A Theory of Stopping Time Games with Applications to Product Innovations and Asset Sales

Dutta, Prajit K. and Aldo Rustichini

Working Paper No. 263
January 1991

A Theory of Stopping Time Games with Applications to Product Innovations and Asset Sales

Prajit Dutta and Aldo Rustichini

Rochester Center for Economic Research
Working Paper No. 263

A Theory of Stopping Time Games With
Applications to Product Innovation and Asset Sales*

Prajit Dutta
Columbia University

Aldo Rustichini
Northwestern University

Rochester Center for Economic Research
Working Paper No. 263

January 1991

Abstract. In this paper, the pure strategy subgame perfect equilibria of general stopping time games are studied. It is shown that there always exists a natural class of Markov perfect equilibria, called stopping equilibria. Such equilibria can be computed as a solution of a single agent stopping time problem. A complete characterization of stopping equilibria is presented. Conditions are given under which the outcomes of such equilibria span the set of all possible outcomes from perfect equilibria. Two economic applications of the theory, product innovations and the timing of asset sales, are discussed. Finally, we show that if players can commit themselves for some length of time, subgame perfect Nash equilibrium may fail to exist.

1. Introduction

In this paper we study the pure strategy subgame perfect equilibria of a general class of *stopping time games*. A stopping time game is described by a stochastic process $\{X_t : t \geqslant 0\}$ and payoff functions l and f. Without intervention by the players, a state variable evolves according to the given stochastic process. At any time, either of two players can "stop" the process. If player i stops the process at state x then his payoff is $l(x)$ and that of the other player $f(x)$. Simultaneous moves result in a payoff of $\alpha l+(1-\alpha)f$, for some fixed $\alpha \in (0,1)$[1]. Our formulation in fact includes the possibility that the process continues to evolve after i's move, that j optimally selects a move thereafter and the stochastic process continues to evolve after both of their moves. In this case the payoff functions are the present discounted values of future returns.

*Parts of this paper were presented at seminars at Chicago, Rochester and Rutgers and we acknowledge helpful comments.

Almost any of the well known *games of timing* is an example of a stopping time game, with the added simplification that the "stochastic process" is just time itself, i.e. $x(t) \equiv t$ for all t. Two classic examples of games of timing are the *war of attrition* and *pre-emption games* (see, Fudenberg-Tirole (1989), chapter 6). In the (symmetric) war of attrition, $f(x) > l(x)$, for all x, whereas in pre- emption games, $f(x) < l(x)$ at least over some states. Games of timing have been extremely useful in understanding issues as diverse as patent races (Fudenberg et al. (1983), Dasgupta-Stiglitz (1980)), exit from declining industries (Ghemawat-Nalebuff (1985), Fudenberg-Tirole (1986)) and the adoption and diffusion of new technologies (Fudenberg-Tirole (1985)). Stopping-time games were introduced in Bensoussan-Friedman (1977).[2] In many applications, a pure game of timing framework is unduly restrictive. Generalizations in at least those directions appear essential. Firstly, one needs to incorporate a (possibly) *stochastic* mechanism to capture the effect of uncertainty on the stopping problem.

For example, in evaluating a trading mechanism, price uncertainty is a crucial element of the decision problem, when to buy or sell an asset. Secondly, many problems involve, naturally, a *number of dimensions* other than time alone. For example, it may be important to analyze the effect of quality or production costs or size of the market, in addition to time, in evaluating the potential of a new product innovation. Further, in some applications, the assumption that the state variable is monotonically increasing, which is true when the state is nothing but time, is unrealistic. Our formulation of the problem allows such generality. Moreover, we do not make any qualitative assumptions on the payoff functions other than continuity. This contrasts with all of the above literature, in which concavity and/or monotonicity restrictions are typically imposed on l and f.

Formally, stopping time games are examples of stochastic games (see, for example, Parthasarathy (1973) or Fudenberg-Tirole (1989)) with the proviso that each player has an irreversible action ("stop") that either terminates the game or fixes the player's payoff in the consequent subgame. In the absence of such an action, the game environment evolves according to the exogeneous process $\{X_t : t \geqslant 0\}$. In particular, unlike a game of timing, although there is only one relevant action, there are different environment histories on which players can coordinate in choosing that action. So, in principle, there are many perfect equilibria of a stopping time game, with a multiplicity of equilibria sustainable by perhaps folk theorem-like logic.

In our first result (Theorem 1), we completely characterize a sub-class of the perfect equilibria in a stopping time game, a class we call *stopping equilibria*. Stopping equilibria are a subset of the set of Markov Perfect Equilibria (MPE) of the game. These equilibria have several useful properties. First, they are solutions to *single agent stopping problems* and hence can be computed by solving an optimization problem (rather than more intractable fixed point exercises). Secondly, different stopping equilibria can be naturally ordered by the implied *degree of cooperation* for different regions. Let us clarify this point. We let $C_k, k \in K$ denote the countable family of (connected) sets on which[3] $l > f$. Each set C_k should be thought of as a potential cooperation region. On this set each player has a private incentive to stop and earn a reward l, which exceeds the reward to the follower, f. A player will cooperate and desist from such move only if he is sure that the other player will not preempt him. Could players desist from moving on a potential *cooperation region*

C_k, so that they both benefit from the state evolving, as a consequence, to a mutually more advantageous position? Each stopping equilibrium corresponds to a particular subset of K, that on which the two players do cooperate. We give a simple *incentive compatibility* condition to determine whether any such subset can be cooperated on, and the condition is in terms of the stopping value generated by *not* cooperating on the complement of this subset. Naturally, different stopping equilibria generate a *partial order on payoffs*. Further, the number[4] of such equilibria is clearly bounded by the cardinality of the power set of K. Finally, a stopping equilibrium always exists.

Next, we turn to the issue: how large a subset of equilibrium payoffs or outcomes is it, that stopping equilibria generate? Clearly, it would be of interest to know under what circumstances the computational ease of stopping equilibria is not compromised by ruling out too many other outcomes. Proposition 3 shows that if $\{X_t : t \geqslant 0\}$ is a single dimensional, increasing process,[5] the set of *outcomes* to stopping equilibria is almost the same set of outcomes, that the more general class of MPE generate. In a sense made precise in the sequel, almost here means that the outcomes generated by a variant of stoping equilibria, stopping* equilibria, are precisely those generated by MPE. Hence, from an observational viewpoint, the two classes of equilibria are equivalent. Further, in Proposition 4 we show that if the environment is single- dimensional and evolves deterministically, then the set of outcomes under stopping* equilibria is the same set as under that generated by *any* perfect equilibrium. It follows from Proposition 3 and 4 that all equilibrium outcomes in games of timing are, in fact, outcomes of stopping* equilibria. In particular, all equilibrium outcomes can be derived by solving a stopping time optimization problem, rather than through a fixed point argument.

All of these results are in Section 4 which follows the game formulation in Section 2 and the existence and characterization of stopping equilibria in Section 3.

We consider two economic applications of the theory in Section 5. In the first one, we briefly present some results from a *product innovation* problem, which we have developed in detail elsewhere (Dutta-Lach-Rustichini (1990)). In the second, we analyze trading of an asset, in particular exploring the intuition: when there is more than one informed trader in a market, the implied competition might lead to early sales. The two models are simple abstractions, but serve to illustrate the applicability of the results and framework in the general case.

Finally, in Section 6 we turn briefly to a related question. In the stopping time game formulation we explicitly assume that no commitments can be made. Since the equilibria turn out to have interesting implications for the degree of cooperation that is achievable, we next ask how the introduction of commitment horizons changes the equilibrium set. Such incorporation of commitment possibilities has a surprising consequence. Equilibria may *fail to exist* for long commitment horizons. The intuition is very similar to the cause of non-existence in Hotelling spatial models (for example, see D'Aspremont-Gabsewicz-Thisse (1979)) and we discuss this briefly. That commitment may destroy equilibrium existence, seems to be a new and hitherto, unnoticed phenomenon.

2. Stopping Time Games: The Model

This section formulates a stopping time game. A detailed explanation of this framework may be found in Bensoussan-Friedman (1977) who analyzed the issue of the existence of

Nash equilibria in such games, without the additional constraint of the equilibria being subgame perfect.

2.1 Notation and A Preliminary Result

Let (Ω, F) be a measurable space and X be a topological space. A progressively measurable stochastic process is a triple (X_t, F_t, P^x) where F_t is an increasing family of sub σ-fields of a σ-field F and $X_t : \Omega \to X$ is measurable with respect to F_t, for all t. A stochastic process with range $\mathbf{R}^n$ is said to be *increasing* if for all x and t, $P^x(x_t \geqslant x) = 1$. A *Markov Time* in a map $T : \Omega \to \mathbf{R}_+$, with $\{T \leqslant t\} \in F_t$ for every t, and $T < +\infty$ a.s. *Extended Markov Times* are Markov times with $T \leqslant +\infty$ a.s.; they are denoted by M.

An optimal discounted stopping time problem is the choice of Markov times to maximize discounted lifetime returns; i.e. is given by

$$\max_{T \in M} e^{-\delta T} g(x_T).$$

The following result is used repeatedly in the sequel (see for instance Fakeev, (1973)).

THEOREM 0. *Let X_t be continuous with probability 1 and suppose g is a continuous function. Then the optimal discounted stopping time problem has a solution given by*

$$T \equiv \inf\{s : s \geqslant 0, \quad x_s \in A\}$$

for some fixed Borel subset A of $\mathbf{R}^n$.

In the game there are two players, 1 and 2. The index i will refer to a generic player. In all statements involving player i, j will refer to the other player.

2.2 Stopping Time Game

An abstract *symmetric* leader-follower game or stopping time game is given by a triple $((X_t, F_t, P^x); l, f)$ where

i) (X_t, F_t, P^x) is a stochastic process with range given by a set $X \subseteq \mathbf{R}^n$, with the Borel σ-field $B(X)$.

ii) l, f are real valued functions on X (and stand for the payoff to the leader and to the follower respectively).

For the present, we analyze the symmetric game. We return to the more general asymmetric formulation at the end of Section 3. The game proceeds as follows: in the absence of any action by either player, the game environment evolves according to the stochastic process $(X_t, t \geqslant 0)$. Each player has a single action, which we shall interchangeably call "stop" or "entry." If player i stops or enters at x, then i's payoff is $l(x)$, and that of $j, f(x)$. One could either think of play terminating at x, with l and f as terminal payoffs. Alternatively, after i's entry, j solves a single agent stopping time problem. The value for that problem is $f(x)$, and $l(x)$ would then be the implied return for the first mover. Simultaneous entry yields $b(x) \equiv \alpha l(x) + (1-\alpha) f(x), \alpha \in (0,1)$.

We shall always assume the following conditions:

(A1) The stochastic process is a Markov process with continuous sample paths.

(A2) l, f are continuous functions.

In all of the applications mentioned above these conditions are satisfied.

2.3 Strategies

At time $t \geqslant 0$, the *history* of the play has two components: the sample path of the state (which depends on the action of "nature") and the actions of the two players. Formally at time t the set of histories is

$$H_t \equiv C([0,t],X) \times ([0,t) \cup \{+\infty\})^2.$$

A typical element of H_t is $h = (x, t_1, t_2)$ where x is the path of the state until t, and $t_i \leqslant t$ denotes the time at which player i has entered. With the convention $t_i = +\infty$ if player i has not entered, this gives a complete description of the play until t.

We denote by G_t the smallest σ-algebra on $C([0,+\infty),X)$ generated by the finite cylinders of the type

$$\{Y \in C : Y(t_i) \in B_i, \quad B_i \in \mathcal{B}(X), \quad i = 1, 2, \ldots, m; \quad t_i \leqslant t\}.$$

where $\mathcal{B}(X)$ are the Borel subsets of X. A *strategy* for a player is a family of functions $\{\sigma_t\}_{t \geqslant 0}$, where each $\sigma_t : H_t \to \{0,1\}$ is progressively measurable with respect to G_t. Here 0 stands for "no entry" and 1 for "entry" respectively. Since a player can move only once, a further restriction on a strategy is that, for every h_t s.t. $t_i < \infty$, $\sigma_t(h_t) \equiv 1$.[6]

A strategy is said to be *stationary Markovian* if actions depend only on the current state. In particular, if a player i has not moved yet at t then whether or not a move happens is decided solely as a function of x_t. Formally, a stationary Markovian strategy is a measurable function, $\varphi : X \to \{0,1\}$, such that $\sigma_t(h) = \varphi(x_t)$, for all t.

As in other stochastic games, Markovian strategies incorporate all payoff relevant factors (and hence imply a minimal degree of rationality), yield equibria which are equilibria even when the larger class of history dependent equilibria are admissible and are computationally and analytically more tractable (for these and additional arguments, see Maskin-Tirole (1988) or Fudenberg-Tirole (1989, chapter 5).

To any pair of strategies and any initial state is associated a stopping time $T : \Omega \to \mathbf{R}_+ \times \{0,1,2\}$. The first coordinate indicates the moment in which some player enters; the second coordinate indicates the identity of the player who enters: 0 denotes simultaneous entry. We write therefore $T = (t, i)$ to denote the time of the move, and the player who moved, respectively.

The payoff in the game with $X_0 = x$ a.s. is defined for player 1, say, by:

$$\pi_1(\sigma^1, \sigma^2)(x) = E_x e^{-\delta t}\{l(x_t)\chi_{\{i=1\}} + f(x_t)\chi_{\{i=2\}} + b(x_t)\chi_{\{i=0\}}\}.$$

2.4 Equilibrium

The game is one of complete information and the equilibrium concept we adopt is that of subgame perfect equilibrium. Note that a proper subgame starts at the end of every partial history h_t, i.e. in the node originating with x_t. Let $\sigma_i(h_t)$ denote the continuation after history h_t, of strategy σ_i. Then, $(\sigma_1^*(h_t), \sigma_2^*(h_t))$ form a *Nash equilibrium* after h_t if

$$\pi_i(\sigma_i^*(h_t), \sigma_j^*(h_t))(x_t) \geqslant \pi_i(\sigma_i, \sigma_j^*(h_t))(x_t), \text{ for every } \sigma_i.$$

The pair (σ_1^*, σ_2^*) forms a *subgame perfect equilibrium* if $\sigma_1^*(h_t), \sigma_2^*(h_t)$ forms a Nash equilibrium for all histories h_t and all t. If σ_1^*, σ_2^* are (stationary) Markovian, they are said to form a *Markov Perfect Equilibrium* (MPE).

3. Stopping Equilibria

As was clear from the discussion of the stopping time game and admissible strategies in Section 2, the set of equilibria is potentially large.

In this section we characterize a particular sub-class of MPE, which we call *stopping equilibria.*

Stopping equilibria will be shown to be solutions to a single agent stopping problem. The intuition behind these equilibria is simple. Suppose player j is a "passive" player, in that he never unilaterally moves. He does move if the other player does, when faced with an immediate loss, i.e. if i moves at x such that $l(x) > f(x)$, then so does j. Given such a strategy of j, i has to decide his unilateral moves, if any. What these will be is completely determined by which of the potential cooperation regions C_k, the two players can (implicitly) agree not to move on. Such an "agreement" is enforceable for i if his unilateral moves are consistent with the degree of cooperation and vice-versa. For j, they are enforceable if cooperation is incentive compatible.

All this is shown to be captured by a single stopping problem and a pair of associated incentive conditions.

3.1 A Modified Stopping Problem

Let $\{C_k\}_{k\in K}$ be the connected components of $\{x \in X : f(x) < l(x)\}$. Clearly K is at most countable, since l and f are continuous functions.

For a finite subset of K, I say, we define the modified stochastic process $\{y_t^I, t \geqslant 0\}$ as follows. Let H denote the hitting time of the set $\cup_{i\in I} C_i$, that is:

$$H(x) = \inf\{t \geqslant 0 : \quad x_t \in \cup_{i\in I} C_i; x_0 = x\}.$$

Now define the stochastic process $\{y_t, F_t, P_x\}$ as

$$y_t = x_t \chi_{\{t<H\}} + x_H \chi_{\{t \geqslant H\}}$$

In other words, the process $\{y_t, t \geqslant 0\}$ is "absorbed" at the boundary of the region $\cup_{i\in I} C_i$. For notational convenience, denote $C_I \equiv \cup_{i\in I} C_i$ and $C_{-I} = \cup_{i\in K\setminus I} C_i$. C_{-I} is a potential cooperation region.

We then define the optimal stopping time problem : choose a Markov time M so as to

$$PL_I \quad \max_T E_x e^{-\delta T} l(y_T^I).$$

As recalled above (Theorem 0), the problem PL_I has a solution with an optimal stopping time T defined by $T = \inf\{t \geqslant 0, X_t \in A_I\}$ for some Borel subset A_I of X.

We shall denote $V_l^I(x)$ the solution of the PL_I problem. Recall $V_l^I(x) \geqslant l(x)$ for $x \in C_I$, and $A_I = \{x : V_l^I(x) = l(x)\}$.

3.2 Stopping equilibria

THEOREM 1.

i) *The pair of (stationary) Markovian strategies* (σ_1^I, σ_2^I), *defined as*

$$\sigma_1^I(x_t) = \begin{cases} 0 & \textit{if } x_t \notin A_I \cup C_I \\ 1 & \textit{otherwise} \end{cases}$$
$$\sigma_2^I(x_t) = \begin{cases} 0 & \textit{if } x_t \notin C_I \\ 1 & \textit{if } x_t \in C_I \end{cases}$$

is a subgame perfect equilibrium of the stopping time game iff

$$V_l^I(x) > l(x), \quad x \in C_{-I} \tag{1}$$

ii) *A subgame perfect equilbrium for the game always exists, and is given by the pair of strategies* (σ_1^K, σ_2^K).

iii) *If* K *is a finite set of, say,* n *elements, there are at most* 2^{n+1} *equilibria of the type described in 1.*

PROOF: We drop the index I for simplicity.

Notice that from the condition $V_l(x) > l(x)$ for $x \in C_{-I}$, we have $x_T \in X \backslash C_{-I}$ almost surely, or $f(x_T) \geqslant l(x_T)$ almost surely. We now prove that the pair (σ_1, σ_2) gives a subgame perfect equilibrium. Notice first that since each component of the pair only considers the present states (so that all subgames starting from the same state are identical), we only need prove that they give a Nash equilibrium on any subgame beginning at any $\xi \in X$. If $\xi \in C_I$, then clearly (σ_1, σ_2) is an equilibrium pair.

For any $\xi \in X \backslash C_I$ we have

$$l(\xi) \leqslant V_l(\xi) = E_\xi e^{-\delta T_l} l(x_{T_l}) \leqslant E_\xi e^{-\delta T_l} f(x_{T_l}) = V_f(\xi). \tag{3}$$

The first inequality follows from the hypothesis, if $\xi \in C_{-I}$ (and here the inequality is strict), and from the basic properties of V_l for every other x. The inequality $V_l(x) \geqslant l(x)$ (with a strict inequality if $x \notin A$) insures that the leader will not deviate and make his move before the state reaches A. On the other hand, the inequality $E_\xi e^{-\delta T_l} f(x_{T_l}) \geqslant l(\xi)$ if $\xi \in X \backslash C_I$ insures that the designated follower will not preempt before the state reaches the set A.

The converse is obvious: for $x \in C_{-I}$, if $V_l(x) = l(x)$ then the designated leader would enter, and since $f(x) < l(x)$ the best response of the follower is to enter too. If C_{-I} is empty, (1) is vacuously satisfied. From this ii) follows. ∎

It follows from the theorem that the equilibria can be given a partial order induced by the inclusive partial order over the subsets of K. More precisely:

COROLLARY 2. *If* $I_1 \subset I_2$, *and the corresponding pairs* $(\sigma_1^{I_i}, \sigma_2^{I_i})$, $i = 1, 2$, *constitute an equilibrium, then* $V_l^{I_1} \geqslant V_l^{I_2}$. *In particular* V_l^K *is the worst equilibrium for the leader.*

PROOF: Immediate from the fact that the equilibrium payoff for the leader is determined by the solution of the PL_{I_1} problem, and in PL_{I_1} the stopping time $T^* = \inf\{t \geqslant 0 : x_t \in A_{I_2} \cup \cup_{i \in I_2} C_i\}$ gives $E_x e^{-\delta T^*} l(x_{T^*}^{I_1}) = V_l^{I_2}(x)$. ∎

4. Outcomes in Equilibrium

How restrictive is the behavior implied by stopping equilibria? The question is important, of course, because the number of equilibria, in this as in all games, may be very large. Since anyway we can only observe the *outcomes* of an equilibrium, it is natural to restrict the analysis to them. In this section we ask how large is the set of all possible equilibrium outcomes. Further we want to determine how many of these equilibrium outcomes correspond to the outcomes of stopping equilibria. The set of *outcomes* for an initial state x and a pair of strategies $\sigma = (\sigma_1, \sigma_2)$, $O_x(\xi)$ can be identified with the implied stopping set

$$O_x(\sigma) = \{z \in X : \sigma_i(z) = 1, i = \text{ 1 or 2, and } P^x(x(T) = z) > 0\}. \tag{4}$$

Notice that we impose the condition $P^x(x(T) = z) > 0$, i.e. that the game goes to state z with positive probability.

It is easy to see that in the general class of stopping games, with no restriction imposed on the nature of the stochastic process, the equilibrium outcomes may have a very arbitrary nature, more complicated than the simple form imposed by the stopping equilibria. Consider however the following simplifying hypothesis:

(A3) $\{X_t, t \geqslant 0\}$ is a single dimensional, increasing stochastic process.

Games of timing are, of course an interesting instance which satisfies this condition. Under this simplification, the set $O_x(\sigma)$ for a markovian strategy σ consists of a singleton; $O_x(\sigma) = \inf_{z \geqslant x} \{z \in X : \sigma_i(z) = 1, i = 1 or 2\}$.

The following example shows that stopping equilibria are not quite sufficient (although Proposition 3. will show that they almost are) to generate the equilibrium outcomes of all MPE.

Example. $X \equiv R_+$, $\dot{x} = 1$. Suppose that $0 < \alpha < \beta$, and let $(0, \alpha)$ be an interval on which $l > f$, and then on (α, β), $l < f$. Finally, $l = f$ on $\{0, \alpha, \beta\}$and outside $[0, 1]$.

Suppose that

$$e^{-\delta(\alpha - x)} l(\alpha) > l(x), \quad x \in (0, \alpha) \tag{5}$$

$$e^{-\delta(\beta - x)} l(\beta) > l(x), \quad x \in [\alpha, \beta) \tag{6}$$

$$e^{-\delta(x - \beta)} l(x) < l(\beta), \quad x \in (\beta, \infty). \tag{7}$$

It is clear that $\sigma_i(x) = 1$ for $i = 1, 2$ if and only if $x = \{0, \alpha, \beta\}$ constitutes a MPE. For any $x \in (0, \alpha)$, the equilibrium outcome is α. On the other hand the outcomes corresponding to the only two stopping equilibria are either x, or β. Notice that in this example the outcome α cannot be generated by stopping equilibria because they require the players to either cooperate on the entire interval $(0, \alpha)$ or not at all.

This example suggests a minor modification of the definition of stopping equilibria, which we now introduce. Stopping equilibria are indexed by the degree of potential cooperation that is "offered" by a passive player on the set where $l > f$. An immediate extension of this concept would be to consider, additionally, potential cooperation on the set of states where leading and following give the same payoff, that is where $l = f$. Denote $S \equiv \{x \in X : l(x) = f(x)\}$. We shall define a *stopping* equilibrium* through strategies which induce non-cooperation over the region C_I and also on an additional region of non-cooperation, $S^* \subset S$. In other words, let

$$H^*(x) = \inf\{t \geqslant 0 : x_t \in C_I \cup S^*; X_0 = x\}$$

and as in Section 3.1 define

$$y_t = z_t \chi_{\{t<H^*\}} + x_H \chi_{\{t \geqslant H^*\}}.$$

Notice that the cardinality of the set of stopping* equilibria may be in principle very large. It is easy to see, however, that if the set of points where the two payoff functions l and f are equal is finite then the number of stopping* equilibria is again finite.

Let $V_l^{I^*}$ define the value to stopping the stochastic process $\{y_t : t \geqslant 0\}$ and suppose that A_I^* is the stopping set (well-defined by Theorem 0). Then, consider the strategies

$$\sigma_1^*(x) = \begin{cases} 0 & x \notin A_I^* \cup C_I \cup S^* \\ 1 & \text{otherwise} \end{cases}$$
$$\sigma_2^*(x) = \begin{cases} 0 & x \notin C_I \cup S^* \\ 1 & \text{otherwiese} \end{cases}$$

The above strategies are a candidate for a stopping* equilibrium. The incentive compatibility condition is precisely (1). Stopping* equilibria are clearly a more general class of equilibria than stopping equilibria, in that they reduce the latter when $S^* = \emptyset$. Further, they have the same conceptual and analytical clarity that stopping equilibria possess. We can now prove:

PROPOSITION 3. *Suppose (A3) holds, and consider any MPE σ and an initial state x. Then, there is a stopping* equilibrium σ^* such that $O_x(\sigma) = O_x(\sigma^*)$.*

The proof of proposition 3. can be found in Appendix A.

Notice that a generic MPE of the stopping game can be described as follows: there are stopping sets $S_i, i = 1, 2$ such that $\sigma_i(x) = 1$ if and only if $x \in S_i$. Stopping* equilibria are constrained to stopping sets of the special kind: $S_j = C_I \cup S^*$ and consequently $S_i = A_I^* \cup C_I \cup S^*$). It should come as a surprise to the reader that although stopping sets can in principle be extremely arbitrary subsets of the state space, in equilibrium they can be restricted to the very regular stopping sets of the stopping* equilibria. The crucial result we use in the proof is Lemma A.1 in Appendix A. This result shows that on any interval on which $\{f > l\}$, only one of the players can be active. This yields the conclusion that without loss of generality $S_i \cap (C_k \cup S)^c$ is non empty for either $i = 1$ or 2 but not both. ¿From that Proposition 3 follows.

MPE are, in turn, a subset of the whole class of subgame perfect equilibria in stopping time games. We present a result now which gives conditions under which the set of outcomes under this much more general class of equilibria is precisely the same set of outcomes as under stopping equilibria.

(A4) $X \subseteq \mathbf{R}$ and $\{X_t; t \geqslant 0\}$ is a deterministic process.

Note again that games of timing are a case in which this condition is satisfied. Note that the process need not be increasing and, indeed for a deterministic increasing process, there is really no distinction between history dependent and memoryless strategies. Given that the process is deterministic, there is a unique history for any initial state x_0. Denote this history $h_t(x_0)$. Then, denote $T \equiv \inf_{t \geqslant 0}\{\sigma_i(h_t) = 1, i = 1 \text{ or } 2\}$, and

$$O_{x_0}(\sigma) = \{x(T)\}.$$

PROPOSITION 4. *Suppose that (A4) holds. Then, for any initial state x_0 and a subgame perfect equilibrium σ, there is a stopping* equilibrium σ^* such that*

$$O_{x_0}(\sigma) = O_{x_0}(\sigma^*)$$

The proof of Proposition 4. can also be found in Appendix A.

5. Product Innovation and Asset Sales

We present now two economic applications of the general theory. In the first we study product innovation.

5.1 Product Innovation

A technology or idea or quality level arrives into an industry at date zero. Over time, the basic technology or idea grows and matures, on account of the firms' own efforts in assimilating knowledge and also possibly from a flow of exogeneous information. One of the interesting issues in this context is the extent of maturation and the pace of diffusion of the original idea. The basic idea eventually generates a host of *differentiated products*, characterized by the maturation and improvement that each undergoes. The precise questions that one is interested in, in this context, are: i) how much maturation precedes the first introduction into the market, the *level of innovation*, and ii) how *diverse* are *subsequent innovations*? It should be noted that this approach is different from, and we believe a valuable complement to, the usual view of the innovation process. In the latter, the role of the first breakthrough is emphasized and all analysis relates to homogeneous products [(see, e.g. Reinganum (1989) for a survey)]. Greater detail on our approach and a more comprehensive model can be found in Dutta-Lach-Rustichini (1990). For completeness we report a detailed analysis in Appendix C.

There we show that this game has two stopping equilibria. The first one, which is the classic *preemption equilibrium*, involves a race by all firms in the industry to be the first to innovate. This results in rent equalization among firms, and comparatively little maturation of the basic idea. In the second and more interesting equilibrium, which we call *maturation equilibrium*, there is no preemptive race. The leader chooses optimally the time of entry, anticipating a later entry of the follower, after a maturation period. Entries are staggered, no rent is dissipated, but the leader has an equilibrium payoff lower than that of the follower.

5.2 Asset Sales

The price of an asset is appreciating. What is the best time to sell it? If there were a single trader, the answer to this question is simply found by solving a stopping time problem, that trades off discount costs to waiting against the benefits of a higher price. Suppose instead that there is more than one trader, for simplicity two. When i sells, it drives the price down for j. The asset price will rise again, but from this lower base. Does competition among traders lead to ***pre-emptive sales***, and a quick unloading of the asset before its price has sufficiently improved? We do not attempt anything like a comprehensive analysis of this problem, but merely offer a simple model in which to explore the question.

Again we refer the reader to Appendix B for a detailed discussion of the model and its analysis . Here we outline the main results.

We may distinguish two basic disinct possibilities, according to the loss in value of the asset after the first sale. In the first case, the asset depreciates heavily: indeed we consider an extreme case in which the price falls to zero. This results in preemptive "fire sales" in which the asset is offloaded early before its value can appreciate.

In the second case, the fall in the price of the asset is dependent on its current price. Again there may be preemptive equilibria. Now, however, if the fall in price is proportionally lower at higher valuations of the asset, then we show the existence of equilibrium outcomes in which traders are willing to forego short term gains. By cooperating initially they build up the asset value before the sale.

6. Commitment Horizons.

In the analysis so far we have allowed players to move at any time of their choosing. In several applications, however, there are both natural constraints (like the length of the day, or of the year, legal and institutional conditions) which make this complete freedom of choice impossible. Also, players may have available institutional arrangements that allow a committment to a specific course of action in the future.

What is the effect of allowing players in a stopping time game to possibly commit to their actions? For instance, suppose that a player could credibly commit to not move over the next τ length of time. More generally, suppose every τ periods, players can make "announcements" or "prepare for a move." At t if such preparation is *not* made, a player *cannot* move between t and $t+\tau$. Having made such a preparation a player can always forego a planned move at some $\hat{t} \in [t, t+\tau)$, but a move can only be made at $\hat{t}$. How does such commitment possibilities affect the set of equilibria? Conventional wisdom suggests that one should be able to enlarge cooperation possibilities. Somewhat surprisingly, we show that if long commitments are possible, equilibrium may fail to exist. The reason is not very different[8] from the non-existence problem that plagues the Hotelling spatial location problem. To keep matters simple, consider $\tau = \infty$, and a pure game of timing, i.e. $x(t) \equiv t$, $t \geqslant 0$.

Is there a Nash equilibrium if player i picks t_i, $i = 1, 2$, his move time, and commits to it in the sense that he can pass up his move but cannot move at another time? The choice of such a time is much like the choice of a location in a spatial model. If i's choice is "low" (small t_i), j prefers taking the "high" end, $t_j > t_i$ and vice-versa. But a middle ground is never a best response.

Let the payoffs l and f be further decomposed as follows. The flow payoff to a first entrant is a function of the time of entry. Denote this $\delta R(t_i)$. After the second entry, the two players earn returns which depend on t_i and t_j. For simplicity assume that these returns only depend on the difference in entry times $t_j - t_i$. Denote these returns $\delta r(t_i - t_j)$

and $\delta r(t_j - t_i)$ respectively, for players i and j. The optimal second entry is then given by

$$\max_{\theta \geq 0} e^{-\delta\theta} r(\theta)$$

Assume that the solution to this problem is unique, and denote it by θ^*. Then

$$f(t) = e^{-\delta\theta^*} r(\theta^*)$$
$$l(t) = (1 - e^{-\delta\theta})R(t) + e^{-\delta\theta^*} r(-\theta^*)$$

(see Appendix B for further details.) Let $r(t_i, t_j) = r(t_i - t_j)$. Strategies are choices of t_i, t_j and the payoffs, with $t_i < t_j$ are given by e.g.,

$$\pi_i = e^{-\delta t_i}\{(1 - e^{-\delta(t_j - t_i)})R(t_i) + e^{-\delta(t_j - ti)} r(t_i - t_j)\}$$
$$\pi_j = e^{-\delta t_j} r(t_j - t_i)$$

PROPOSITION 5. *There are concave increasing functions r and R, such that there is no Nash equilibrium to the commitment game, for some length of the commitment.*

PROOF: We shall briefly outline the proof.

Let s be the choice of entry time of player 1 say, and compute the best response for player 2. We consider two cases.

Case 1. If he decides to follow, the best entry time t, with $t \geqslant s$ is the solution of $\max_{t \geqslant s} r(t - s)e^{-\delta t}$, which is a constant, Δ say. Call $t_f(s)$ the best entry time, and $F(s)$ the value of being a follower.

Case 2. If he decides to be a leader, then he is solving $\max_{\{t: t \leqslant s\}} W(t, s)$, where $W(t, s) = R(t)(e^{-\delta t} - e^{-\delta s}) + r(t - s)e^{-\delta s}$.

To choose between the leader or the follower position, finally, the player chooses $\max\{F(s), L(s)\}$.

If we now set $R(x) = (\max\{x, 0\})^\gamma; r(x) = (\max\{x, 0\})^\alpha$ we find

1. A solution to the case 1 and case 2 problems exists.
2. There exists an s^* such that $F(s) \geqslant L(s)$ if and only if $s \in [0, s^*]$.

Fig. 1 here

The best response correspondence is upper semicontinuous but not convex valued.

One can prove (an unattractive exercise in calculus) that for some values of the parameters (δ, γ, α) no equilibrium exists.

Appendix A: Proof of Proposition 3. and 4.

To prove Proposition 3 we first establish an additional property satisfied by any MPE. Pick an $x_0 \in X$ and let $O_{x_0}(\sigma) = a$. Clearly, $a \in \{f \geqslant l\}$ since a is an equilibrium outcome.

LEMMA A.1. *Suppose (A1)–(A3) hold. If* $\sigma_1(a) = 1$ *for* $a \in \{f > l\}$, *then* $\sigma_2(x) = 0$ *for every* x *in the connected component of* $\{f > l\}$ *which contains* a.

PROOF: Define

$$b = \begin{cases} \inf\{x : f(x) < l(x), \quad x > a\} \\ +\infty, \text{ if this set is empty.} \end{cases}$$

Let $M_i \equiv \{x \in [a,b) : \sigma_i(x) = 1\}$; since $a \in M_1$, we claim $M_2 = \emptyset$. Letting $d(M_1, M_2)$ denote the distance between M_1 and M_2, with the convention $d(M_1, M_2) = +\infty$ if $M_2 = \emptyset$, we claim first that $+\infty > d(M_1, M_2) > 0$ is impossible. Supposing otherwise we conclude in fact that there exist two points C_i, $i = 1, 2$ such that $\sigma_i(C_i) = 1$, $\sigma_i(x) = 0$ for any x in the interval with extreme points C_1 and C_2. We may assume w.l.o.g. that $C_1 < C_2$. Then

1. $l(C_1) \geqslant E_{C_1} e^{-\delta T_{C_2}} f(C_2)$ (because $\sigma_1(C_1) = 1$);
2. $l(C_1) < E_{C_1} e^{-\delta T_{C_2}} l(C_2)$ (because $\sigma_2(x) = 0, \quad x \in [C_1, C_2]$).

But 1 and 2 give a contradiction.

On the other hand, $d(M_1, M_2) \neq 0$. Supposing otherwise, we can find $b_1 \in M_2$ and, for any $\delta > 0$, a pair (C_1, C_2) such that $C_i \in M_i$, $C_i \in [a, b_1]$, $i = 1, 2$, and (again w.l.o.g.) $\delta > C_2 - C_1 > 0$. But this contradicts, for δ small enough, that $\sigma_1(C_1) = 1$ is a best response, as it follows immediately from Proposition 1. and the fact that $f(x) > l(x)$, $x \in [a, b_1]$. ∎

PROOF OF PROPOSITION 3: Let σ be a MPE and denote $O_x(\sigma) = a, a \neq x$. It must be that $f(a) \geqslant l(a)$. Else, by continuity of the payoffs, there is a state z in some left neighborhood of a, such that

$$l(z) > E_z e^{-\delta T}\{\alpha l(a) + (1-\alpha) f(a)\}$$

where $T = \inf\{t : x_t = a; x_0 = x\}$ In this case, some player can deviate and enter by unilaterally moving at z. If player i is the entrant at a,then a is an optimal move in anticipation of j's next "state of entry," say $b \geqslant a$. Suppose $l(b) < f(b)$ (and hence $b > a$). Then, let us denote (C, D) to be the interval such that $b \in (C, D)$, and $l(C) = f(C)$. Consider now a pair of stopping* strategies such that, $\sigma_i(C) = \sigma_j(C) = 1$, there is cooperation on all potential cooperation regions C_k, between a and b, and no cooperation on any cooperation region beyond b.

For states between x and a, the prospect is exactly as in σ, so the previous behavior (which was in accordance with stopping* strategies) is still an equilibrium. Moreover,

$$l(a) \geqslant E e^{-\delta T} f(b) > E e^{-\delta T} l(b) \geqslant E e^{\delta \tau} l(C)$$

where

$$T = \inf\{t \geqslant 0 : x_t = b; x_0 = a\}$$
$$\tau = \inf\{t \geqslant 0 : x_t = C; x_0 = a\}$$

The last inequality follows from Lemma A.1 and the fact that player j given a choice of any entry date in (C, b), actually picked b.

The argument for $l(b) = f(b)$ is identical.

We now turn to the proof of Proposition 4.

PROOF: Suppose $O_{x_0}(\sigma) = \{a\}$. Then, it follows that along $h_t(x_0)$, $x_t \neq a$, $t < T$. Else, a player who moves at T, could move before T at the same state a, and given discounting, make strictly more in payoffs. In particular, then no state repeats between x_0 and a. Otherwise, the process would merely cycle and never get to a. So x_t converges either from above or from below from x_0 to a.

As in Proposition 3, consider stopping strategies which do not stop over potential cooperation regions in (x_0, a) or (a, x_0) and stop over all other cooperation regions C_k. They constitute an equilibrium. ∎

Appendix B: Continuity of the payoff functions

In this Appendix we explore some "primitive" formulations which are covered by our stopping time game framework. In particular we have in mind cases in which the payoff functions l and f are not given as primitive elements of the game, but arise as solution of the decision problem faced by the follower.

Let the game environment evolve, as before, according to the stochastic process $\{X_t : t \geqslant 0\}$. Suppose now that player i moves at a state x. Consequent to this action, i gets a flow payoff per period of $\delta R(x)$, till such time as j moves. Let the state at j's move by y. Then, in the subsequent stage, i gets a payoff of $\delta r(x, y)$ per period, and j gets $\delta r(y, x)$. We shall always assume that the functions R and r are continuous. These last payoffs may be interpreted in some contexts as the outcomes of one shot games in which i and j have "states" x and y respectively. We refer to Section 5. for one concrete application.

Consider j's stopping problem,

$$\max_{T \in M} E e^{-\delta T} r(y_T, x)$$

The value of this problem can be denoted by $f(x)$, and the total value that the player i will receive as a consequence is then $l(x)$. So in either an "exit" or an "entry" interpretation, stopping time games covers a variety of economic problems.

Now we proceed to demonstrate sufficient conditions on r and R which guarantee the continuity of l and f. In the following we consider special cases.

I I.1 $r(x, y) = s(x - y)$;
I.2 $X = \mathbf{R}^m$;
I.3 The stochastic process is spatially homogeneous.

Note the I. 1,2,3 above imply that

a) $f(x) = f(0) = \sup_{T_f} e^{-\delta T_f} s(x_{T_f})$, so T_f is a stopping time independent of x, and f is a constant;

b) $l(x) = E_x\{R(x)(1 - e^{-\delta T_f}) + s(x - X_{T_f})e^{-\delta; T_f}\} = R(x)C_1 + s(x - X_{T_f})C_2$ where C_1 and C_2 are constants, and $x_{T_f} - x = \Delta \in \mathbf{R}^m$ a.s.

II II.1 $X = \mathbf{R}$
II.2 The stochastic process is increasing, with the drift coefficient of the infinitesimal generator given by a continuously differentiable function M, and with a derivative M' satisfying $0 \leqslant M' < \delta$.
II.3 The function $x \to r(x, y)$ is strictly concave for every y.

III III.1 $X = \mathbf{R}$, or $X = \mathbf{R}_+$, and a reflecting barrier at the origin;
III.2 the stochastic process is given as solution of the stochastic differential equation $dx = Mdt + \sigma dW$, with $M \in \mathbf{R}$, W a standard Brownian motion;
III.3 the function $x \to r(x, y)$ is concave for every y; and $r(x, y) = 0$ if $x < y$.

For both cases II and III above, the proof of the continity of l and f consists of observing that the optimal stopping time of the follower is determined by

$$T_f(x) = \inf\{t > 0 : x_t = A(x)\}$$

and the function $x \rightarrow A(x)$ is continuous because it is determined as the tangency points of two concave functions: $r(x,y)$ and the solution u of the equation $Lu = 0$, where L is the infinitesmal generator associated with the stochastic process.

It follows that

$$f(x) = r(A(x), x)E_x e^{-\delta T_f(x)}$$

and

$$l(x) = R(x)E_x(1 - e^{-\delta T_f(x)}) + r(A(x), x)E_x e^{-T_f(x)}$$

are continuous.

Collecting all of the above, it follows that

PROPOSITION B.1. *Under the conditions I, or II or III, the payoffs to leader and follower in the associated stopping time game are continuous functions of the state.*

Appendix C: Product Innovation and Asset Sales

C.1: The Product Innovation Model

An initial idea or technology, with a payoff relevant real valued variable x_0 is available at date 0. Over time this technology evolves. Suppose that the technology process $\{X_t; t \geqslant 0\}$ is increasing.[9] For expositional ease, let us also take it to be deterministic and single dimensional,[10] and hence by renormalization, $X(t) \equiv t$, $t \geqslant 0$. At time t either firm can introduce a product, incorporating attribute t, and consequently earn monopoly returns per period of $\delta R(t)$, till the other firm enters the market. After the moment of the two innovations, at t_i and t_j say, i and j respectively make flow profits of $\delta r(t_i - t_j)$ and $\delta r(t_j - t_i)$.[11] This specification means, of course, that returns only depend on relative quality or vintage. If firms attempt to move simultaneously, each actually moves with probability of $\frac{1}{2}$.

Stopping Equilibrium

Suppose firm i moves at t. Firm j's stopping problem is

$$\max_{\theta \geqslant 0} e^{-\delta\theta} r(\theta). \tag{1}$$

Suppose the maximand in (1) has a unique solution $\theta^* > 0$. Then, with l and f denoting the payoff to the leader and the follower as usual,

$$\begin{aligned} f(t) &= e^{-\delta\theta^*} r(\theta^*), \quad t \geqslant 0 \\ l(t) &= (1 - e^{-\delta\theta^*}) R(t) + e^{-\delta\theta^*} r(-\theta^*). \end{aligned} \tag{2}$$

If the monopoly returns $\delta R(t)$ improve with the monopolist's quality t, then the first mover's payoffs increase with the state. The follower's payoffs are clearly constant, say ϕ.

Fig.C.1 here

Let us suppose that $l(0) < \phi < l(\infty)$,[12] and define t_m as $l(t_m) = \phi$. Clearly, both firms move beyond t_m, if neither has moved till then (alternatively, $C_k = (t_m, \infty)$, $\#K = 1$). Note,

$$\begin{aligned} V_l(t) &= \max_{\theta \in [0, t_m - t]} e^{-\delta\theta} l(t + \theta), \quad t \in [0, t_m) \\ &= e^{\delta t} \max_{t+\theta} e^{-\delta(t+\theta)} l(t + \theta). \end{aligned} \tag{3}$$

Suppose finally that $e^{-\delta x} l(x)$ is quasi-concave (and hence single peaked) on $[0, t_m]$. Denote the position of the peak to be t^*. We have two cases:

Case 1. $t^* < t_m$. Then, $A = \{t \in [t^*, t_m]\}$

$$\begin{aligned} V_l(t) &= \begin{cases} e^{\delta t}[e^{-\delta t^*} l(t^*)] & t \leqslant t^* \\ l(t) & t > t^* \end{cases} \\ V_f(t) &= \begin{cases} e^{\delta t}[e^{-\delta t^*} f(t^*)] & t \leqslant t^* \\ f(t) & t > t^*. \end{cases} \end{aligned} \tag{4}$$

In this case Theorem 1 and Proposition 4 yield

PROPOSITION C.1. *i) If $t^* < t_m$, the unique*[13] *subgame perfect equilibrium in pure strategies is the stopping equilibrium,*

$$\sigma_i(t) = \begin{cases} 1 & \text{on } [t^*, \infty) \\ 0 & \text{else} \end{cases}, i = 1,2, i \neq j$$
$$\sigma_j(x) = \begin{cases} 1 & \text{on } [t_m, \infty) \\ 0 & \text{else.} \end{cases}$$

In particular, the outcome in equilibrium (from initial state 0) is i moves at t^*, j moves at $t^* + \theta^*$. Further, payoffs are respectively $V_l(t^*)$ and $V_f(t^*)$, with $V_l(t^*) < V_f(t^*)$.

Elsewhere, (Dutta-Lach-Rustichini (1990)) we have called this a *maturation equilibrium* to emphasize the fact that a second mover waits optimally to innovate, makes higher payoffs by doing so and consequently does not dissipate the monopoly rents of the first mover.

On the other hand,

Case 2. $t^* = t_m$. Then, $A = \{t_m\}$ and (4) holds with $t^* = t_m$.

PROPOSITION C.2. *ii) If $t^* = t_m$, the unique subgame perfect equilibrium in pure strategies is the stopping equilibrium*

$$\sigma_i(t) = \sigma_j(t) = \begin{cases} 1 & \text{on } [t_m, \infty) \\ 0 & \text{else.} \end{cases}$$

The outcome is t_m, $t_m + \theta^*$. Payoffs are $l(t_m) = \phi$, for each firm.

A fuller discussion of the issues pertaining to product innovation may be found in Dutta-Lach-Rustichini (1990).

C.2: Asset Sales

The price of an asset appreciates over time. For example we have in mind a situation in which the market demand for the asset increases over time. Any sale satisfies some surrent demand and lowers the market price. Subsequently the price rises again due to long term demand increases.

Suppose that the price process is increasing, single dimensional, and purely for expositional ease, deterministic. Let price at time t be t itself.[14] If trader i sells at time t, his (present discounted) value of returns is $R(t)$ [15]. Let $m(t)$ denote the market price after i's sale at time t, $m(t) < t$. Trader j can then decide how much longer he waits before he sells. As before, simultaneous attempts at selling result in each trader selling with a probability of $\frac{1}{2}$.

The Stopping Equilibria of the Asset Sales Game

Trader j's stopping problem, after i's sale at t, is

$$\max_{\theta \geqslant 0} e^{-\delta\theta} R[m(t) + \theta] = e^{\delta m(t)} \max_{m(t)+\theta} e^{-\delta[m(t)+\theta]} R[m(t) + \theta]$$

Suppose $e^{-\delta x}R(x)$ is quasi-concave (and hence single peaked) on $\mathbf{R}_+$. Let the peak be at $t^* > 0$. Then,

$$l(t) = R(t)$$
$$f(t) = \begin{cases} e^{\delta m(t)}[e^{-\delta t^*}R(t^*)], & m(t) \leqslant t^* \\ R(m(t)) & m(t) > t^*. \end{cases} \tag{5}$$

Fig. C.2 here

As is clear from (5), the features of the payoff f are completely determined by the market price adjustment function m. We do not model that explicitly here but merely point to a couple of interesting possiblities.

Case 1. $m(t) = 0$, for all t.[16] Then, $f_1(t) \equiv e^{-\delta t^*}R(t^*) = \phi$, $t \geqslant 0$. Let t_m define, $l(t_m) = \phi$. It follows then that $t_m < t^*$ and we are in the second case of the Product Innovation Model.

PROPOSITION C.3. *i) If $m(t) = 0$, $t \geqslant 0$, the unique subgame perfect equilibrium in pure strategies is the "pre-emptive sales" stopping equilibrium*

$$\sigma_i(t) = \sigma_j(t) = \begin{cases} 1 & \text{on } (t_m, \infty), \sigma_i(t_m) = 1 \\ 0 & \text{else}, \sigma_i(t_m) = 0. \end{cases}$$

The outcome is sales at $t_m (< t^*)$ and $t_m + t^*$. Payoffs are $l(t_m) = \phi$ for each firm.

Hence, a single trader would wait till the price rose to t^* and earn a present discounted value *at period* 0 of ϕ. A competitive trader sells at $t_m < t^*$, and makes returns, evaluated at period 0, of $e^{-\delta t_m}\phi < \phi$.

On the other hand, suppose the market adjustment process is subject to increasing returns for intermediate prices and hence leads to f_2.

Fig.C.2 here

Then, the cooperation regions are $C_1 = [t_1, t_2]$ and $C_2 = [t_3, \infty)$. Clearly, there cannot be cooperation over C_2. So, the only question is, could the traders desist from selling the first time first mover advantage in payoffs appears, i.e. over C_1. From Theorem 1, this can only be if

$$V_l(t) = e^{\delta t} \max_{(t+\theta)\in[t_2,t_3]} \{e^{-\delta(t+\theta)}R(t+\theta)\} > R(t), \quad t \in C_1.$$

By the quasi-concavity of $e^{-\delta x}R(x)$ and the fact that $t_i < t^*$, $i = 1, 2, 3$, we have

$$V_l(t) = e^{\delta t}[e^{-\delta t_3}R(t_3)] > R(t).$$

Hence, cooperation over C_1 is supportable.

PROPOSITION C.3. *ii) In Case 2, there are two subgame perfect equilibria. In the "good" equilibrium,*

$$\sigma_i(t) = \sigma_j(t) = \begin{cases} 1 & \text{on } [t_3, \infty) \\ 0 & \text{else} \end{cases}$$

The outcome is sales at t_3 *and* $t_3 + (t^* - m(t_3))$. *Payoffs are equal and evaluated at* $t = 0$, $e^{-\delta t_3} R(t_3) = e^{-\delta t_3}[e^{\delta m(t_3)} \phi]$.

In the "bad" equilibriu m,

$$\sigma_i(t) = \sigma_j(t) = \begin{cases} 1 & \text{on } [t_1, t_2] \cup [t_3, \infty) \\ 0 & \text{else} \end{cases}$$

The outcome is sales at t_1 *and* $t_1 + (t^* - m(t_1))$. *Payoffs are equal but even lower at* $e^{-\delta t} R(t_1) = e^{-\delta t_1}[e^{\delta m(t_1)} \phi]$.

There are always pre-emptive sales. Some pre-emptive sales are worse than others!

Footnotes

1. Later we show how we can accomodate different soecifications of payoffs.

2. Bensoussan-Friedman's formulation also allows for flow payoffs before the process is stopped, that depend on the "state" of the system. This generalization is ignored here.

3. The symmetry assumption here is important. The extension to the non symmetric case seems non trivial.

4. Note that in some games of timing, backward induction leads to a *unique* equilibrium. However, the arguments invariably involve iterated elimination of (weakly) dominated strategies, starting from the "end." Hence, they require that there be an "end" to a game of potentially infinite duration, i.e. there be a time, beyond which a player will "always stop" or "never stop," regardless of the other. This itself implies strong restrictions on l, f which in particular applications may be natural but are unreasonable to impose in general. Further, with uncertainty in the game evolution, conditions ensuring the existence of outcomes to iterated elimination, are even more restrictive.

5. Of course, the process could still be stochastic.

6. Notice that, with $T_a = \inf\{t > 0 : X_t > a\}$, a strategy defined by

$$\sigma(s, w) = \{1 \text{ if } s \in T_a + \cup_n[1/2n, 1/2n + 1); \quad 0 \quad otherwise\}$$

   is progressively measurable because it is left continuous.

   We want to be able to say the player stops at a; in order to exclude situations like the case described above, we impose the following regularity conditions on strategies

$$\limsup_{t \downarrow t_0} \sigma(t, w) \leqslant \sigma(t_0, w) \text{ a.s., for every } t_0.$$

7. Note that the public randomization here involves, in some sense, an even greater degree of sophistication than in the case of repeated games. Players in our game act but once. Hence, correlation is enforced to achieve "abstinence from action," rather than to bring about particular actions as with repeated games.

8. The difference is that non-existence arises in the choice of "location" rather than at the Bertrand price competition phase.

9. The reader could imagine alternately that a firm improves its technology by keeping its laboratory open, and only by doing so. Alternatively, technology improvements happen in the public domain (e.g. at universities and government research laboratories) and are accessible by all firms in an industry. In the former interpretation, there is some periodic cost of doing research which we normalize to zero. That does not affect the results.

10. That the two applications we discuss both involve deterministic, single dimensional, increasing processes, should not of course suggest that the general theory is inapplicable otherwise. Our inquiry was motivated precisely to address applications outside this pure game of timing framework. For expositional clarity we revert to the timing framework. None of the arguments, of course, hinge on these simplifications.

11. Think of these returns as profits under Cournot or Bertrand competition. Note also that a fixed cost of introducing a product could be incorporated without any difficulty.

12. If the first inequality is violated, we have a pure pre-emption game, and if the second, a pure war of attrition. The equilibria in these two cases are obvious: respectively, both move immediately at $t = 0$ and i moves at $\arg\max_{t \geqslant 0} e^{-\delta t} l(t)$, j follows after θ^*.

13. Modulo permutation of the identities of the identical symmetric players.

14. That prices are unbounded does not drive any of the analysis as the reader is invited to check. An alternative formulation, with $X(t) = P(t)$, for a bounded P, would suffice just as well.

15. For instance, $R(t)$ is the lifetime return to t dollars in the best future investment. Note, incidentally that it does not much matter for our purposes as to whether i gets to sell all of his asset at price t or at some price below t. All that is needed is that the *average price* i sells at is strictly *more* than the price *after* his sale, i.e. $m(t)$.

16. As can be checked exactly, the same arguments work when prices adjust to some fraction of the previously prevailing price, i.e. $m(t) = \alpha t$, $\alpha \in (0, 1)$.

REFERENCES

1. Bensoussan, A. and Friedman, A., *Non-Zero Sum Stochastic Differential Games with Stopping Times and Free Boundary Problems*, Transactions of the A.M.S. **231** (1977), 275–327.

2. Dasgupta, P. and Stiglitz, I., *Uncertainty, Industrial Structure, and the Speed of R& D*, Bell Journal of Economics **11** (1980), 1–28.

3. D'Aspremont, C.; Gabscewicz, J. and Thisse, J. F., *On Hotelling's Stability in Competition*, Econometrica **17** (1979), 1145–1151.

4. Dutta, P.K; Lach, S.; Rustichini, A., *Innovation and Product Differentiation*, Northwestern University, Working Paper.

5. Dutta, P.K. and Rustichini, A., *(s,S) Equilibria in Stochastic Games with an Application to Product Innovation*, preprint, University of Rochester.

2. Fakeev, A.G., *Optimal Stopping Rules for Stochastic Processes with Continuous Parameter*, Theory of Probability and Its Applications **18** (1979), 304–311.

7. Fudenberg, D. and Tirole, J., *Preemption and Rent Equalization in the Adoption of New Technology*, Review of Economic Studies **52** (1985), 383–401.

8. ________________, *A Theory of Exit in Duopoly*, Econometrica **54** (1986), 943–960.

9. ________________, *A Course in Game Theory*, MIT Press, forthcoming (1989).

10. Ghemawat, P. and Nalebluff, B., *Exit*, Rand Journal of Economics **16** (1985), 184–194.

11. Maskin, and Tirole, J., *A Theory of Dynamic Oligopoly, I: Overview and Quantity Competition with Large Fixed Costs*, Econometrica **56** (1988).

12. Reinganum, J.F., *The Timing of Innovation: Research, Development, and Diffusion*, Handbook of Industrial Organization North Holland (1989).

13. Parthasarathy, T., *Discounted Positive, Non Cooperative Stochastic Games*, International Journal of Game Theory (1973).

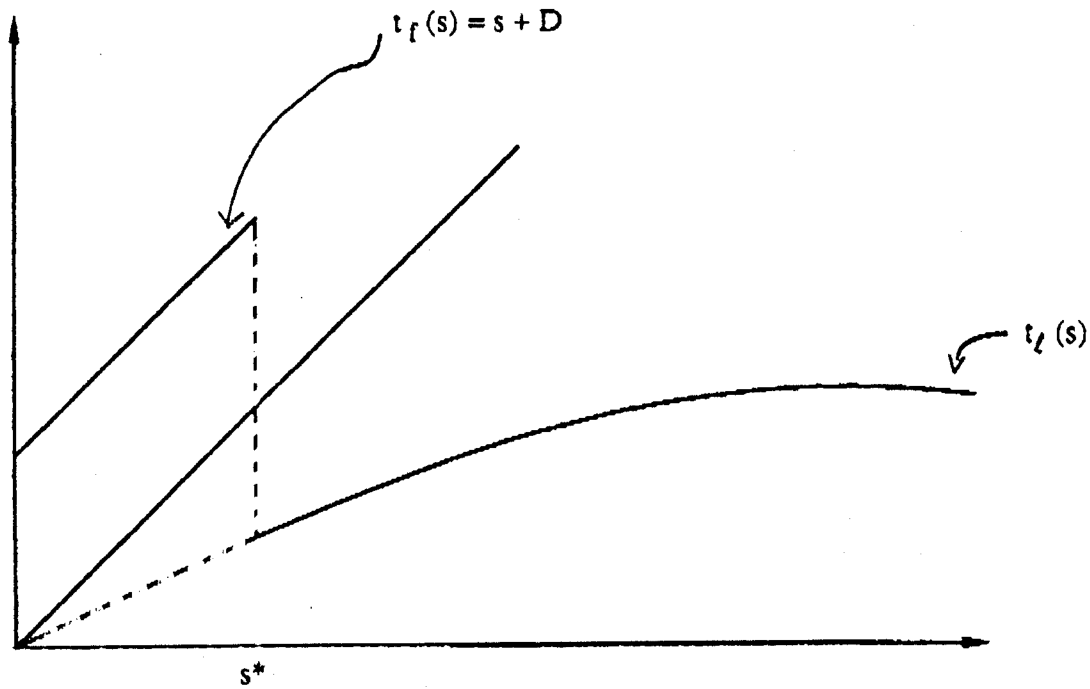

Figure 1

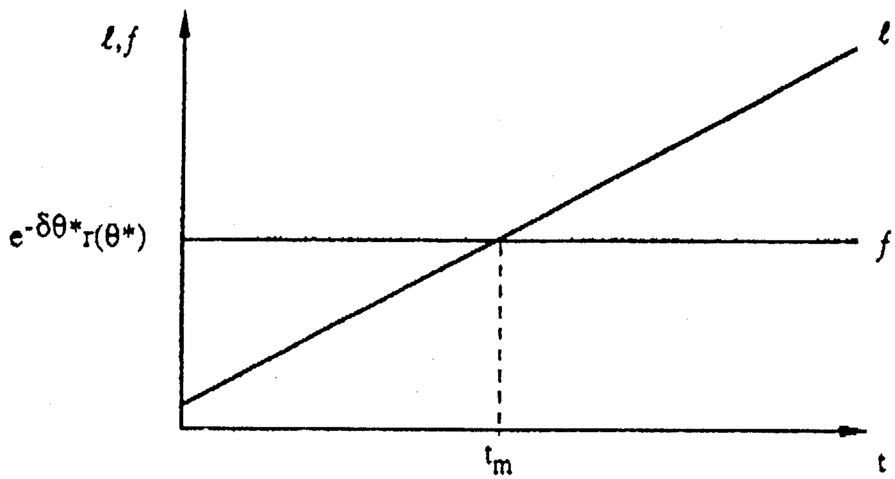

Figure C.1

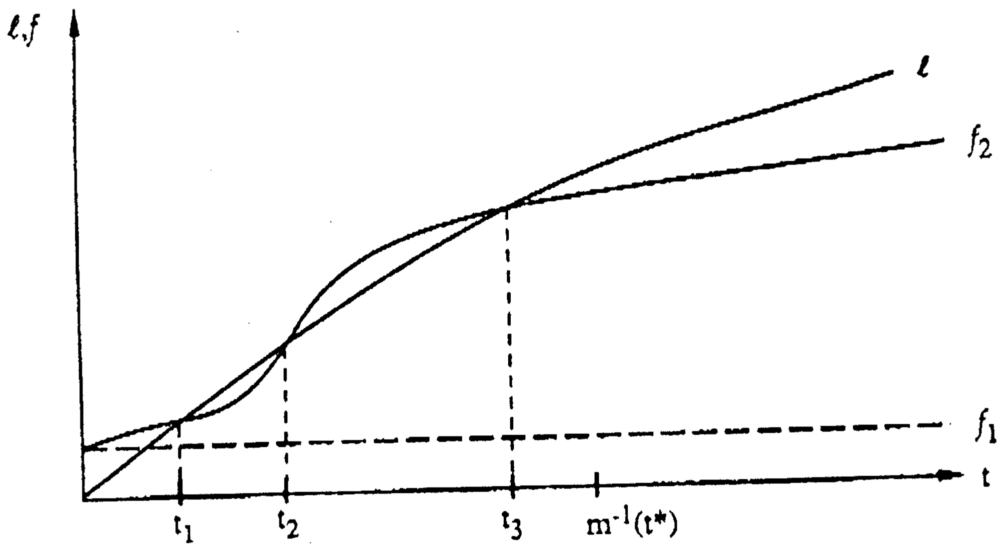

Figure C.2