tend to 0 as r tends to 0, but the proof is not as straightforward as the proof for $rV_1(x)$.

If $V_1'(x)$ tends to 0, then our result would extend as long as $\underline{\pi} < \lim_{r\to 0} rV_i(x) < \bar{\pi}$. Examining the proof of Lemma 1, the key is what happens to the mean of $x(t)$ relative to the standard deviation as $t \to \infty$. More specifically, $x(t)$ will be distributed normally with mean $m(t, x(0), \mu)$ and standard deviation $s(t, x(0), \sigma)$. To retain $\underline{\pi} < \lim_{r\to 0} rV_i(x) < \bar{\pi}$, we need

$$\lim_{t\to\infty} m(t, x(0), \mu)/s(t, x(0), \sigma)$$

bounded. Thus, we need $m(t, x(0), \mu)$ and $s(t, x(0), \sigma)$ to grow at the same rate. In our model, $s(t, x(0), \sigma)$ is of the order $\sqrt{t}$. If we simply assumed that $\mu(x) = \mu$ (i.e., a constant), we would have $m(\mu, t) = \mu t$ and $\lim_{t\to\infty} m(\mu, t)/s(t, x(0), \sigma) = \infty$. One natural way to deal with this issue would be to make the process mean-reverting.

Another extension would be to include a cost function for different variance choices: $c(\sigma_i)$. What shape should such a cost function be? In most applications, the cost of setting either very low variance or very high variance is likely prohibitive. Thus, one might consider a U-shaped cost function. One immediate technical difficulty is existence of a pure strategy ME. Our straightforward proof fails with the addition of cost of variance, but can be rescued by allowing mixing. Specifically, modify the Bellman equation by making flow payoffs $\pi(x) - c(\sigma_1(x))$, which implies a first-order condition $V_1''(x) = c'(\sigma_1(x))$, and a satisfied second-order condition $-c''(\sigma_1(x)) < 0$. Thus, from the first-order condition and the U-shaped cost function, $\sigma_1(x)$ will be monotonically increasing in $V_1''(x)$. Again, $rV_i(x)$ will tend toward the average of the flow payoffs at the extremes. Intuitively, $\sigma_1(x)$ will be increasing in x.

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 not 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 article, 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 chooses low variance.

We also show that, with enough patience, the ME outcome is efficient when players are close enough together and inefficient when players are sufficiently far apart.

Appendix

☐ *Proof of Lemma 1.*

Proof. We shall establish the result for rV_1; the proof is nearly identical for rV_2. Let $f(x' \mid t, x(0), \sigma)$ be the probability that $x(t) = x'$ at time t, given starting value $x(0)$ and Markov control σ. As $x(t)$ is an Ito process, f is a normal density with mean $x(0)$ and some standard deviation $s(t, x(0), \sigma)$, where $\lim_{t\to\infty} s(t, x(0), \sigma) = \infty$. To simplify notation, let

$$g(t \mid x(0), \sigma) = \int_{-\infty}^{\infty} \pi(x')f(x' \mid t, x(0), \sigma)\,dx'.$$

Claim 1. $\lim_{t\to\infty} g(t \mid x(0), \sigma) = (\underline{\pi} + \bar{\pi})/2$.

Proof of Claim 1. Fix any $\bar{x} > 0$, then:

$$\int_{\bar{x}}^{\infty} f(x' \mid t, x(0), \sigma)\,dx' = \frac{1}{2} + \frac{2}{\sqrt{\pi}} \int_0^{\frac{x(0)-\bar{x}}{s(t,x(0),\sigma)\sqrt{2}}} e^{-z^2}dz.$$

So, $\lim_{t\to\infty} \int_{\bar{x}}^{\infty} f(x' \mid t, x(0), \sigma)\,dx'$

$$= \frac{1}{2} + \frac{2}{\sqrt{\pi}} \lim_{t\to\infty} \int_0^{\frac{x(0)-\bar{x}}{s(t,x(0),\sigma)\sqrt{2}}} e^{-z^2}dz$$

$$= \frac{1}{2},$$

where the last line follows from the fact that $\lim_{t\to\infty} s(t \mid x(0), \sigma) = \infty$.

Similar steps establish that:

$$\lim_{t\to\infty} \int_{-\infty}^{-\bar{x}} f(x' \mid t, x(0), \sigma)\,dx' = \frac{1}{2} \quad \forall \bar{x} > 0.$$

Thus,

$$\lim_{t\to\infty} \int_{-\bar{x}}^{\bar{x}} f(x' \mid t, x(0), \sigma)\,dx' = 0 \quad \forall \bar{x} > 0,$$

and so,

$$\lim_{t\to\infty} \int_{-\bar{x}}^{\bar{x}} \pi(x')f(x' \mid t, x(0), \sigma)\,dx' = 0 \quad \forall \bar{x} > 0.$$

Together these imply that

$$\lim_{t\to\infty} g(t \mid x(0), \sigma) = \frac{1}{2}\lim_{x\to\infty} \pi(x) + \frac{1}{2}\lim_{x\to-\infty} \pi(x) = (\underline{\pi} + \bar{\pi})/2,$$

and we have established Claim 1.

Now for any optimal Markov control σ, we have:

$$rV_1(x) = \int_0^{\infty} re^{-rt}g(t \mid x(0), \sigma)\,dt.$$

Integration by parts yields:

$$r V_1(x) = \left[-e^{-rt} g(t \mid x(0), \sigma)\right]_{t=0}^{\infty} + \int_0^{\infty} e^{-rt} g_t(t \mid x(0), \sigma)\, dt$$

$$= 0 + g(0 \mid x(0), \sigma) + \int_0^{\infty} e^{-rt} g_t(t \mid x(0), \sigma)\, dt.$$

So that,

$$\lim_{r \to 0} r V_1(x) = g(0 \mid x(0), \sigma) + \lim_{r \to 0} \int_0^{\infty} e^{-rt} g_t(t \mid x(0), \sigma)\, dt$$

$$= g(0 \mid x(0), \sigma) + \lim_{t \to \infty} g(t \mid x(0), \sigma) - g(0 \mid x(0), \sigma)$$

$$= \lim_{t \to \infty} g(t \mid x(0), \sigma),$$

which by Claim 1 equals $(\underline{\pi} + \bar{\pi})/2$. *Q.E.D.*

References

ATHEY, S. AND SCHMUTZLER, A. "Investment and Market Dominance." *RAND Journal of Economics,* Vol. 32 (2001), pp. 1–26.

BASAR, T. AND OLSDER, G.J. *Dynamic Noncooperative Game Theory.* Philadelphia: Society for Industrial and Applied Mathematics, 1998.

BERGEMANN, D. AND VÄLIMÄKI, J. "Entry and Vertical Differentiation." *Journal of Economic Theory,* Vol. 106 (2002), pp. 91–125.

BERGIN, J. AND MACLEOD, B. "Continuous Time Repeated Games." *International Economic Review,* Vol. 34 (1993), pp. 21–37.

BHATTACHARYA, S. AND MOOKHERJEE, D. "Portfolio Choice in Research and Development." *RAND Journal of Economics,* Vol. 17 (1986), pp. 594–605.

BOLTON, P. AND HARRIS, C. "Strategic Experimentation: The Undiscounted Case." In P. Hammond and G. Myles, eds., *Incentives Organization and Public Economics: Essays in Honour of James Mirrlees.* Oxford: Oxford University Press, 2000.

BUDD, C., HARRIS, C., AND VICKERS, J. "A Model of the Evolution of Duopoly: Does the Asymmetry between Firms Tend to Increase or Decrease?" *Review of Economic Studies,* Vol. 60 (1993), pp. 543–573.

CABRAL, L.M.B. "Increasing Dominance with No Efficiency Effect." *Journal of Economic Theory,* Vol. 102 (2002), pp. 471–479.

———. "R&D Competition When Firms Choose Variance." *Journal of Economics and Management Strategy,* Vol. 12 (2003), pp. 139–150.

——— AND RIORDAN, M.H. "The Learning Curve, Market Dominance and Predatory Pricing." *Econometrica,* Vol. 62 (1994), pp. 1115–1140.

CYERT, R.M. AND MARCH, J.G. *A Behavioral Theory of the Firm.* Englewood, NJ: Prentice-Hall, 1963.

DIXIT, A. AND PINDYCK, R. *Investment under Uncertainty.* Princeton, NJ: Princeton University Press, 1994.

DUTTA, P. "What Do Discounted Optima Converge To?" *Journal of Economic Theory,* Vol. 55 (1991), pp. 64–94.

ERICSON, R. AND PAKES, A. "Markov-Perfect Industry Dynamics: A Framework for Empirical Work." *Review of Economic Studies,* Vol. 62 (1995), pp. 53–82.

FERSHTMAN, C. AND PAKES, A. "A Dynamic Oligopoly with Collusion and Price Wars." *RAND Journal of Economics,* Vol. 31 (2000), pp. 207–236.

HARRIS, C. "Generalized Solutions of Stochastic Differential Games in One Dimension." Working Paper, Nuffield College, 1993.

——— AND VICKERS, J. "Racing with Uncertainty." *Review of Economic Studies,* Vol. 54 (1987), pp. 1–21.

HÖRNER, J. "A Perpetual Race to Stay Ahead." *Review of Economic Studies,* Vol. 71 (2004), pp. 1065–1088.

JUDD, K.L. "Closed-Loop Equilibrium in a Multi-Stage Innovation Race." *Economic Theory,* Vol. 21 (2003), pp. 673–695.

KHANNA, T. AND IANSITI, M. "Firm Asymmetries and Sequential R&D: Theory and Evidence from the Mainframe Computer Industry." *Management Science,* Vol. 43 (1997), pp. 405–421.

KLETTE, T.J. AND DE MEZA, D. "Is the Market Biased against Risky R&D?" *RAND Journal of Economics,* Vol. 17 (1986), pp. 133–139.

ØKSENDAL, B. *Stochastic Differential Equations.* Berlin: Springer-Verlag, 1998.

SIMON, L. AND STINCHCOMBE, M. "Extensive Form Games in Continuous Time: Pure Strategies." *Econometrica,* Vol. 57 (1989), pp. 1171–1214.