

UC Berkeley

Other Recent Work

Title

Extensive Form Games in Continuous Time: Pure Strategies

Permalink

<https://escholarship.org/uc/item/03x115sh>

Authors

Simon, Leo K.
Stinchcombe, Maxwell B.

Publication Date

1987-07-16

Peer reviewed

UNIVERSITY OF CALIFORNIA, BERKELEY

Department of Economics

Berkeley, California 94720

Working Paper 8746

EXTENSIVE FORM GAMES
IN CONTINUOUS TIME:
PURE STRATEGIES


Leo K. Simon and Maxwell B. Stinchcombe

July 16, 1987

Key words: Game theory, continuous-time games, subgame perfect equilibrium, coordination games, patent races.

Abstract

We propose a new framework for modelling games in continuous time. Continuous time is viewed as "discrete-time, but with a grid that is infinitely fine." We apply the model to several standard problems from the industrial organization literature. We then address questions such as: "is discrete time with a very fine grid a good proxy for continuous time?" and "does every continuous-time equilibrium in our model have a discrete-time analog?"



ACKNOWLEDGEMENT

The authors have benefited greatly from discussions with numerous colleagues, but especially Robert Anderson, Karl Iorio, Ariel Rubenstein, Hugo Sonnenschein and Bill Zame. We are grateful to three referees who provided extremely helpful comments on an earlier draft.

Leo K. Simon

Department of Economics
University of California, Berkeley

Maxwell B. Stinchcombe

Department of Economics
University of California, San Diego



UNIVERSITY OF CALIFORNIA, BERKELEY

Department of Economics

Berkeley, California 94720

Working Paper 8746

EXTENSIVE FORM GAMES
IN CONTINUOUS TIME:
PURE STRATEGIES

Leo K. Simon and Maxwell B. Stinchcombe

July 16, 1987

Key words: Game theory, continuous-time games, subgame perfect equilibrium, coordination games, patent races.

Abstract

We propose a new framework for modelling games in continuous time. Continuous time is viewed as "discrete-time, but with a grid that is infinitely fine." We apply the model to several standard problems from the industrial organization literature. We then address questions such as: "is discrete time with a very fine grid a good proxy for continuous time?" and "does every continuous-time equilibrium in our model have a discrete-time analog?"

ACKNOWLEDGEMENT


The authors have benefited greatly from discussions with numerous colleagues, but especially Robert Anderson, Karl Iorio, Ariel Rubenstein, Hugo Sonnenschein and Bill Zame. We are grateful to three referees who provided extremely helpful comments on an earlier draft.

Leo K. Simon

Department of Economics
University of California, Berkeley

Maxwell B. Stinchcombe

Department of Economics
University of California, San Diego



This is the first of a series of papers proposing a new framework for modelling complete information games in continuous time.¹ Our objective is to construct a theory that conforms as closely as possible to discrete-time game theory.² Indeed, there is a precise sense in which our continuous-time formulation can be interpreted as "a discrete time model, but with a grid that is *infinitely* fine."

A fundamental problem for game theory in continuous time is that there is no natural notion of a "last time before t ."³ Therefore induction cannot be applied. Since induction is fundamental to discrete-time game theory, the most primitive concepts of this theory have no direct continuous-time analogs. These concepts--decision trees, strategies and outcomes--must therefore be defined from first principles.

To see the problem, consider the following strategy for a one-person game-form in discrete time:⁴

"I play '*left*' at time zero; at every subsequent time, play the action I played last period." (0.1)

We will attempt to play this strategy in continuous time, say on the interval $[0, 1)$. The first question that arises is: does this strategy have any meaning in continuous time, when there is no longer any notion of "last period?" Possibly, we can resolve this problem by replacing "last period" with "in the immediate past." But then there is a second problem: what "outcome"--i.e., path through the game tree--do we associate to this strategy? This outcome must specify a choice for the player at every time in $[0, 1)$. However, at every positive time, t , there have been times strictly between zero and t . The action the agent should choose at t depends on the choices he made at these earlier times, so these choices must be determined before we can determine his action at t . But the same thing is true for every t ! Thus, at *every* positive time, we have insufficient information to determine what the agent should be doing. We conclude from this that the very language of discrete time game theory ceases to be meaningful when extended to continuous time.

¹ In the present paper, agents are restricted to choose pure strategies. In subsequent papers, (Simon [21], Simon [22], Simon-Stinchcombe [24]) we extend the model to allow for randomization.

² Stinchcombe [25] proposes an alternative approach to modelling continuous-time games, using techniques from the theory of stochastic processes. His model admits a very rich (but noncomparable) class of strategies and allows for incomplete information. On the other hand, the connection between his formulation and conventional discrete-time game theory is more tenuous.

³ Or a "first time after t ." This is because neither "greater than" nor "less than" well-order the continuum. See Anderson [1] for an related discussion of this issue.

⁴ Anderson [1] contains a related example.

Questions like these must be addressed before we can even begin to formulate a continuous time model. Specifically, we must identify a class of strategies with the property that when any member of the class is combined with any other (or several others), some outcome is uniquely identified. Moreover, there must be some sense in which the rule that defines outcomes is an "appropriate" one. This is a nontrivial task. Even for so simple a strategy as (0.1) above, there are a vast number of outcomes to choose from. Consider, for example, the outcome: "play 'left' on the interval $[0, t]$, and 'right' thereafter," where t is any time in the interval $[0, 1)$. This outcome is "compatible" with strategy (0.1), in the strongest possible sense. At every $s \leq t$, the player has played 'left' in the immediate past, so that according to his strategy, he should play 'left' at s . The outcome has him doing so. At every $s > t$, there have been times in the recent past at which he has played 'right', so that according to his strategy, he should play 'right' at s . The outcome has him doing so. The problem here is that there is no *first* s beyond t at which the agent jumps from 'left' to 'right', and therefore no point at which the strategy and outcome are incompatible.

In spite of these and related problems, we believe that the development of a continuous-time formulation will repay the effort, for several reasons. First, in many actual economic situations, agents can make decisions virtually whenever they wish. A model that restricts agents to, say, one move a day may therefore yield misleading predictions, especially if agents stand to gain a great deal by moving just before their allotted times. Second, a variable that frequently proves important in dynamic models is the length of reaction or information lags. Economists are often particularly interested in the polar case of negligible lags. In discrete-time, "negligible" can only mean "the length of one period." In continuous-time, agents can react instantaneously, so that lags may truly be negligible. It turns out that the difference between one period lags and no lags at all can be dramatic: there are problems for which we can obtain unique and intuitive solutions that contrast sharply with the predictions of discrete time models. (See sections II and V for examples.) A third reason for studying continuous time is that in certain contexts, it is much more convenient. Economists usually model quantities and prices as continuous variables, because these are easier to work with than discrete ones. In particular, differential calculus techniques can be used to make marginal calculations. These reasons may be equally valid when time is the economic variable being analyzed. (See section V for an example.)

Aside from the particular virtues of continuous-time, it is useful to have an alternative perspective from which to view dynamic problems. The predictions of continuous- and discrete-time models can then be checked against each other for robustness. Our continuous-time alternative will be all the more useful if it is highly compatible with the traditional, discrete-time model, because then sharper comparisons can be made. Our model has been designed with these considerations in mind. Indeed, there is a sense in which our model and the traditional one can be viewed as special cases of the same overall structure. Accordingly, we can ask questions such as, "Is discrete-time with a very fine grid a good proxy for continuous time?" and "Does every equilibrium in our model have a discrete-time analog?" Some partial answers to these questions are provided in section VI.

The paper is organized as follows. We begin with three informal and heuristic sections. Section I introduces the model, II illustrates it and III explains various conceptual issues that arise in modelling continuous time. Section IV sets out the formal model. (We cannot, of course, prove existence until randomization is introduced.⁵) Section V considers two examples. In particular, we obtain a striking uniqueness result for a class of repeated games. In section VI, we study the relationship between our model and the conventional discrete-time one. Proofs are gathered together in the Appendix.

I. Introduction to our continuous-time model.

There is a finite set of agents. Each agent can choose from a finite set of actions. There is a set of times at which agents can move. In a discrete-time game, this set would be a finite subset of the unit interval. In continuous-time, it is the interval $[0, 1)$.⁶ Our agents can change their actions at any point in this interval. A decision node is a point in time, paired with a complete description of the moves that agents have made in the past. A strategy is a function that assigns an action to each decision node. A strategy profile is a list of strategies, one for each agent. An outcome is a complete record of the decisions made throughout the game, that is, a path through the game tree. An outcome function is a mapping from decision nodes and strategy profiles to outcomes. To define a continuous-time game-form, we must identify (a) a set of

⁵ See Simon [21] for an existence theorem for a simple class of games.

⁶ With virtually no modifications, the time interval could be changed to $[0, \infty)$. We could also work with the closed unit interval, but chose not to because the last point in the game must be treated specially.

continuous-time outcomes, (b) a family of strategy profiles and (c) a consistent way to associate outcomes to decision nodes and profiles.

Unless some limitations are imposed on agents' behavior, the program outlined above is daunting. As we have seen, things can go wrong with even the simplest conceivable strategies. These problems become much more severe when strategies are more complex. Accordingly, we will simplify matters by restricting the options available to our agents. In particular, we will require that agents can make only finitely many changes in the actions they play.⁷ This simultaneously restricts the set of admissible strategies *and* the set of outcomes that we need to consider. It does, of course, also limit the applicability of our model. However, there remains a large variety of problems for which our simplified model is applicable. (This discussion is taken up in section III).

We now introduce our outcome function. First observe that the decision nodes for each discrete-time game are contained by the set of continuous-time decision nodes. Therefore, a continuous-time strategy profile can be restricted to the decision nodes for a given discrete-time game, to define a strategy profile for this game, and hence an outcome starting from each node. Thus, a continuous-time strategy can be interpreted as a set of instructions about how to play the game on every conceivable discrete-time grid. There is, now, a natural candidate procedure for defining outcomes: (a) fix a continuous time decision node and strategy profile; (b) restrict the profile to an arbitrary, increasingly fine sequence of discrete-time grids and play the restricted profiles from the given node; (c) define the continuous time outcome to be the limit of the discrete-time outcomes.

For example, the procedure works nicely for example (0.1) above (p. 1). On any discrete time-grid, this strategy generates by induction the unique outcome: "play '*left*' at every grid point." The limit of these outcomes as the grids become finer is the uniquely defined outcome: "play '*left*' at every time in $[0, 1)$." In this example, therefore, the approach just described selects the only outcome that is intuitively sensible. In general, however, the procedure may break down. For some strategies, the limit of discrete time outcomes may not exist. For others, it may exist but depend on the particular sequence of grids. Finally, for some profiles, a unique limit may exist, but it may be "incompatible" with the original strategies in a fundamental

⁷ Stinchcombe [25] allows for countably many changes in action and infinite action spaces. This enables him to model time paths that depend smoothly on time.

way. Accordingly, for our procedure to be coherent, we need to identify a class of strategies with the following property: starting from any decision node, the discrete time outcomes defined by playing the restrictions of a given profile on an arbitrary sequence of grids must converge to a unique limit that is independent of the particular sequence of grids and is compatible with the original strategies. Once we have identified this family of strategies, and by implication, the universe of possible outcomes, we will have completed the specification of our continuous-time game-form.

To define a continuous-time *game*, we need to add a valuation function. In all of our applications, this function will be included in the "intuitive" specification of the problem. The natural solution concept for our model is Selten's [20] subgame perfection (SGP). In an SGP equilibrium, each agent's strategy must be a best response to other agents' strategies, starting from every node in the game. In some applications this equilibrium notion will be too weak and we will invoke a much stronger solution concept, involving iterated elimination of weakly dominated strategies.⁸ This concept uses the special properties of continuous-time to obtain sharp predictions that simply cannot be obtained from discrete-time models.

II. A Simple Example.

In this section we study two variants of a simple preemption game.⁹ The first illustrates various aspects of our model. The second is a "negative" example in the sense that it has no pure-strategy equilibrium. It does, however, highlight an important difference between discrete- and continuous-time.

We now describe the first variant. There are two firms, who may enter a market at some time in the interval $[0, 1)$. Earlier entry is more costly: if a firm enters at t , it incurs a cost of $(1 - t)^2$. Once a firm has entered the market, it cannot leave again. So long as only one firm is active in the industry, it earns monopoly flow profits $\$ \pi^m$ per unit time. In the first variant of the problem, we set $\pi^m = 2$. If both firms are active, industry flow profits are $\$3$ per unit time. The earlier entrant acquires loyal customers over time and hence

⁸ This idea has a long history that goes back to Luce and Raiffa [15]. Moulin formalized the concept as "dominance solvability" [16]. Our concept is closely related, but not identical to Moulin's. We proceed in the following way: we first eliminate weakly dominated strategies from each agent's strategy set; once agents' strategy sets have been shrunk, new domination relations appear, so that further strategies can be eliminated. And so on. Alternatively, there is a natural way to extend Myerson's "properness" [17] to games such as ours (See Simon [23]). This concept is strictly more stringent than the one we use in this paper, and thus yields uniqueness also.

⁹ They are related to a version of the game discussed in Fudenberg-Tirole [7].

maintains a larger market share once the second firm enters. Specifically, if j enters at t_j and i follows at $t_i \geq t_j$, then i receives the fraction $[\frac{1}{2} - \frac{1}{2}(t_i - t_j)]$ of industry flow profits.

Entering an Industry with Loyal Customers.

We begin by analyzing the problem from an "intuitive" standpoint. While the logic presented below is convincing, it is a nontrivial matter to formalize it within a rigorous game-theoretic model. Indeed, one of the main goals of our research program is to provide a simple framework within which these arguments can be made formal.

There is a unique "intuitive" solution: both players enter at time $t^* = \frac{1}{4}$. Though a monopolist would enter at time zero (since $\pi^m = 2$), neither duopolist will attempt preemptively to monopolize the market, because it anticipates that the other will follow suit immediately if it enters. We then show how the problem can be formulated as a continuous-time game. The game has multiple SGP equilibria. However, the unique equilibrium that satisfies our stronger solution criterion (see fn. 8) implements the intuitive solution. By contrast, when the problem is formulated as a discrete-time game, there is a unique equilibrium, but it is quite different from the intuitive one. Finally, we use the example to illustrate the relationship between continuous- and discrete-time games.

First observe that if j enters the market at t_j and i waits for a while before entering, his market share will fall at a faster rate than his entry cost declines. Therefore, i should respond by entering "immediately."¹⁰ Precisely, let $\Pi_i(t_i, t_j)$ denote i 's payoff when he and j enter at, respectively, t_i and t_j and let $\tau_i(t_j)$ denote the maximizer of $\Pi_i(\cdot, t_j)$ on $[t_j, 1)$. It is easy to check that for each t_j , $\tau_i(t_j) = t_j$.¹¹

~~Now consider an agent's problem before either firm has entered the market. Assume that j never enters first, but follows suit immediately if i enters. Firm i 's problem is to choose t_i to maximize $\Pi_i(t_i, \tau_j(t_i))$ on $[0, 1)$. This function is strictly concave and is maximized when $t_i = \frac{1}{4}$. By symmetry, each~~

¹⁰ A priori, the meaning of "immediately" is unclear in continuous time. It will have a precise meaning once we have specified our model, but for the moment, we will remain vague.

¹¹ We have: $\Pi_i(t_i, t_j) = \int_0^{t_i} 0 dt + \int_{t_j}^1 3[\frac{1}{2} - \frac{1}{2}(t_i - t_j)] dt - (1 - t_i)^2 = 3/4(1 + t_j - t_i)(1 - t_i) - (1 - t_i)^2$. For every t_j , $\Pi_i(\cdot, t_j)$ is strictly decreasing on $[t_j, 1)$ and therefore attains a constrained maximum at $\tau_i(t_j) = t_j$. Moreover, $\Pi_i(t_j, t_j)$ is strictly positive. This verifies that $\tau_i(t_j) = t_j$.

agent's optimal action is to wait until time $\frac{1}{4}$ and then enter. This, then, is our "intuitive" solution.

We now explain how to formulate this problem as a continuous time game. The time interval is $[0, 1)$. Each agent has two actions, 'enter' and 'wait'. Once an agent has entered the market, he cannot leave again. A decision node in this game is a point in time, paired with a complete description of past activity in the game. In this context, past activity can be summarized by a list of the times--if any--that agents have entered the market in the past. A strategy is a function mapping each of these nodes to either 'wait' or 'enter'.

The "intuitive" outcome is a SGP equilibrium outcome for this game, but not the unique one. For every $t < \frac{1}{4}$, there is a symmetric SGP strategy profile --[II.1(t)] below--that implements the outcome in which both players enter simultaneously at t . For each i , the strategy is:

"play 'wait' at time zero; at every time $0 < s < t$, play 'wait' if both players have always [II.1(t)] played 'wait' in the past, otherwise play 'enter'; play 'enter' at every $s \geq t$, regardless of the history."

We now show that these strategies indeed implement the specified outcome and form a subgame perfect equilibrium for our model. Fix $t \leq \frac{1}{4}$. If the strategies [II.1(t)] are played on an arbitrary discrete-time grid, the resulting outcome is that both players jump simultaneously at the first grid-point weakly beyond t .¹² In the limit, they move *exactly* at t . Each player's payoff from this outcome is $\Pi_i(t, \tau_j(t)) = (\frac{3}{2} + t - 1)(1 - t)$. To see that [II.1(t)] is an equilibrium, we need to check that, starting from each decision node in the game, each agent's action is a best response to the other's. The only nodes that are nontrivial to check are those of the form $(t, \text{nobody has entered})$. If $s \geq t$, then j 's strategy has him entering at s , regardless of i 's action. Since $\tau_i(t) = s$, entering simultaneously is as good as any response for i . Now suppose that $s < t$ and that i deviates by entering at this time instead of waiting. When i 's deviant strategy is combined with j 's and played on any discrete-time grid, the resulting outcome will be: i enters at the first grid-point weakly beyond s , and j follows suit at the next. In the limit, i 's termination at s will result in an *instantaneous* response by j , i.e., the two agents will move "consecutively, but at the same instant of time." Since, $\Pi_i(\cdot, \tau_j(\cdot))$ is strictly concave, i 's defection at $s < t \leq \frac{1}{4}$ will result in a lower payoff

¹² We will use the expression "weakly" beyond t to mean greater than or equal to t . Similarly, "weakly before" t means less than or equal to t .

than had he waited until t . This establishes that i 's strategy is indeed optimal, given j 's (credible) strategy of immediate response.

After iterated elimination of weakly dominated strategies (see fn. 8), the only outcome that survives is our "intuitive" solution: both players move at $1/4$. The elimination process proceeds as follows. Any strategy that has i waiting for an interval before entering, once j has already entered, is weakly dominated by the strategy that is identical to it, except that i follows suit immediately. After one round of elimination, the strategy [II.1($1/4$)] weakly dominates all other "instantaneous response" strategies. The proof of this parallels our earlier argument (pp. 6-7) that $t = 1/4$ maximizes $\Pi_i(\cdot, \tau_j(\cdot))$.¹³

Summarizing, our model yields the strong and intuitive prediction that each agent will be deterred from preempting by the other's credible declaration of intent to follow suit immediately. By contrast, in every discrete-time version of this game, the unique SGP outcome is that each player enters at time zero and earns the suboptimal payoff $\Pi_i(0, \tau_j(0))$. To see this, observe that there can be no equilibrium in which termination occurs at any positive time t , because either player could gain the slight advantage of one-period's monopoly profits--as well as an increased market share--by preempting one period before t .¹⁴

This result is counterintuitive. When reaction lags are very small, the downside potential of the strategy "follow suit as quickly as possible if j enters, otherwise enter at $1/4$ " is virtually zero, while its upside potential can be as large as $\Pi_i(1/4, 1/4) - \Pi_i(0, 0)$. Informally, therefore, this strategy "ought to" dominate early entry. This intuition cannot be formalized in a discrete-time model: when it comes to the elimination of "inferior" strategies, "virtually zero" is not small enough. We cannot eliminate a strategy on the grounds that it is "nearly weakly dominated." On the other hand, when lags are *literally zero*, the downside risk associated with "wait and see" is exactly zero. Weak dominance can then be invoked to reduce the set of equilibria to a singleton set. The example thus illustrates our earlier observation (pp. 2-3) that from a modelling standpoint, there is a significant difference between very short reaction lags and no lags at all.

¹³ More precisely, compare [II.1($1/4$)] against the instantaneous response that has i entering unilaterally at some $t < 1/4$. If j enters weakly before t , then i 's response is the same in either case: enter immediately. If j does *not* enter unilaterally until after t , then if i plays [II.1($1/4$)], *both* players will enter strictly later than they would have, if i had preempted.

¹⁴ The above argument is entirely familiar. It closely parallels the argument that "always defect" is the unique equilibrium in the finitely repeated Prisoners' Dilemma.

The example also illustrates one of our results relating continuous- to discrete-time equilibria. This result states that under weak conditions, any continuous-time equilibrium is close to some approximate equilibrium for any nearby discrete-time game. In this example, these discrete-time equilibria are just the restrictions of the continuous-time strategies [II.1(t)] to the appropriate grids. To see this, observe that if j plays [II.1(t)] on any grid, then j can gain only a one-period advantage by entering before t , because i will follow at the very next grid point. For any ε , we can choose a grid sufficiently fine that the gain to deviating will not exceed ε .

Entering an Industry with Fickle Customers.

This second variant of our game differs from the first in two respects. First, monopoly profits π^m are smaller, lying in the interval $(\frac{3}{4}, 2)$. Second, duopolists share the market equally, regardless of their relative entry times. This second difference radically alters the nature of the problem. Specifically, i 's payoff if he enters at t_i , given that j has entered at t_j , is $\Pi_i'(t_i, t_j) = \frac{3}{4}(1 - t_i) - (1 - t_i)^2$. If $t_j < \frac{1}{4}$, i will choose to wait until $\frac{1}{4}$ and then enter. Now consider the problem facing i at the beginning of the game. Since $\pi^m \in (\frac{3}{4}, 2)$, the optimal time for a monopolist to enter is $t^m = 1 - \frac{1}{2}\pi^m \in (0, \frac{1}{4})$. Ideally, player i would also choose to enter at this time, and enjoy monopoly profits until j enters at $\frac{1}{4}$. When there are two potential entrants, however, monopoly rents will be dissipated by preemptive entry. The only possible equilibrium outcome in this model (modulo relabelling of agents) is that one of the players enters at the "preemption time" $t^P \in (0, t^m)$, where t^P is defined as follows. If player i would strictly prefer to enter the market at time zero and monopolize it until $\frac{1}{4}$, rather than enter as a duopolist at $\frac{1}{4}$, then $t^P = 0$. Otherwise, t^P is the unique time at which i is indifferent between being the first entrant at t^P and the second at $\frac{1}{4}$.

On any discrete-time grid, the following pair of "chattering" strategies form an equilibrium that implements this outcome:

Player #1: "play 'wait' before t^P ; at the k 'th grid point in the interval $(t^P, \frac{1}{4})$, play 'enter' if k is *odd* and both players have always played 'wait' in the past, otherwise play 'enter'; at every grid point after $\frac{1}{4}$, play 'enter' regardless of the history." [II.2-1]

Player #2's strategy, [II.2-2], is identical to #1's, except that *odd* is replaced by *even*. These strategies implement the outcome: #1 enters at the first grid point after t^P , while #2 enters at $\frac{1}{4}$. To see that these strategies

form an equilibrium, observe that if $t^P = 0$, then #2 cannot preempt; if $t^P > 0$, then by definition of t^P , #2 does not want to preempt.

These "chattering" or "alternating" strategies have no analogs in continuous time. Indeed, the continuous-time version of the game has no pure strategy equilibrium.¹⁵ The reason is that the chattering strategies depend on the fact that discrete-time is well-ordered, i.e., that there is a first, second, third, ..., grid point after t^P . In the equilibrium just described, player #1 is willing to enter if he reaches an odd grid point in the interval (t^P, t^m) only because if he didn't enter, then #2 would enter at the *next* grid-point. Since there is no notion of "next" in continuous-time, it is not possible to formulate continuous-time counterparts of [II.2]. The strategies thus depend *intrinsically* on the special properties of discrete time. One cannot, therefore, have confidence in the predictions of a model that involves strategies like [II.2], *unless* one believes that "in actuality," the notion of a "next time" *is* really meaningful. If not, these "chattering" equilibria must be viewed as artifacts of the discrete-time formulation.

III. Specifying a continuous-time game form: the conceptual issues.

This section contains examples illustrating various conceptual problems that we confront in this paper. The first subsection focuses on issues that arise because induction cannot be applied in continuous time. The second is concerned with problems specific to our particular view of continuous time as "discrete-time, but with an infinitely fine grid." In each of our examples, each agent has two actions, 'left' (lf) and 'right' (rt). We emphasize at the outset that the issues raised in this section are more primitive than those discussed in section II. The difficulties that concern us at the moment are at the level of defining a *game-form*, rather than a game. That is, we discuss strategies and outcomes, but make no mention of payoffs or equilibria.

Defining a game-form without induction on the set of time nodes.

We began with the observation that since induction cannot be applied in continuous time, there is no obvious way to specify an outcome function. Our example (0.1) on (p. 1) showed that there may be a con-

¹⁵ Once behavior strategies are introduced into the model, there is a unique equilibrium (modulo relabelling of agents). See Simon [21] and Simon-Stinchcombe [24].

tinuum of outcomes that are all "compatible" with a given strategy. We now show that for some strategies, there may be no conceivable outcome that is sensible. More precisely, we will say that an outcome is consistent with a given strategy if except at a finite set of time points, agents' strategies are "compatible" with the outcome, i.e., they are actually playing what their strategies specify they should play. (A stronger definition of consistency would require agreement between strategies and the outcome at *every* point in time. We show below that this more stringent notion is too strong a requirement.) For the strategy profile that follows, there is no consistent outcome.

The example is called the Odd Couple profile.¹⁶ There are two players. The first attempts to "match" the second, while the second attempts to "mismatch" the first.

Player #1: "play 'left' at time zero; at every positive time t , play 'left', if player #2 was playing 'left' just before t ; otherwise, play 'right'." (III.1-1)

Player #2: "play 'left' at time zero; at every positive time, play 'left', if player #1 has been playing 'right' just before t ; otherwise, play 'right'." (III.1-2)

In discrete-time, "just before t " can only mean "at the last grid point before t ." When these strategies are played on any discrete-time grid, the resulting outcome is the finite-length cycle that begins with (lf lf, lf rt, rt rt, rt lf, lf lf, ...). If we *could* apply induction in continuous time--i.e., identify a "last continuous-time moment before t "--then the uniquely identified continuous-time outcome would be the infinite-length analog of this cycle, with the cycles occurring infinitely fast. Since there is no such last moment, the Odd Couple strategies make sense only if they are fundamentally reinterpreted. For example, in continuous time, "just before t " might mean "at some point in every open interval $(t - \delta, t)$." However, when the Odd couple profile is reinterpreted in this way, there is *no conceivable* outcome that is consistent with it.¹⁷ We conclude from this that agents in *any* continuous-time model must be prohibited from choosing

¹⁵ The above assertion remains true if the notion of consistency is weakened by replacing "finite" with "nowhere dense."

¹⁶ This example is related to one discussed by Krishna [13]. It was brought to our attention by Karl Iorio.

¹⁷ Suppose there exists an outcome that is consistent with the Odd Couple profile. Assume first that this outcome has player i playing a single action throughout some interval, apart from a finite number of exceptions (henceforth a.f.e.). Since the outcome is consistent with j 's strategy, j must also be playing a single action on the interval (a.f.e.). Since #1 "matches" and #2 "mismatches," however, at least one of the players' actions on this interval must be incompatible (a.f.e.) with his strategy. It follows, therefore, that each player must play both 'left' and 'right' infinitely often on every open interval in $[0, 1)$. In particular, for every positive s and δ , player #1 must play 'right' at some point in the interval $(s - \delta, s)$. But in this case, #2's strategy profile instructs him to play 'left' at every point in $[0, 1)$. Since the outcome is consistent with #2's strategy, he must indeed be playing 'left' throughout (a.f.e.). But we have just established that this cannot happen. This contradiction completes the proof.

strategies of this kind.

Our second example illustrates a rather different kind of problem. Quite apart from whether or not an outcome can be uniquely identified, there is the question of how it should be represented as a mathematical object. In particular, we need a way to represent the limit of discrete-time outcomes in which agents move at consecutive grid points.¹⁸ Without induction, there is no notion of "consecutive." Therefore, we *must* think of the limit chain of consecutive moves as occurring at a *single* instant of time. On the other hand, it is important that we keep track of the *order* in which the moves occur. This means that we need a way to represent the apparently paradoxical idea of moves that occur consecutively but at the same moment in time.

Our third example, called "Follow the Leader," illustrates the point. There are three players: #1 leads, #2 follows #1 and #3 follows #2. The strategies for players #1 and #*i* = 2, 3 are:

Player #1: "At every time $t < \frac{1}{2}$, play 'left'; at every t beyond $\frac{1}{2}$, play 'right'." (III.2-1)

Player #*i*: "play 'left' at time zero; at every positive time, play 'left' if #*i*-1 has played 'left' at every point in the past; otherwise, play 'right'." (III.2-i)

In discrete-time, the outcome generated by this profile is: #1 jumps at the first grid point weakly beyond $\frac{1}{2}$; #2 follows suit at the second, and then #3 at the third. In the limit, the time intervals between these moves collapse and we must find a mathematical form to represent this limit. It cannot be represented as a function on $[0, 1)$ without some loss of information. Any such function would have to have exactly one discontinuity point, i.e., at $t = \frac{1}{2}$. However, a single discontinuity point cannot "carry" more than two pieces of "order" information, one fewer than we need. Thus, any functional representation must necessarily confound at least some of the information that is available in the passage to the limit.^{19,20}

We resolve this issue is by representing limit histories as vectors of pairs. Each pair corresponds to a distinct jump that occurs in the passage to the limit. For example, the limit outcome for "Follow the Leader" is represented as the "4-length" vector of pairs: $[(0, lf\ lf\ lf), (\frac{1}{2}, rt\ lf\ lf), (\frac{1}{2}, rt\ rt\ lf), (\frac{1}{2}, rt\ rt\ rt)]$. The first pair denotes the actions that agents take at the start of the game. The others list their jumps and jump

¹⁸ We saw a simple example of this in section II. If one agent entered the industry, the other would follow suit "immediately." At that point, we were being sufficiently vague that the issue could be avoided.

¹⁹ "Confound" means "to mix up or mingle so that the elements become difficult to distinguish or impossible to separate." (O.E.D.)

²⁰ If the function were right-continuous at $\frac{1}{2}$, then the order in which the agents move would be suppressed completely. If #1 moves left-continuously, and the others right continuously, then the fact that #2 moves before #3 would be suppressed. Etc.

destinations in the correct order. We interpret this outcome as: each player plays 'left' on the interval $[0, \frac{1}{2})$; at time $\frac{1}{2}$, they consecutively switch from 'left' to 'right'; each of them then plays 'right' on $(\frac{1}{2}, 1)$. Note that this outcome is "consistent" with the profile (III.2) above, in the sense defined on pp. 10-11. At every $t \neq \frac{1}{2}$, the outcome has agents are playing what their strategies specify they should play. At $t = \frac{1}{2}$, agents are not playing any *single* action, so the notion of "compatibility" defined above (pp. 10-11) cannot be applied.

Continuous time outcomes as the limit of discrete time outcomes.

In this subsection, we focus on problems that are specific to the particular modelling approach outlined in section I. We need to identify a class of strategies that yield well-defined, consistent outcomes, in the sense defined on pp. 10-11. To accomplish this, we will impose three restrictions on strategies. We will introduce these assumptions informally, and illustrate what goes wrong when they are violated.

The first class of strategies that must be excluded are those in which agents move unboundedly often. The Odd Couple profile (Example (III.1-2)) illustrates the reason. We observed that no conceivable continuous time outcome can sensibly be associated with this profile. Not surprisingly, therefore, our specific proposal for defining outcomes breaks down in this instance: the cyclic outcomes generated by the profiles in discrete time have no limit in the topology we impose on outcomes.²¹ To exclude strategies of this kind, we will impose the assumption (F1 below) that for every strategy an agent chooses, there is a uniform upper bound on the number of jumps that the agent makes.

From an abstract mathematical standpoint, F1 is a seriously restrictive assumption. Moreover, we believe it is--and will remain--the most serious practical limitation of our model. For example, we see no prospect of ever being able to formulate the idea of an infinitely repeated game on the unit interval.²² There are, however, many applications in which our finiteness condition is satisfied automatically. For example, are games in which there is a given, ordered progression of moves that players can make. In a duel, for instance,

²¹ Indeed, there is no Hausdorff topology for which they converge.

²² This is because we insist on assigning outcomes to strategies and then valuations to outcomes. If one is willing to go directly from strategies to payoffs without specifying what actually happens when the strategies are played, then the frequency with which agents move need not be a problem.

agents might have a finite number of bullets to fire. Once these bullets have been fired, the agent has no more moves left to make. In other models, agents may incur costs every time they change their actions. If agents' financial resources are bounded in these models, then the constraint will again be satisfied.

Our second regularity condition on strategies is that they should not fluctuate too wildly with respect to time. The following example illustrates why this restriction is needed. There is one agent, whose strategy is:

"play 'left' at time zero; if t is a rational number larger than $\frac{1}{2}\sqrt{2}$, and 'left' has always been played in the past, play 'right'; otherwise, play 'left'." (III.3)

When (III.3) is restricted to a discrete-time grid, the outcome that results will be very sensitive to the particular grid. If the grid contains no rational numbers beyond $\frac{1}{2}\sqrt{2}$, the agent will play 'left' throughout the game. Otherwise, he will jump at the first rational grid point beyond $\frac{1}{2}\sqrt{2}$ (which could be anywhere between $\frac{1}{2}\sqrt{2}$ and 1). To eliminate strategies like this, we will require (as F2 below) that strategies depend piecewise continuously on time.

Even piecewise continuous strategies may be very sensitive to grid structure. Consider for example:

"play 'left' at time zero; if $t = \frac{1}{2}\sqrt{2}$, and 'left' has always been played in the past, play 'right'; otherwise, play 'left'." (III.4)

If (III.4) were restricted to a discrete-time grid without further modification, the resulting outcome would clearly depend on whether or not $\frac{1}{2}\sqrt{2}$ were contained in the grid. It is easy, however, to neutralize this sensitivity by "moulding" each strategy profile to each grid, before "playing" it. For example, before playing (III.4) on a grid, we will reinterpret it to mean: "play 'right' at the first grid point beyond $\frac{1}{2}\sqrt{2}$." The modified strategy will be called a "graph preserving restriction." Clearly, the outcomes generated by the reinterpreted strategies will have the well-defined limit: $[(0, lf), (\frac{1}{2}\sqrt{2}, rt)]$.

We now explain our most severe restriction on strategies. In general, the jump-times of the discrete-time histories generated by a given strategy will differ, but converge "from above," as the grids become finer. (Example (III.4) above illustrates this.) To ensure the existence of a unique, consistent limit history, we must guarantee that agents' actions *later* in the game are not too sensitive to the precise times at which *earlier* jumps occurred. Now, our model is constructed so that the only possible sequences of jump-times that can arise are nonincreasing ones. Therefore we will impose the restriction that agents' strategies be continuous with respect to sequences of histories whose jump-times converge from above (i.e., from "the right").

The precise form of continuity we need turns out to be very strong indeed. Roughly, we require that there exists some countable partition of the universe of possible "pasts" into sets that are "closed from above," such that the strategy does not distinguish between any two "pasts" that are members of the same set in the partition.

This is a severely restrictive assumption. Fortunately, it plays a less fundamental role in our model than the other two assumptions. A perfectly sensible game form can be defined for the class of strategies that satisfy only F1 and F2. In particular, we can construct an outcome function for this set of strategies that is consistent in the sense above (pp. 10-11). If F3 is violated, however, our interpretation of continuous-time outcomes as limits of discrete-time outcomes, will no longer be valid. Strategies may generate discrete-time outcomes that have no limits. Alternatively, a unique limit may exist, but differ from the (consistent) outcome identified by the constructed outcome function. A related point is that F3 plays an essential role in the proof of our theorems relating continuous- and discrete-time equilibria.

There is, nonetheless a sense in which F3 is dispensable. Subgame perfect equilibrium profiles that satisfy F1 and F2 can be arbitrarily closely approximated by approximate equilibrium profiles that satisfy F1-F3. This means that if F3 proves too restrictive, we can relax it and still provide a satisfactory interpretation of the model and its predictions.²³

Example (III.5) below illustrate the role of the condition. Player #1's strategy is independent of the past, while #2's depends linearly on histories. The strategies are thus extremely well behaved, and *ought* to be allowable. The example thus illustrates why such a stringent condition is needed and the importance of being able to relax it in the way described above.

Player #1: "Play 'left' at all times before $\frac{2}{3}$; otherwise play 'right'." (III.5-1)

Player #2: "play 'left' before $\frac{1}{3}$, play 'right' between $\frac{1}{3}$ and twice the last time some agent jumped; beyond this point, jump back to 'left', provided that player #1 has always played 'left' in the past; otherwise, continue to play 'right'." (III.5-2)

When these strategies are played on a discrete-time grid, the outcome depends on whether $\frac{1}{3}$ is contained in the grid. To see this, play the profile on a sequence of grids that contain $\frac{1}{3}$. In the discrete-time outcomes, player #2 makes his first jump at exactly $\frac{1}{3}$. At the first grid point weakly beyond $\frac{2}{3}$, #2 observes that #1

²³ See Simon-Stinchcombe [24].

has always played 'left' in the past, and so jumps back again to 'left'. At the same instant, player #1 jumps to 'right'. The limit outcome, therefore, is $[(0, lf lf), (1/3, lf rt), (2/3, rt lf)]$. Now assume that the n 'th grid in the sequence contains the points $1/3 + 1/n$ and $2/3$, but not $1/3$. When the profile is played on the n 'th grid, player #2's first jump will occur at $1/3 + 1/n$ and #1 will jump at $2/3$. Player #2 does not consider jumping again until $2/3 + 2/n$. By this time, however, player #1 has already jumped and is no longer playing 'left'. According to #2's strategy, therefore, he should continue to play 'left'. For this sequence of grids, then, the limit outcome is $[(0, lf lf), (1/3, lf rt), (2/3, rt rt)]$.

IV. The Formal Model.

Our game is played on the interval $[0, 1)$ (see fn. 6). Let I denote a finite set of agents, with generic element i . At any point in time, each agent can choose from a finite set of actions, A_i , which is called the action set for player i . Let $A = \prod_{i \in I} A_i$ denote the set of possible action profiles, with generic element $a = (a_i)_{i \in I} \in A$.²⁴ A profile a will frequently be written as (a_i, a_{-i}) , where $a_{-i} = (a_j)_{j \neq i}$.

Histories.

A history of the game is a string of pairs, representing the "jumps" that agents have made during the play of the game. For an explanation and interpretation of the formalism below, the reader is referred back to our earlier discussion on pp. 12-13. Each "jump" is denoted by a pair $(t, a) \in [0, 1) \times A$, where ' t ' is called a jump-time and ' a ' is called a jump-destination. A " k -length-history," h , is a string of k jumps, i.e., $h = [(t^1(h), a^1(h)), \dots, (t^\kappa(h), a^\kappa(h)), \dots, (t^k(h), a^k(h))]$, where $(t^1(h), a^1(h))$ is the initial "jump," necessarily taken at the beginning of the game, and $(t^\kappa(h), a^\kappa(h))$ is the κ 'th jump. We let $k(h)$ denote the length of the history h but will usually denote the last (i.e., the $k(h)$ 'th) jump of a history by $(t^{last}(h), a^{last}(h))$. We will also represent a history h as a pair of vectors $(t(h), a(h))$, where $a(h) \in A^{k(h)}$ is called the action vector of h and $t(h) \in [0, 1)^{k(h)}$ the vector of jump times of h . We remind the reader

²⁴ In practice, the actions available to an agent at any point in time may depend on the past evolution of the system. For example, in certain simple games, it is natural to assume that jumps cannot be reversed. For instance, in a duel, there is no obvious way to reverse the action "fire a bullet." We will enforce such restrictions within our model by an appropriate choice of payoff function: an inadmissible sequence of jumps will result in a prohibitively low payoff.

that a novel feature of histories in our model is that agents can jump consecutively, but at the "same" instant in time. For this reason, $t(h)$ need not be *strictly* increasing. Formally, h must satisfy:

- (i) $t^1(h) = 0$;
- (ii) the vector of jump times is nonincreasing, i.e. for $1 \leq \kappa < k(h)$, $t^\kappa(h) \leq t^{\kappa+1}(h)$;
- (iii) adjacent actions are distinct, i.e. for $1 \leq \kappa < k(h)$, $a^{\kappa+1}(h) \neq a^\kappa(h)$.

Let H^k denote the set of k -length histories. In addition to the positive-length histories, there will also be an artificial "zero-length" history, denoted by the null symbol " \emptyset ." Let $H = \bigcup_{k=0}^{\infty} H^k$ denote the universe of finite length histories. Since we are restricting agents to move only finitely often, H is exactly the mathematical space we need in order to represent outcomes.

We now define a metric, d^H , on H . Two histories will be at distance one from each other if their lengths are different or if they have different action vectors. Otherwise, the distance between them will be the sum of the absolute values of the differences between the corresponding jump times of the two histories,

$$\text{i.e., } d^H(h, h') = \begin{cases} \sum_{\kappa=1}^{k(h)} |t^\kappa(h) - t^\kappa(h')| & \text{if } k(h) = k(h') \text{ and } \mathbf{a}(h) = \mathbf{a}(h') \\ 1 & \text{otherwise} \end{cases} .$$

In this metric, a sequence of histories (h^n) converges to h if and only if, for n sufficiently large, h^n is an $k(h)$ -length history whose action vector agrees with h 's and the vectors of jump-times of the h^n 's, i.e., the $t(h^n)$'s, converge to $t(h)$, in the usual sense of convergence on $\mathbf{R}^{k(h)}$. Thus, the set of admissible histories is the disjoint union of a collection of finite dimensional subsets of Euclidean space. Note that in this metric, a history in which i moves just before j is far away from the one that is identical except that j moves before i . Moreover, both these histories are far away from the one in which i and j move simultaneously.

We will denote by $h_{1\kappa}$ the κ -length truncation of h : $h_{1\kappa}$ is the history obtained from h by truncating it after the first κ jumps. For example, if $h = [(0, lf\ lf), (\frac{1}{2}, rt\ lf), (\frac{1}{2}, rt\ rt)]$, then h_{11} is the 1-length history $(0, lf\ lf)$, and h_{12} is $((0, lf\ lf), (\frac{1}{2}, rt\ lf))$.

Decision Nodes.

A decision node is a point in time, paired with a complete description of past activity in the system. Just as in discrete-time, the set of decision nodes is the domain on which a strategy in our model is defined. Because our class of admissible histories is small, the set of possible decision nodes is also small. This is important, because it bounds the complexity of the model. Pure strategies are functions with a simple domain and finite range and are therefore relatively easy to deal with.

There is a distinguished node, denoted by $(0, \emptyset)$, that represents the start of the game. A regular decision node is a pair (t, h) , where t is a point in time, and h is a positive-length history whose last jump-time is no greater than t . For example, in a two person game, the regular decision node $(t, [(0, lf lf), (\frac{1}{2}, rt rt)])$, is interpreted as follows, for every $t \geq \frac{1}{2}$: the present time is t ; agents chose 'left' at the beginning of the game, and simultaneously jumped to 'right' at $\frac{1}{2}$.

Observe that we allow decision nodes of the form (\bar{t}, \bar{h}) , where \bar{t} coincides with the last jump-time of \bar{h} . In such cases, \bar{t} is interpreted as the first available opportunity after the last jump in \bar{h} occurred. On every discrete-time grid, this first opportunity is well-defined. Since we view continuous-time as "discrete-time, but with an infinitely fine grid," it is natural--and extremely convenient--to allow for a corresponding "first" moment in our continuous time model.

Let DN denote the set of all decision nodes, i.e., $DN = \bigcup_{h \in H} \{(t, h) : t \in [t^{last}(h), 1)\} \cup \{(0, \emptyset)\}$.

We denote the generic element of DN by (t, h) . We emphasize that DN is *not* the cartesian product of $[0, 1]$ and H . Whenever we refer to $(t, h) \in DN$, we are implicitly asserting that *either* $(t, h) = (0, \emptyset)$ *or* that t weakly exceeds the last jump-time of h .

Discrete-time games.

We now explain how to represent a discrete-time game in our framework. For each discrete subset, R , of T , we will say that a history in H is R -admissible if its jump-times are *strictly* increasing and contained in R .²⁵ Let H^R denote the set of such histories, i.e.,

²⁵ There is, clearly, no sensible way to interpret in discrete time a history in which consecutive jumps that occur at the same instant.

$H^R = \{h \in H: \text{for all } k < k(h), t^k(h) \in R \text{ and } t^k(h) < t^{k+1}(h)\}$. Now let DN^R denote the set of R -admissible decision nodes, i.e., points in R , paired with R -admissible histories whose last jump times occur *strictly* before r . That is, $DN^R = \{(r, h) \in DN: r \in R, h \in H^R \text{ and } t^{\text{last}}(h) < r\} \cup \{(0, \emptyset)\}$. A R -admissible strategy is any function from DN^R to A_i .

Profiles of discrete-time strategies uniquely define outcomes in the usual way. Specifically, fix a discrete-time grid R , a strategy g and a regular decision node $(\underline{t}, \underline{h})$ for the game-form on R .²⁶ The profile g defines by induction a unique history from $(\underline{t}, \underline{h})$ --call it \bar{h} --in exactly the usual way. The first $k(\underline{h})$ jumps of \bar{h} are the same as those of \underline{h} . If agents play $a^{\text{last}}(\underline{h})$ from \underline{t} until the end of the game, then the specification of \bar{h} is complete. Otherwise, $(t^{k(\underline{h})+1}(\bar{h}), a^{k(\underline{h})+1}(\bar{h}))$ is defined as follows: $t^{k(\underline{h})+1}(\bar{h})$ is the first grid-point after \underline{t} at which agents choose some action profile that differs from $a^{\text{last}}(\underline{h})$. Next, set $a^{k(\underline{h})+1}(\bar{h})$ equal to the actions chosen by agents at $t^{k(\underline{h})+1}(\bar{h})$, i.e., $a^{k(\underline{h})+1}(\bar{h}) = f(t^{k(\underline{h})+1}(\bar{h}), \underline{h})$. Now inductively define the remaining jumps of \bar{h} in the corresponding way. It should be clear that this formulation of a discrete-time game-form is *exactly* to the conventional one (see, for example Fudenberg-Levine [4]) *except* that our representation of histories is terser than the usual one: we record only the *jumps* that agents make, rather than listing the actions they choose at *every* discrete-time node.

Continuous-time pure strategies.

We now state precisely the three assumptions discussed in the preceding section. The formal description will be rather terse. The reader is urged to refer back to the corresponding heuristic discussion in section III at each point.

Assumption F1: uniformly bounded number of jumps (cf. pp. 13-14).

This condition states that for each strategy, there exists some upper bound n such that if i has jumped more than n times in the past, he never jumps again. Unless this condition is carefully stated, it will be hopelessly restrictive. The problem is that each strategy must instruct the agent how to play at every possible decision node, including ones that would never be reached had the agent been playing this strategy in the

²⁶ If $(\underline{t}, \underline{h}) = (0, \emptyset)$, modify the procedure below in the obvious way.

past. In particular, there will be "off the equilibrium path" nodes at which the agent has made jumps that were not specified by his strategy. For example, consider the constant strategy "always play 'left'." For any given n , there are "off the equilibrium path" nodes at which the agent has already jumped more than n times and has ended up playing, say, 'right'. At such nodes, the constant strategy instructs the agent to jump once more. Thus, a restriction that imposed an unqualified upper bound on the number of jumps that an agent can make would exclude even constant strategies!

To avoid excluding such strategies, we will require that our upper bound apply only in situations where i 's last jump could conceivably have occurred, had i been playing f_i in the past. To formalize this idea, we let $k_i(h)$ denote the largest κ at which i jumps, i.e., such that $a_i^\kappa(h) \neq a_i^{\kappa-1}(h)$. For example, if $\bar{h} = [(0, lf), (1/2, rt lf), (1/2, rt rt)]$, then $k_1(\bar{h}) = 2$ and $k_2(\bar{h}) = 3$. We now say that a decision node (t, h) is "compatible with f_i " if there exists some time s either at or immediately after $t^{k_i(h)}$ such that at the node $(s, h_{1_{k_i(h)-1}})$, i 's strategy actually calls for the jump to $a_i^{k_i(h)}$. In symbols, (t, h) is compatible with f_i if for all δ , there exists $s \in [t^{k_i(h)}(h), t^{k_i(h)}(h) + \delta)$ such that $f_i(s, h_{1_{k_i(h)-1}}) = a_i^{k_i(h)}(h)$. (For example, suppose that f_1 is the constant strategy "always play 'left'," and \bar{h} is defined as above. For $s \geq 1/2$, the node (s, \bar{h}) is *incompatible* with f_1 because for every $s \in [1/2, 1)$, $f_1(s, \bar{h}_{11}) = lf$.) Assumption F1 requires that there is some upper bound on jumps that applies to all decision nodes that are compatible with f_i . Precisely, a strategy f_i satisfies F1 if

there exists $n \in \mathbf{N}$ such that for all $(t, h) \in DN$, if (a) (t, h) is compatible with f_i and

(b) the number of κ 's such that $a_i^{\kappa+1}(h) \neq a_i^\kappa(h)$ exceeds n , then $f_i(t, h) = a_i^{k_i(h)}(h)$.

The order of quantifiers is important. There are strategies with arbitrarily large upper bounds, but *some* uniform upper bound must exist for each strategy.

Assumption F2: piecewise continuity with respect to time (cf. p. 14).

Our next condition states that for each history h , the function $f_i(\cdot, h)$, must be piecewise continuous on that interval on which it is defined, i.e., on $[t^{last}(h), 1)$. That is, we require that $f_i(\cdot, h)$ be discontinuous on at most a finite set of points. Note that this and the previous assumptions are unrelated. F1 concerns the

number of jumps an agent has made in the past, while the present assumption restricts the way strategies can vary with time in the future.

Assumption F3: strong right continuity with respect to histories (cf. pp. 14-16).

Our final condition restricts the way agents can condition their actions on the past. The requirement is that strategies be insensitive to slight differences in the jump times of histories, provided these times are converging "from above." More precisely, the condition is: for every h , there exists a positive δ such that if h' is within δ of h in the metric d^H and if all of the jump times of h' weakly exceed the corresponding jump times of h , then for every $s > t^{last}(h')$, the nodes (s, h) and (s, h') , are treated in *exactly* the same way. Also, $(t^{last}(h), h)$ and $(t^{last}(h'), h')$ are treated identically. In symbols, we say that a strategy f_i is strongly right continuous w.r.t. h if for all $h \in H$, there exists $\delta > 0$ such that if $d^H(h', h) < \delta$ and $t(h') \geq t(h)$, then $f_i(t^{last}(h), h) = f_i(t^{last}(h'), h')$ and, for all $s > t^{last}(h')$, $f_i(s, h) = f_i(s, h')$. A simple but useful class of strategies that satisfy F3 are "jump-time-independent" strategies. Strategies of this kind depend on histories only through their action vectors. That is, f_i is jump-time-independent if $f_i(\cdot, h)$ depends on $a(h)$, but not on $t(h)$. A strategy will be called admissible if it satisfies conditions F1-F3. A profile of such strategies will be called an admissible strategy profile.

Continuous-time outcomes as the limits of discrete-time outcomes.

We first show how a profile of strategies generates a well-defined discrete-time outcome when it is restricted to an arbitrary discrete-time grid. As we explained on p. 14, before "playing" a profile f on a given grid, we first "mould" it to take account of the particular structure of the grid. The "moulded" strategy is called the "graph preserving restriction" of f and defines by induction a history in H . We then now state precisely the relationship between our continuous-time outcome function and the discrete-time outcome functions. For any strategy profile f and decision node $(\underline{t}, \underline{h})$, there exists a unique history η with the following property: for any sequence of increasingly fine grids, (R^n) , the sequence of outcomes generated by playing

²⁷ This and the following subsection involve considerable technical detail and some additional notation, but, essentially, merely generalize the procedure we have been applying intuitively all along. A reader who has strong intuition but is impatient with detail may wish to skip to expression (IV.1) and Theorem I below.

the graph preserving restrictions of f from (t, h) converges to η .

We now explain how a strategy profile f is "moulded" to a grid R .²⁸ First, set $f_{|R}(0, \emptyset) = f(0, \emptyset)$. Now fix a positive-length history h . Recall that $f(\cdot, h)$ is a piecewise function defined on the interval $[t^{last}(h), 1)$. Since the range of f is finite, the function consists of a finite number of constant segments. The "moulded" function, $f_{|R}(\cdot, h)$, will be defined on the finite set $(t^{last}(h), 1) \cap R$. If R is sufficiently fine, the graph of $f_{|R}(\cdot, h)$ will be very similar to that of $f(\cdot, h)$, *except* that the *beginning* of each constant segment of $f(\cdot, h)$ will be shifted slightly to the right. We will now be more precise. For each point t in $[0, 1)$, let $\lfloor t \rfloor^R$ ($\lceil t \rceil^R$) denote its immediate strict predecessor (strict successor) in R . For each grid point r strictly exceeding the successor of $t^{last}(h)$ in R , define $f(r, h)$ as follows: If the interval $(\lfloor r \rfloor^R, r)$ contains one or more discontinuity points of $f(\cdot, h)$, then set $f_{|R}(r, h)$ equal to the value of f at the largest such point.²⁹ Otherwise, set $f_{|R}(r, h) = f(\lfloor r \rfloor^R, h)$. Now define $f_{|R}$ at $(\lceil t^{last}(h) \rceil^R, h)$ in exactly the same way, *except* that the relevant interval is $(t^{last}(h), \lceil t^{last}(h) \rceil^R)$ rather than of $(\lfloor r \rfloor^R, r)$. Note that if R has few grid points relative to the number of discontinuity points of $f(\cdot, h)$, then $f_{|R}(\cdot, h)$ may bear little relation to $f(\cdot, h)$. For any given h , however, we can choose a grid R sufficiently fine that the graphs of $f_{|R}(\cdot, h)$ and $f(\cdot, h)$ will be very similar.³⁰ For this reason, we call $f_{|R}$ the graph preserving restriction of f to R .

The following example illustrates the various aspects of our moulding procedure (see figure I). As usual, suppose there are two players, and each can choose between 'left' and 'right'. Fix the one-length history

$$\bar{h} = (0, lf\ lf) \text{ and define } f(\cdot, \bar{h}) \text{ as follows: } f(t, \bar{h}) = \begin{cases} rt\ lf & \text{if } t = 0 \\ lf\ rt & \text{if } t = 1/5 \text{ or } t = 1 \\ lf\ lf & \text{if } t \in (0, 1/5) \cup (3/5, 1) \\ rt\ rt & \text{if } t \in (1/5, 3/5] \end{cases} \quad \cdot \text{ Figure I illus-}$$

trates the graph preserving restrictions of this profile to two grids, $R^5 = \{0, 1/5, \dots, 4/5\}$ and $R^{10} = \{0, 1/10, \dots, 9/10\}$. The restrictions are:

²⁸ We will always assume that discrete-time grids contain the point zero.

²⁹ A largest discontinuity point exists because $f(r, h)$ is piecewise continuous.

³⁰ This is another consequence of piecewise continuity.


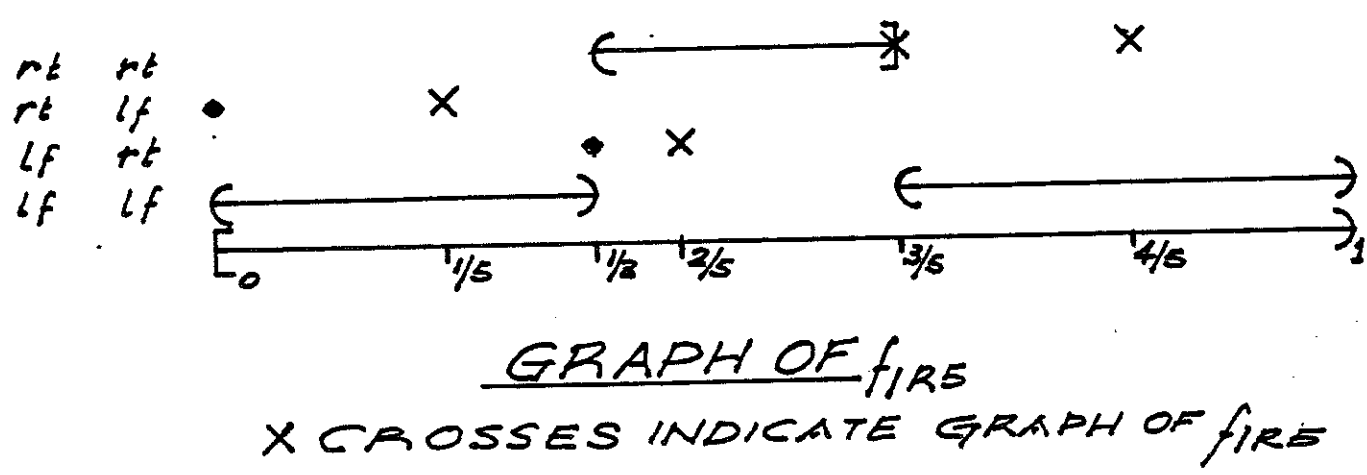



FIGURE 1

$$f_{|R^s}(r, \bar{h}) = \begin{cases} rtlf & \text{if } r = {}^1\mathcal{A} \\ lf rt & \text{if } r = {}^2\mathcal{A} \\ lf lf & \text{if } r \in \{{}^3\mathcal{A}, {}^4\mathcal{A}\} \end{cases} ; f_{|R^{10}}(r, \bar{h}) = \begin{cases} lf lf & \text{if } r = \{{}^1\mathcal{A}, {}^3\mathcal{A}_0, {}^4\mathcal{A}, {}^9\mathcal{A}_0\} \\ rt lf & \text{if } r = {}^1\mathcal{A}_0 \\ lf rt & \text{if } r = {}^2\mathcal{A} \\ rt rt & \text{if } r \in \{\frac{1}{2}, {}^3\mathcal{A}, {}^7\mathcal{A}_0\} \end{cases}$$

Note that neither restriction is defined at $(0, \bar{h})$, since 0 is the last jump-time of \bar{h} . The first grid is insufficiently fine to capture all of the details of $f(\cdot, \bar{h})$. In particular, the segments $(0, {}^1\mathcal{A})$ and $({}^3\mathcal{A}, 1)$ are not represented; on the other hand, the graphs of $f(\cdot, \bar{h})$ and $f_{|R^{10}}(\cdot, \bar{h})$ are very similar.

Summarizing, the graph preserving restriction of f to R is defined by, for $(r, h) \in DN$ such that $r \in R$ and strictly exceeds $t^{last}(h)$:

$$f_{|R}(r, h) = \begin{cases} f(0, \emptyset) & \text{if } (r, h) = (0, \emptyset) \\ f(t^{last}(h), h) & \text{if } r = \lceil t^{last}(h) \rceil^R \text{ and } f(\cdot, h) \text{ is continuous on } (\lceil t^{last}(h) \rceil^R, r) \\ f(\lfloor r \rfloor^R, h) & \text{if } r \in (\lceil t^{last}(h) \rceil^R, 1) \text{ and } f(\cdot, h) \text{ is continuous on } (\lfloor r \rfloor^R, r) \\ f(\max\{s < r: f(\cdot, h) \text{ is dis-} \\ \text{continuous at } s\}, h) & \text{otherwise} \end{cases}$$

Note that $f_{|R}$ is not an R -admissible strategy, because its domain is too big. Specifically, $f_{|R}(\cdot, h)$ is defined for *all possible* histories, not just R -admissible ones. This is immaterial, however. Clearly, $f_{|R}$ can be played from any continuous-time decision node and will generate a unique history by the inductive procedure described above (p. 19). For our purposes, this is all that matters.

To illustrate the process described above, we consider three profiles, f^a , f^b and f^c . The examples illustrate that certain differences between strategy profiles matter a great deal while others are unimportant. As usual, there are two players, each with two strategies, 'left' and 'right'. The first two profiles are identical, except that f^a is discontinuous from the left, while f^b is discontinuous from the right. As one would expect, the limit outcome is insensitive to this difference. In the third example, player #1's strategy has a discontinuity from the left, while #2's is discontinuous from the right. In this case, the difference between the discontinuities matters a great deal: the graph preserving restrictions of f^c preserve this distinction and the resulting outcome has player #1 jumping alone.

The first profile, f^a , is defined as follows: #1 begins by playing 'left', then switches to 'right' at $\frac{1}{2}$, regardless of the past history; #2 begins by playing 'left', but switches to 'right' if #1 ever plays 'right'.

$$f_1^a(t, h) = \begin{cases} lf & \text{if } t < 1/2 \\ rt & \text{if } t \geq 1/2 \end{cases} \quad f_2^a(t, h) = \begin{cases} rt & \text{if #1 has ever played } rt \\ lf & \text{otherwise} \end{cases}$$

Now consider the graph preserving restriction of f^a to a grid R . We have: $f_{1R}^a(0, \emptyset) = lf\ lf$ and

$$f_{1R}^a(r, (0, lf\ lf)) = \begin{cases} lf\ lf & \text{if } r \leq 1/2 \\ rt\ lf & \text{otherwise} \end{cases} . \text{ Also, for all } t \text{ and } r > t, f_{1R} \left(r, [(0, lf\ lf), (t, rt\ lf)] \right) = rt\ rt. \text{ Fi-}$$

nally, for all $r > t' \geq t \geq 1/2$, $f_{1R}^a \left(r, [(0, lf\ lf), (t, rt\ lf), (t', rt\ rt)] \right) = rt\ rt$. Applying the algorithm above,

the discrete-time outcome generated by f_{1R}^a from $(0, \emptyset)$ is: player #1 jumps at the first grid point strictly greater than $1/2$; player 2 follows suit at the next grid point; there are no further jumps. Now define f^b identically to f^a , except that f_1^b has player 1 shifting to 'right' immediately after $1/2$, i.e.,

$$f_1^b(t, h) = \begin{cases} lf & \text{if } t \leq 1/2 \\ rt & \text{if } t > 1/2 \end{cases} . \text{ The discrete-time outcome generated by } f_{1R}^b \text{ has player #1 jumping at the}$$

second grid point strictly greater than $1/2$, and player #2 following suit at the third. It is easy to see that the two profiles converge to the same limit, i.e., $[(0, lf\ lf), (1/2, rt\ lf), (1/2, rt\ rt)]$. Thus, player #1 jumps at $1/2$ and player #2 follows suit immediately afterwards.

Finally, we consider a third profile, f^c . The two players' strategies are identical, except that #1 shifts from 'left' to 'right' exactly at $1/2$, while #2 shifts immediately after $1/2$.

$$f_1^c(t, h) = \begin{cases} rt & \text{if } t \geq 1/2 \text{ and } a^1(h) = lf\ lf \\ lf & \text{otherwise} \end{cases} ; \quad f_2^c(t, h) = \begin{cases} rt & \text{if } t > 1/2 \text{ and } a^1(h) = lf\ lf \\ lf & \text{otherwise} \end{cases}$$

Our "moulding" procedure preserves the one difference between the two players' strategies. For any grid R , f_{1R} has player #1 jumping alone at the first grid point strictly later than $1/2$, and both jumping together beyond this time. The limit outcome is that player #1 jumps exactly at $1/2$; player #2 never jumps.

An explicit algorithm for computing continuous-time outcomes.

There is a simple inductive algorithm for determining limit histories directly, without ever constructing graph preserving restrictions or computing discrete-time outcomes. The algorithm closely parallels the conventional one that determines discrete-time outcomes, but induction plays a fundamentally different role. In discrete time, we do induction on the set of possible times that agents can move. In continuous time, we in-

duct on the set of times that they *choose* to move. Our restriction F1 ensures that this latter set is finite, and so induction is possible.

Given a profile, f , and a decision node, $(\underline{t}, \underline{h})$, the formula below identifies a unique history, $\bar{\eta}$, defined as follows. Its first $k(\underline{h})$ jumps correspond to those of \underline{h} . If f has agents playing $a^{last}(\underline{h})$ at every t between \underline{t} and the end of the game, then the specification of \bar{h} is complete. Otherwise, $t^{k(\underline{h})+1}(\bar{\eta})$ is defined as the infimum of the times beyond \underline{t} at which agents choose an action profile that differs from $a^{last}(\underline{h})$. If f has some player jumping at *exactly* $t^{k(\underline{h})+1}(\bar{\eta})$, then $a^{k(\underline{h})+1}(\bar{\eta})$ is the action profile chosen at this time, i.e., $f(t^{k(\underline{h})+1}(\bar{\eta}), \underline{h})$. Otherwise, $a^{k(\underline{h})+1}(\bar{\eta})$ is the profile chosen *immediately after* $t^{k(\underline{h})+1}(\bar{\eta})$, i.e., $\lim_{\delta \downarrow 0} f(t^{k(\underline{h})+1}(\bar{\eta}) + \delta, \underline{h})$. The remaining jumps of $\bar{\eta}$ are defined in the corresponding way. Summarizing, for each f and regular node $(\underline{t}, \underline{h})$, the history generated by f from $(\underline{t}, \underline{h})$ is the unique history, $\bar{\eta}$, identified by the following conditions:

$$\text{for } 1 < \kappa \leq k(\underline{h}), \quad (t^\kappa(\bar{\eta}), a^\kappa(\bar{\eta})) = (t^\kappa(\underline{h}), a^\kappa(\underline{h})) \quad (\text{IV.1})$$

for $\kappa > k(\underline{h})$, if $t^{\kappa-1}(\bar{\eta})$ has been defined and there exists $s > \max(\underline{t}, t^{\kappa-1}(\bar{\eta}))$ such that $f(s, \bar{\eta}_{|\kappa}) \neq a^\kappa(\bar{\eta})$

$$t^\kappa(\bar{\eta}) = \inf\{1 > s > \max(\underline{t}, t^{\kappa-1}(\bar{\eta})) : f(s, \bar{\eta}_{|\kappa-1}) \neq a^{\kappa-1}(\bar{\eta})\}$$

$$a^\kappa(\bar{\eta}) = \begin{cases} f(t^\kappa(\bar{\eta}), \bar{\eta}_{|\kappa-1}) & \text{if } f(t^\kappa(\bar{\eta}), \bar{\eta}_{|\kappa-1}) \neq a^{\kappa-1}(\bar{\eta}) \\ \lim_{\delta \downarrow 0} f(t^\kappa(\bar{\eta}) + \delta, \bar{\eta}_{|\kappa-1}) & \text{otherwise} \end{cases}$$

If $(\underline{t}, \underline{h})$ is the start of the game, the formula is identical to the one above except that $(t^1(\bar{\eta}), a^1(\bar{\eta})) = (0, f(0, \emptyset))$.

To illustrate the formula, we return to the two examples, f^a and f^b , discussed above (pp. 23-24) Let

$\bar{\eta}^a$ and $\bar{\eta}^b$ denote, respectively, the outcomes generated by f^a and f^b from $(0, \emptyset)$. Clearly, the two histories have the same first jumps: $(t^1(\bar{\eta}^a), a^1(\bar{\eta}^a)) = (t^1(\bar{\eta}^b), a^1(\bar{\eta}^b)) = (0, lf\ lf)$. The second jump-time is the infimum of the times beyond zero at which some player plays 'right'. In each case, this time is $t = 1/2$, so that $t^2(\bar{\eta}^a) = t^2(\bar{\eta}^b) = 1/2$. Now f^a has some player jumping exactly at $1/2$, so that $a^2(\bar{\eta}^a) = f^a(1/2, \bar{\eta}_{|1}^a) = rt\ lf$. On the other hand f^b has both players continuing to play 'left' at $1/2$, so that $a^2(\bar{\eta}^b)$ is determined by their choices immediately after $1/2$. Once again, this limit choice is $rt\ lf$, i.e.,

$a^2(\bar{\eta}^b) = \lim_{\delta \downarrow 0} f^b(\frac{1}{2} + \delta, \bar{\eta}_{11}^b) = rt \text{ lf}$. Finally, consider the third jumps. Both f^a and f^b have player #2 following suit immediately after #1's jump, so that $(t^3(\bar{\eta}^a), a^3(\bar{\eta}^a)) = (t^3(\bar{\eta}^b), a^3(\bar{\eta}^b)) = (\frac{1}{2}, rt \text{ rt})$.

We have now completed the program set out on p. 4 for specifying our outcome function. Our first theorem verifies that the history identified by our explicit construction (IV.1) is indeed the limit of the discrete-time histories it induces. Say that a discrete-time grid is δ -fine if it contains zero, its largest member exceeds $1 - \delta$ and the largest distance between any two adjacent grid points is at most δ . Theorem I states that for any profile and any decision node, the history defined by (IV.1) is the limit (in the metric d^H) of the histories generated by playing the graph-preserving restrictions of the profile, starting from the given decision node, on any sequence of increasingly fine grids.

Th'm I: Let f be a profile satisfying F1-F3. From any decision node $(\underline{t}, \underline{h})$ and for $\varepsilon > 0$, there exists δ such that if R is a δ -fine grid, then the discrete-time history generated by $f|_R$ from $(\underline{t}, \underline{h})$ is within ε of the history generated by f from $(\underline{t}, \underline{h})$.

Valuation functions, payoff functions and equilibrium notions.

The valuation function, $V = (V_i)_{i \in I}$, for a game assigns to each history a vector of payoffs. $V_i(h)$ is the value player i assigns to the history h . In most of the applications we consider, we will assume that $V(\cdot)$ is uniformly continuous with respect to d^H . At first sight, this seems like a strong assumption, because most games of timing are thought of as highly discontinuous. In fact the assumption is extremely weak, because the topology induced by d^H is very fine (see p. 17).

In many games, the valuation function is obtained by integrating with respect to time some instantaneous flow payoff matrix, $u: [0, 1) \times A \rightarrow \mathbf{R}^H$, which depends only on time and the current state of

the system. To define payoffs in such cases, it is convenient to have an alternative representation of histories that can be integrated. We will define the time path of h to be the function ρ^h , defined from h by: $\rho^h(t)$ is be the last action chosen weakly before t . For example, the time path of the history

$h = [(0, \text{lf lf}), (\frac{1}{2}, \text{rt lf}), (\frac{1}{2}, \text{rt rt})]$ would be ρ^h , where $\rho^h(t) = \begin{cases} \text{lf lf} & \text{if } t < \frac{1}{2} \\ \text{rt rt} & \text{if } t \geq \frac{1}{2} \end{cases}$. This representation

suppresses the information that player #1 jumped before #2. With integral payoffs, however, this information is irrelevant. Now, if the instantaneous payoff matrix is u , the value to i of the history h is

$V_i(h) = \int_{[0,1]} u_i(s, \rho^h(s)) ds$. For illustrations, see the following section. It is straightforward to show that whenever $u(\cdot, a)$ is an integrable function, for every a , the derived valuation function will be uniformly continuous in the metric d^H .

The continuous-time payoff function, $P = (P_i)_{i \in I}$, assigns a payoff vector to each strategy profile and decision node. The payoff function is derived from the valuation function in the obvious way: if η is the outcome generated by f from (t, h) , then $P_i(f, t, h) = V_i(\eta)$, i.e., player i 's payoff if agents play f from the subgame beginning at (t, h) .

We will say that strategy profile is an ϵ -best reply from a decision node, if in the subgame starting from this node, no agent can gain more than ϵ by deviating from his part of the profile. That is, f is an ϵ -best reply from (t, h) if for all i , and all f'_i , $P_i(f, t, h) \geq P_i((f'_i, f_{-i}), t, h) - \epsilon$. A profile f is an ϵ -subgame perfect equilibrium (ϵ -SGP equilibrium) if it is an ϵ best reply from every decision node. Finally, f is a subgame perfect equilibrium if it is an ϵ -SGP equilibrium, for every $\epsilon > 0$.

For certain kinds of games, the set of SGP equilibria may be extremely large. (We saw one example in section II. The following section contains another.) In some instances, this set can be reduced considerably by iterative elimination of dominated strategies. To formalize this idea, we define the notion of an "iteratively undominated equilibrium" (see fn. 8). Let F_i^0 denote player i 's set of admissible strategies. Let $F_{-i}^0 = \prod_{j \neq i} F_j^0$. Say that f_i is zero-th order undominated if $f_i \in F_i^0$. Now suppose that for each agent, the set of $(k-1)$ -th order undominated strategies has been defined and is denoted by F_i^{k-1} . We will say that a strategy f_i for i is k -th order undominated if there exists no $(k-1)$ -th order undominated strategy for i that is at least as good as f_i against any $(k-1)$ -th order undominated strategies by the other players and strictly better against at least one. Let F_i^k denote the k -th order undominated strategies for i . In symbols,

$$F_i^k = \{f_i \in F_i^{k-1} : \text{there exists no } f'_i \in F_i^{k-1} \text{ satisfying: for all } f_{-i} \in F_{-i}^{k-1}, \text{ and all } (t, h) \in DN \\ P_i((f'_i, f_{-i}), t, h) \geq P_i((f_i, f_{-i}), t, h), \text{ with strict inequality holding for some } f_{-i}\}.$$

We now say that a profile f is an iteratively undominated equilibrium if it is a subgame perfect equilibrium

and survives an infinite number of rounds of iterated elimination, i.e., if for all i , $f_i \in \bigcap_{k=0}^{\infty} F_i^k$.

V. Applications.

This section considers two applications.³¹ The first is a stylized model of a deterministic race, in the spirit of Fudenberg et. al. [3] and Harris-Vickers [9] (see also [8]). The second is a "one-sided matching game": at least one agent strictly prefers to match the action being played by the other. We believe that these applications are interesting in their own right. Their main purpose, however, is to illustrate various aspects of our model. Both applications illustrate the usefulness of a framework that allows zero reaction lags. In the first example, this is simply convenient from a computational standpoint. In the second, the consequence is more far reaching. It enables us to isolate the cooperative outcome as the *unique* solution for the game.

Brinkmanship in a deterministic race.

We consider a race run over the interval $[0, 1)$. A prize of π^* is won by the first player to accumulate a certain stock of "knowledge." We assume π^* is *not* an integer and that $\pi^* > 2$.³² If no agent completes the course in the allotted time, the prize is forfeited; in the event of a tie, the prize is divided evenly among the winners. Participants in the race choose over time a *integer* level of investment. Knowledge is accumulated at a rate proportional to the agent's level of investment. Investment levels can be upgraded, but not downgraded.³³ If an investment level of n is chosen, a lump sum cost of $\$n$ is incurred. The cost of upgrading from level n to n' is $\$(n' - n)$. Agent i is endowed with $\$i$ at the start of the game. Borrowing is prohibited, so that i 's maximum possible investment level is i .

We now formulate this problem as a continuous-time game. There is a set of agents, $I = \{1, \dots, \bar{i}\}$. Player i 's action set A_i is $\{0, \dots, i\}$. At any decision node, (t, h) , i 's progress up to t can be calculated from i 's component of the time path of h , i.e., from ρ_i^h .³⁴ In particular, i 's current stock level is $\int_0^t \rho_i^h(s) ds$. (For example, if i invests at level 1 at time zero, and upgrades to 2 at $1/3$, then by time $1/2$ his accumulated

³¹ For more applications, involving behavior strategies, see [24], [21] and [22].

³² The first assumption avoids complications arising from indifference. The second ensures that ties will not occur in equilibrium.

³³ This assumption ensures that our finite move constraint (F1) will be satisfied.

³⁴ Recall that ρ^h was the alternative representation of h , defined on pp. 26-27. The information lost in the conversion from the history h to ρ^h --i.e., instantaneous sequences of upgrades--is not relevant for computing the agent's progress.

stock of knowledge is $2/s$.) For any h , let $\tau_i(h)$ denote the time--if any--at which player i crosses the finish line, i.e., $\tau_i(h) = \inf\{t > 0: \int_0^t \rho_i^h(s) ds \geq 1\}$. Let $I(h)$ the set of winners of the race, if the history is h , i.e., $I(h)$ contains the j 's whose $\tau_j(h)$'s are minimal. The game is now summarized by the valuation function, V , below. Player i is disqualified from the race if ρ_i^h ever decreases. Disqualified players earn $-\infty$. Otherwise, i 's payoff is his share of the prize, if he is a winner, minus his terminal investment level. Thus, V_i is defined

$$\text{by } V_i(h) = \begin{cases} (\#I(h))^{-1} \pi^* - a_i^{\text{ast}}(h) & \text{if } h_i \text{ is nondecreasing and } i \in I(h) \\ -a_i^{\text{ast}}(h) & \text{if } h_i \text{ is nondecreasing and } i \notin I(h). \\ -\infty & \text{otherwise} \end{cases}$$

To completely specify a solution to this problem, we would need to introduce behavior strategies. Randomization is, however, *necessary* for existence only at decision nodes that are so far away from the equilibrium path to be "irrelevant."³⁵ There are, therefore, equilibria in which agents play pure strategies along the equilibrium path. We will characterize these equilibria in an informal way.

Apart from its continuous-time setting, our model differs in several respects from those in [9] and [3]: there are many agents instead of two; investment is irreversible; investment costs are lump-sum rather than flows; agents have different resources available for investment.³⁶ As a result of these differences, our game has some interesting properties that, we believe, have not been discussed in the patent race literature. There are two kinds of equilibria. In the first kind--which we call preemptive equilibria--some agent *other than* \bar{i} invests at the outset at a level so high that no other agent can profitably compete. In these equilibria, the rents to winning are almost entirely offset by the cost of preemption. In addition, there is one more solution, which we call a brinkmanship equilibrium. In this solution, player \bar{i} maintains the lowest possible level of investment throughout the race. Rents are barely dissipated at all. In each kind of equilibrium, the outcome is completely determined by the actions agents take at the start of the game.

³⁵ Like Fudenberg et. al. [3], we need to introduce randomization only at nodes where agents have, in a certain sense, equal chances to win. Since our agents are unevenly matched at the outset, these nodes can be reached only if the stronger agents fall behind a little way, but not far enough that they cannot profitably catch up. The nodes of this kind are bounded away any equilibrium path for the game.

³⁶ [9] have a related assumption, but in [3], the competitors are identical.

The equilibria are discussed in detail in the Appendix. The following is a brief summary. Only \bar{i} , the strongest agent, can win the race without committing significant resources to it. This is because \bar{i} is the only player with the resources to sustain a fight in which the stakes of the game escalate. The irreversibility of investment plays an important role here: because "sunk costs are sunk," \bar{i} is willing to match any investment that any other agent is willing to make. If \bar{i} sinks \$1 into the project at time zero, he will maintain his competitive edge if he merely matches the expenditure of his closest competitor. (At the beginning of the game, no stocks have been accumulated, so the agent with the "edge" is the one who has already committed to the highest rate of investment.) Moreover, no challenger will be willing to spend more than $[\pi^*]$ to challenge \bar{i} ,³⁷ while \bar{i} is always willing to spend up to $[\pi^*]$ to maintain his lead, regardless of the size of his sunk costs. For this reason, all other agents are deterred from challenging, once \bar{i} has a minimal "edge" in the race.

For example, suppose that $I = \{1, 2, 3\}$ and $\pi^* = 2.5$. This game has two equilibria. In the preemptive equilibrium, player #2 invests 2 units at the outset and maintains this level throughout. Player #3 can either force a tie or win, but in either case, the cost of doing so would exceed the benefit. In the brinkmanship equilibrium, #3 invests one unit. Player #2 will certainly not be willing to spend more than \$2 on the race; if #2 challenges by investing either \$1 or \$2, #3 will immediately match this expenditure and maintain his lead. Knowing this, #2 refrains from challenging.

To conclude this discussion, we characterize the equilibrium set for this game and partially describe the equilibrium strategies. (What is missing is a description of what happens off the equilibrium path, if the designated leader has fallen behind.)

³⁷ We will denote by $[\pi^*]$ the largest integer amount less than π .

There is a brinkmanship equilibrium in which \bar{i} wins the race. The equilibrium has the following properties. At the start of the game, \bar{i} chooses an investment level of one and no other agent invests. Along the equilibrium path, no further investment occurs by any agent. Off the equilibrium path, \bar{i} upgrades in response to a challenge by another player if and only if the upgrade is (a) minimal (i.e., no smaller response would suffice); (b) feasible (i.e., within \bar{i} 's budget constraint); (c) profitable (i.e., yields \bar{i} a higher payoff than he would achieve by not responding). In addition to this equilibrium, there is for every $i \in [[\pi^*], \bar{i})$, a preemption equilibrium in which i invests at level $[\pi^*]$ at the beginning of the game, and maintains this level throughout. The other properties of preemption equilibria are similar to those of the brinkmanship equilibrium. There are no other equilibria. (V.1)

A 'one-sided' coordination game.

Our second application is related to the one in section II. As in that example, we iteratively eliminate weakly dominated strategies to obtain a unique equilibrium. In this case, however, the elimination process is more extensive and delicate.

The example is a two person game in which payoffs are obtained by integrating a time-independent flow payoff matrix. We will assume that there is some exogenously fixed number n^* such that neither agent can change his action more than n^* times.³⁸ Each player has two actions, 'good' (gd) and 'bad' (bd). Player #1 strictly prefers to be playing whatever #2 is playing. Player #2 strictly prefers to match #1, if #1 is playing 'bad'. Each player strictly prefers the state ' $gd\ gd$ ' to ' $bd\ bd$ '. When played in continuous-time, any game in this class has multiple subgame perfect equilibria. In discrete-time, the "bad" outcome is always an equilibrium; there may or may not be others. We will establish, however, that in continuous time, the Pareto superior outcome is the *unique* equilibrium that survives iterated elimination of weakly dominated strategies.

The two flow payoff matrices below illustrate the qualitative features of the game.

³⁸ One could endogenize this bound by supposing that each move is costly, and that agents have finite resources. This gain in realism complicates the problem but adds no new insights.

	'good'	'bad'
'good'	(3, 3)	(0, 7)
'bad'	(2, 0)	(1, 1)

	'good'	'bad'
'good'	(3, 3)	(0, 0)
'bad'	(0, 0)	(1, 1)

The right-hand matrix is an example of a class of games that has recently been the focus of considerable attention. Games like this, in which agents' payoffs are identical, are known as "pure coordination" games. Since agents' interests coincide exactly, one would expect that in such games, noncooperative behavior by individuals would *always* lead them to the Pareto superior outcome. Traditional models have been unable to predict this. In particular, both the discrete- and continuous-time versions of the game defined by the right-hand matrix have a continuum of subgame perfect equilibria. For example, for any subinterval T of $[0, 1)$, there is an equilibrium in which both agents play 'bad', when $t \in T$ and otherwise play 'good'.

It has proved surprisingly difficult to construct alternative models and solution concepts that yield unique predictions in such games. The first such results has only recently been obtained:³⁹ in particular, Aumann-Sorin [2] show that if a two-person pure coordination game is "perturbed" in an appropriate way, the modified game has a pure strategy equilibrium that is close (in payoffs) to the cooperative outcome.⁴⁰ Our uniqueness result is therefore striking, since for the class of games we consider, it may be that the cooperative outcome cannot even be sustained as a one-shot equilibrium.

We now formulate the problem as a continuous-time game. Each player is allowed to change his action no more than n^* times, where $n^* \in \mathbf{N}$. There is a time-independent flow payoff matrix for the game, u , satisfying the following inequalities: (i) for each i , $u_i(gd\ gd) > u_i(bd\ bd)$, (ii) $u_1(gd\ gd) > u_1(bd\ gd)$ and $u_1(bd\ bd) > u_1(gd\ bd)$; (iii) $u_2(bd\ bd) > u_2(bd\ gd)$. Observe that "bd bd" is a one-shot Nash equilibrium, while "gd gd" may or may not be, depending on whether or not $u_2(gd\ gd)$ weakly exceeds $u_2(gd\ bd)$. The

valuation function V is defined by
$$V(h) = \begin{cases} \int_{[0,1)} u(\rho_i^h(s)) ds & \text{if } i \text{ has jumped no more than } n^* \text{ times} \\ -\infty & \text{otherwise} \end{cases}$$

³⁹ Kreps et al. [12] obtain uniqueness in the repeated Prisoners' Dilemma by allowing for a small possibility that some agent is a compulsive tit-for-tat player. On the other hand, Fudenberg Maskin [6] show that when the model is expanded to allow for arbitrary kinds of compulsiveness, then virtually anything can happen.

(As usual, ρ^h is the time path of the history h .)

Prop'n II: The unique payoff vector that can be implemented by an Iteratively Undominated equilibrium for a game in the class described above is $u(gd\ gd)$.

We now give a brief sketch of the proof. (Details are deferred to the Appendix.) The idea is to prove that if players start out the game by playing anything other than ' $gd\ gd$ ', then after iterated elimination, the only surviving strategies have them cycling from state to state until they eventually settle at the state ' $gd\ gd$ '.⁴¹ Moreover, the cycle must occur in zero time!

As usual, we begin at terminal subgames of the game and work forwards. First suppose that #1 has one move remaining and #2 has none. If #1 is not already matching #2's action then, obviously, he should switch immediately. Now suppose that both players have one move remaining and are currently matching each other. If both are playing ' bad ', then #2 *must* lead #1 out of the bad state by switching to ' $good$ ', because he *knows* that #1 *must* follow suit. Since #1 can respond instantly to #2, there will be no interval of time during which the players will be mismatched. Therefore, the cost to #2 of leading in this way is zero.

The next step of the argument is more delicate. Assume that #1 has two moves remaining and is playing ' bad ', while #2 has one move left and is playing ' $good$ '. We need to show that #1 must switch immediately to ' $good$ '. The problem here is that #2 might switch to ' bad ' at the very instant that #1 switches. Nevertheless, we can show that #1 must indeed switch. We can then proceed by induction.

Note that this argument depends critically on the fact that reaction lags are literally nonexistent. In discrete-time, the argument just given would break down. Certain anti-cooperative strategies will be "almost weakly dominated," but, as we pointed out in section II, this is not a sufficient reason to eliminate them.

VI. Equilibria of discrete- and continuous-time games.

This section studies the relationship between our model and the conventional discrete-time one. We first say what it means for a discrete-time profile to approximate a continuous-time one.⁴² We then study the

⁴⁰ There is a very easy version of the argument given below that yields uniqueness in the model considered by Aumann-Sorin [2], provided that agents can move only finitely many times. Aumann-Sorin's result, however, does not depend on finiteness.

⁴¹ Agents may cycle through a large number of steps before settling down, or they might settle down immediately. However, cycles that are longer than necessary would not occur in equilibrium if movement costs were introduced (see fn. 38).

⁴² For a parallel discussion of the relationship between finite- and infinite-horizon games, see Fudenberg-Levine [4], [5] and Harris [10]. [5] is also compares discrete- and continuous time. See also Hendriks-Wilson [11] for a comparison of discrete- and continuous-time equilibria in a very different context.

continuity properties of the SGP equilibrium correspondence. We emphasize that this section says nothing about *iteratively undominated equilibria*. As we have emphasized repeatedly, the equilibria are very special to continuous time and have no discrete-time analogs.

There are two general kinds of questions that are of interest. Is discrete time with a very fine grid a good proxy for continuous time? Conversely, do the equilibria for our continuous-time model have discrete-time analogs? The first question has arisen frequently in the literature on dynamic games. In the absence of an established continuous-time framework, many authors have studied sequences of discrete-time equilibria, allowing the grids to become increasingly fine, and have interpreted the limit of such equilibria as proxies for the equilibria of an unspecified continuous time version of the game.⁴³ With a continuous-time model and a notion of closeness for strategies in place, we can pose the question precisely. If a sequence of ϵ^n -equilibria increasingly closely approximates a continuous time profile, f , with $\epsilon^n \rightarrow 0$ as the period length shrinks, will f be an equilibrium? The answer to this question is "not in general." There is, however, a positive qualification that can be added. If the discrete-time profiles happen to be the graph-preserving restrictions of f , then f will be an equilibrium. We will state this formally as Theorem III.

From our particular perspective, the second question--do our continuous-time equilibria have discrete-time analogs?--is of more direct concern. The answer is "yes, provided that agents' payoffs are insensitive to the actions other agents choose near the end of the game."⁴⁴ When this condition is satisfied, we have a compelling, if conservative, validation of the equilibria of our model: they do not depend intrinsically on the special properties of continuous time. When it fails, continuous time is *intrinsically* different from discrete time. Once again, the source of this difference is that discrete time is well-ordered, while continuous time is not.⁴⁵ (See below for an illustration.)

⁴³ Among the best known of these studies are Rubinstein [19] and Kreps-Wilson [14].

⁴⁴ This is the discrete- to continuous-time counterpart of the now familiar continuity results linking finite- and infinite-horizon discrete time games. See Radner [18] and Fudenberg-Levine [4], [5].

⁴⁵ This intrinsic difference becomes much more apparent once randomization is introduced. For example, Simon [21] constructs a continuous-time SGP equilibrium in behavior strategies that is far away from any approximate equilibrium of any nearby discrete-time game. If we are to base predictions on equilibria of this kind, we must believe that time in reality is better described by the continuum than by a finite set.

We now define what it means for a discrete-time strategy profile to approximate a continuous-time one.⁴⁶ We will say that an R -admissible profile g ε -approximates a continuous-time profile f if for each R -admissible decision node (r, h) such that $r < 1 - \varepsilon$, there exists a *nearby* continuous-time decision node such that the outcomes generated by g and f from the respective decision nodes are close to each other. Two points about this notion need to be highlighted. First, we add the caveat " $r < 1 - \varepsilon$ " because in general it will not be possible for a discrete-time strategy to closely approximate a continuous-time strategy at the very end of the game. The reason is that any discrete time grid must have a second-to-last, third-to-last, ..., grid point, but there are no corresponding last points "at the end" of the continuum. Second, we have not defined a *metric* on profiles, but rather a "one-way" notion of closeness. This is because the domain of a continuous-time profile is so much more complex than that of a discrete-time one. Therefore, in general, it will be difficult if not impossible to represent all of the strategic detail of a continuous-time profile on a given discrete-time grid. We now state the definition precisely. Let $o^{f, dn}$ denote the outcome generated by f from the decision-node ' dn '. Given a discrete-time grid R , we say:

an R -admissible strategy profile g ε -approximates a continuous-time strategy profile f if for every R -admissible decision node (r, h) such that $r < 1 - \varepsilon$, there exists $(t, \eta) \in DN$ such that (i) $|t - r| + d^H(h, \eta) < \varepsilon$ and (ii) $d^H(o^{g, r, h}, \eta^{f, t, \eta}) < \varepsilon$

The only node close to $(0, \emptyset)$, in the sense of (i) above, is $(0, \emptyset)$ itself. Therefore, g will approximate f only if the two profiles generate similar outcomes from the start of the game.

We now return to our first question: is discrete-time a good proxy for continuous-time? The example below establishes that the answer is "not in general." There are three players. Each has two actions, '*continue*' (ct) and '*terminate*' (tm). Players #2 and #3 are completely indifferent about the outcome of the game.

Player #1's payoff depends only on the terminal state of the game. His payoff is

$$V_1(h) = \begin{cases} 1 & \text{if } a_1^{last}(h) = a_3^{last}(h) = tm \\ 0 & \text{otherwise} \end{cases} . \quad \text{An exact equilibrium for the discrete-time game played on } R^n \text{ is}$$

the strategy g^n , defined as follows:

⁴⁶ We emphasize that our notion of closeness is defined for *profiles* not for individual strategies. This is because metrics defined on individual agents' strategy sets tend to have very poor continuity properties. In discrete-time game theory, this fact is well-known. See Fudenberg-Levine [4], [5] and Harris [10].

"player #1 never terminates; player #2 terminates iff $t = \frac{1}{2} - \frac{1}{n}$ and nobody has yet terminated; player #3 terminates, iff both #1 and #2 have already terminated and both of them terminated at some time weakly after $\frac{1}{2}$."

The profile (VI.1.n) is clearly an exact equilibrium for the game played on R^n , since no unilateral deviation by player #1 will induce player #3 to terminate. Moreover, g^n clearly ϵ^n -approximates the limit of the (VI.1.n)'s. In this limit, player #2 terminates exactly at $\frac{1}{2}$. The limit is not an equilibrium, however. If #1 deviates and terminates simultaneously with #2 at $\frac{1}{2}$, #3 will terminate also and #1's payoff will increase. Summarizing, even though player #1's payoff function is continuous with respect to histories, player #3's strategy is not (though it is strongly right continuous w.r.t. histories). Because #3 reacts discontinuously to other players' jump-times, #1's strategic opportunities vary discontinuously with the time at which #2 jumps.⁴⁷

In this example, #2's strategy "converges" from below (i.e., his jump-times are strictly increasing in the passage to the limit). On the other hand, a consequence of strong right continuity (F3) is that if a continuous-time strategy can be approximated "from above" by discrete-time approximate equilibrium strategies, then the limit strategy will be an equilibrium. In particular, the discrete-time strategies defined by the graph-preserving restrictions of f do converge "from above" to f .⁴⁸ We have the following result:

Th'm III: Consider a continuous-time game with a d^H -continuous valuation function. Let f be a continuous-time strategy profile satisfying F1-F3. Suppose that there exists a sequence of δ^n -fine grids, (R^n) , where $\delta^n \rightarrow 0$ and a sequence (g^n, ϵ^n) such that $\epsilon^n \rightarrow 0$ and for each n , g^n is an ϵ^n -SGP equilibrium for the game played on R^n . Further suppose that g^n is defined by further restricting $f|_{R^n}$ to the R -admissible decision nodes. Then f is an SGP equilibrium for the continuous time game.

We now come to the main result of the section, which is our lower hemi-continuity result. For this, we need a further restriction on the game and an additional assumption on payoffs. We assume there is an upper bound, that is *uniform over strategies*, on the number of jumps that agents can make.⁴⁹ Our assumption on

⁴⁷ This failure of upper hemi-continuity is entirely consistent with our experience from discrete-time game theory. It is well known that topologies on discrete-time strategies must be essentially discrete to guarantee upper hemi-continuity. See Fudenberg-Levine [4], [5] and Harris [10].

⁴⁸ Recall that graph preserving restrictions are not actually discrete-time strategies, because their domains are too large. However, when $f|_{R^n}$ is further restricted to the R -admissible decision nodes, the resulting profile is R -admissible.

⁴⁹ This is stronger than our assumption F1, which allowed the upper bound on moves to vary with the strategy. We could easily engender the restriction by imposing suitable conditions on payoffs: for example, a cost of moving plus a budget constraint.

agents' payoffs is that they are not too sensitive to other agents' actions at the very end of the game. That is, we assume that player i will assign similar values to two histories if the only difference between them is that other players' actions differ at the very end of the game.⁵⁰ Formally, we say that two histories h and h' agree before s if for every κ such that $t^\kappa(h) < s$, $(t^\kappa(h), a^\kappa(h)) = (t^\kappa(h'), a^\kappa(h'))$. Now set $h_i = (t_i(h), a_i(h))$, where $t_i(h)$ is the vector of times at which i moves and $a_i(h)$ is the corresponding vector of i 's jump destinations. Our assumption is

$$\text{For all } \varepsilon, \text{ there exists } \delta > 0 \text{ such that if two histories } h \text{ and } h' \text{ agree before } 1 - \delta \text{ and } h_i = h'_i, \text{ then } |V_i(h) - V_i(h')| < \varepsilon. \quad (*)$$

If V_i is determined by integrating an instantaneous flow payoff matrix then (*) will automatically be satisfied. In general, however, the restriction is nontrivial. For example, it may well fail if i 's valuation depends significantly on the terminal action taken by some other agent.

The following example illustrates why lower hemi-continuity may fail if (*) is not satisfied. There are two players; each has two actions, 'cooperate' (cp) and 'defect' (df). The game is like a repeated prisoners' dilemma, except that payoffs depend only on the last action that players choose. Specifically, the valuation

$$\text{function for the game is } V(h) = \begin{cases} (1, 1) & \text{if } a^{last}(h) = cp \ cp \\ (2, -2) & \text{if } a^{last}(h) = df \ cp \\ (-2, 2) & \text{if } a^{last}(h) = cp \ df \\ (0, 0) & \text{if } a^{last}(h) = df \ df \end{cases} \text{ This function clearly fails } (*). \text{ In}$$

discrete-time, cooperation cannot be sustained even as an approximate equilibrium, because each agent has a large incentive to defect at the last period of the game. In continuous time, however, there is no "last" period in which an agent can defect.⁵¹ The outcome "cooperate forever" can be sustained as an SGP equilibrium by the threat of immediate retaliation. We now state the result.

⁵⁰ This condition is related to the assumption that Fudenberg-Levine [4] invoke to obtain the analogous relationship between finite- and infinite-horizon games. Our assumption is strictly weaker, since they require that agents be approximately indifferent between any two histories that differ only at the end of the game.

⁵¹ In a two-person game played on a closed time interval, (*) would not be needed to obtain lower hemi-continuity. With three or more players, however, (*) is once again required. In a three player game, the source of the problem will be that there is no second-to-last period before the end of the game.

Th'm IV: Consider a continuous-time game in which payoffs are d^H -uniformly continuous and satisfy assumption (*). Assume there exists n^* such that no agent can move more than n^* times. If f be a continuous-time SGP equilibrium for this game, then for all $\epsilon > 0$, there exists $\delta > 0$ such that for every δ -fine grid R , there exists an ϵ -SGP equilibrium for the game played on R that ϵ -approximates f .

We begin our discussion of this result with an example that shows why it is difficult to prove. The example also illustrates how we go about proving it. We then discuss the proof in more detail. The point of the example is that we cannot obtain our approximating sequence merely by restricting the original profile to the sequence of discrete-time grids.⁵² The reason is that in the process of restricting the profile, we may significantly distort the strategic implications of the strategies.

The example has two players. Each two possible actions, 'continue' (ct) and 'terminate' (tm). If a player ever once plays 'terminate', he must do so for the remainder of the game. Payoffs are obtained by integrating the following instantaneous flow payoff matrix:

	'continue'	'terminate'
'continue'	(0, 0)	(0, $4t - 2$)
'terminate'	(1, $4t - 3$)	(-1, 0)

First consider #1's problem in this game. If he could preempt #2, terminate before $\frac{1}{2}$ and ensure that #2 would *never* follow suit, then he would certainly choose to do so. However, #1 would rather continue forever than terminate and have #2 terminate also. Now consider player #2. From his point of view, the best possible outcome in this game is: #2 terminates at time $\frac{1}{2}$, and #1 continues forever. Notice that so long as #2 can deter #1 from entering before $\frac{1}{2}$, he can attain this outcome, since once #2 terminates, #1 has a strong incentive *not* to follow suit. Moreover, #2 can indeed deter #1, by threatening to terminate immediately after #1 does. This threat is credible so long as $t \leq \frac{1}{2}$, because at this stage of the game, the short-run gains to terminating offset the long term loss. We have verified, therefore, that the following strategies form an equilibrium for the continuous-time version of this game.

⁵² It does not help to take the *graph-preserving* restriction.

Player #1: "play 'continue' if #1 has always played 'continue' in the past; otherwise play 'terminate'." (VI.2-1)

Player #2: "play 'continue' at time zero; at every positive time t , play 'continue' if either $t < \frac{1}{2}$ and no agent has yet terminated or $t \geq \frac{1}{2}$ and #1 has already terminated; otherwise play 'terminate'." (VI.2-2)

The outcome generated by this profile is that player #2 jumps at $t = \frac{1}{2}$, while player #1 never jumps. Player #1's payoff is zero, while #2's is $\frac{1}{2}$.

Now suppose that these strategies are simply restricted to a discrete-time grid. We will show that the strategic opportunities available to #1 are significantly changed. If #1 terminates at the last grid point in R before $\frac{1}{2}$, #2's first opportunity to react to this deviation will not occur until weakly beyond $\frac{1}{2}$. By this time, however, his strategy instructs him to continue. Thus, there is a deviation available to #1 against #2's restricted strategy that yields #1 a payoff exceeding $\frac{1}{2}$. The problem here is that when #2's strategy is restricted, a gap is opened up in his defenses, through which #1 can slip and escape unpunished! Therefore the restricted profile is not even an approximate equilibrium for the discrete-time game.

There is, however, an obvious way to perturb the restricted profile that leaves #2's defenses intact without seriously affecting his credibility. Observe that if #1 terminates just before $\frac{1}{2}$, #2 is *approximately* indifferent between terminating immediately and continuing forever. We can, therefore, slightly extend the period over which #2 threatens to terminate in response to an early termination by #1, and still have an approximate equilibrium. Specifically, the restriction of #1's strategy above, together with the following strategy for #2, form an approximate equilibrium for any discrete-time game, if the grid is sufficiently fine.

Player #2: "Play 'continue' at time zero; at every positive time t , play 'continue' if either $t < \frac{1}{2}$ and no agent has yet terminated or t strictly exceeds the first grid point beyond $\frac{1}{2}$, and #1 has already terminated; otherwise play 'terminate'." (VI.2-2')

The only difference between (VI.2-2) and (6.2-2') is that " $t \geq \frac{1}{2}$ " is replaced by " t strictly exceeds the first grid point beyond $\frac{1}{2}$." If #1 terminates at *any* time before $\frac{1}{2}$ on any discrete-time grid, this strategy has #2 responding by terminating at the next grid point. The strategic flavor of the original profile is thus restored, and the original outcome is implemented as an approximate equilibrium.

The above example illustrates our constructive technique for proving the theorem. We now explain the algorithm that we use for the general case. Fix a continuous-time SGP equilibrium profile f and a fine discrete-time grid R . We will "build" a profile g that ϵ -approximates f and is an approximate equilibrium

for the corresponding game on R . First, let $\bar{\eta}$ denote the outcome generated by f from the start of the game. Let \bar{h} denote the closest R -admissible history to $\bar{\eta}$ whose vector of jump-times exceeds $t(\bar{\eta})$. (Note in particular that if $\bar{\eta}$ has jumps that occur at the same instant of time, the corresponding jumps of \bar{h} will be spread out over consecutive grid points of R .) Our first step in the definition of g ensures it generates \bar{h} from the start of the game. For each r , let $\bar{h}_{\parallel r}$ denote the history defined by truncating \bar{h} after the last jump that occurs strictly before r . Now set $g(r, \bar{h}_{\parallel r}) = \rho^{\bar{h}}(r)$. Clearly, the outcome defined by g from $(0, \emptyset)$ is indeed \bar{h} .

Now suppose that agent i deviates from g_i by playing $a_i \neq g_i(r, \bar{h}_{\parallel r})$ at the node $(r, \bar{h}_{\parallel r})$. We need to ensure that this deviation is no more than slightly profitable. Let $\bar{\eta}_{\parallel r}$ denote the (unique) truncation of $\bar{\eta}$ that is close to $\bar{h}_{\parallel r}$. Find a continuous time node $(t, \bar{\eta}_{\parallel r})$ "close to" $(r, \bar{h}_{\parallel r})$ and let η^{dev} denote the outcome that would result if i played $a_i \neq f_i(t, \bar{\eta}_{\parallel r})$ at $(t, \bar{\eta}_{\parallel r})$ and agents played f thereafter. Construct h^{dev} close to η^{dev} just as we constructed \bar{h} close to $\bar{\eta}$. Now define g so that if i plays a_i at $(r, \bar{h}_{\parallel r})$, the resulting outcome will be h^{dev} . By construction, $h^{dev} \approx \eta^{dev}$, while $\bar{h} \approx \bar{\eta}$. Since i weakly prefers $\bar{\eta}$ to η^{dev} , he cannot greatly prefer h^{dev} to \bar{h} .

We proceed in this way to complete the definition of g , taking into account deviations from deviations, etc. The process is more complicated but the basic principle is the same. Each deviation by agent i from g is treated by other agents *as if* they were playing in continuous-time, and i had made the corresponding deviation from a corresponding node. Since the corresponding deviation is assumed to be unprofitable in continuous time, it can be only barely profitable in discrete time.

APPENDIX.

Proof of Theorem I.

Fix a regular decision node $(\underline{t}, \underline{h})$ and $\varepsilon > 0$. (The proof if $(\underline{t}, \underline{h})$ is the start of the game is essentially identical.) Let η denote the outcome generated by f from $(\underline{t}, \underline{h})$. First note that η must have only finitely many jumps, because f satisfies *F1*. If $\eta = \underline{h}$ then, obviously, there is nothing to prove. Assume therefore, that $\eta \neq \underline{h}$. Pick $\delta < \frac{\varepsilon}{2k(\eta)+1}$ sufficiently small that for every $k(\underline{h}) \leq \kappa \leq k(\eta)$, (a) if $d^H(\eta_{1\kappa}, h') < \delta$ and $t(h') \geq t(\eta_{1\kappa})$, then $f_i(t^{last}(\eta_{1\kappa}), \eta_{1\kappa}) = f_i(t^{last}(h'), h')$ while for all $s > t^{last}(h')$, $f(s, \eta_{1\kappa}) = f(s, h')$; (b) the smallest distance between any two discontinuity points of $f(\cdot, \eta_{1\kappa})$ exceeds 2δ . Such a δ exists, because f satisfies *F2* and *F3*. Now let R be a δ -fine grid and let h denote the outcome generated by f_{1R} from $(\underline{t}, \underline{h})$. We will show that $d^H(h, \eta) < \varepsilon$. First observe that from the definition of f_{1R} , $t^{k(\underline{h})+1}(h) = [([t^{k(\underline{h})+1}(\eta)]^R)^R]^R$, if $f_{1R}(\cdot, \eta_{1k(\underline{h})})$ is left continuous at $t^{k(\underline{h})+1}(\eta)$, and $[t^{k(\underline{h})+1}(\eta)]^R$ otherwise. Also $a^{k(\underline{h})+1}(\eta) = a^{k(\underline{h})+1}(h)$. It follows that $d^H(\eta_{1k(\underline{h})+1}, h_{1k(\underline{h})+1}) < \frac{2\varepsilon}{2k(\eta)+1}$ and $t(h_{1k(\underline{h})+1}) \geq t(\eta_{1k(\underline{h})+1})$. Now fix $\kappa > k(\underline{h})$ and assume that $d^H(\eta_{1\kappa}, h_{1\kappa}) < \frac{2(\kappa - k(\underline{h}))\varepsilon}{2k(\eta)+1}$ and $t(h_{1\kappa}) \geq t(\eta_{1\kappa})$. From our choice of δ , it follows that then $f_i(t^{last}(\eta_{1\kappa}), \eta_{1\kappa}) = f_i(t^{last}(h_{1\kappa}), h_{1\kappa})$ and, for all $s > t^{last}(h_{1\kappa})$, $f(s, \eta_{1\kappa}) = f(s, h_{1\kappa})$. Therefore, proceeding as above, $t^{\kappa+1}(h)$ is at most two grid points beyond $t^{\kappa+1}(\eta)$ and $a^{\kappa+1}(\eta) = a^{\kappa+1}(h)$. Therefore $d^H(\eta_{1\kappa+1}, h_{1\kappa+1}) < \frac{2(\kappa + 1 - k(\underline{h}))\varepsilon}{2k(\eta)+1}$ and $t(h_{1\kappa+1}) \geq t(\eta_{1\kappa+1})$. This completes the proof of Theorem I. \square

A 'proof' of statement (V.1)

The following discussion focuses on the very beginning of the race. Specifically, when we demonstrate that a profile is an equilibrium, we will consider only deviations at decision nodes whose time components are zero. We can do this without loss of generality because if it is not profitable for a contender to challenge the leader at the outset, before the anybody has accumulated any stock of knowledge, then it cannot be profitable latter in the game, since by then, the leader will have a head start.

We now discuss the equilibrium in which \bar{i} wins. We will show that \bar{i} can credibly threaten to respond immediately to any challenge that another player might have an incentive to make at time zero. More precisely, a player j might challenge \bar{i} at time zero, by increasing his investment rate to a level exceeding \bar{i} 's. However, j cannot possibly gain if he upgrades by more than $\lfloor \pi^* \rfloor$ units: even if he ultimately won the race, the incremental cost would exceed the value of the prize. (Recall that $\lfloor \pi^* \rfloor$ is the largest integer smaller than π^* .) On the other hand, we will show that if j upgrades by an amount less than $\lfloor \pi^* \rfloor$, \bar{i} will respond by upgrading his investment level to one unit above j 's, and so maintain his lead. If j retaliates, i will respond yet again. Since \bar{i} resources exceed j 's, any escalation cycle must result in a victory for \bar{i} .

More precisely, we establish that the following statements are true for any decision node $(0, h)$, any player $j < \bar{i}$ and any $0 \leq \kappa \leq j$. In what follows, we will assume w.l.o.g. that j is \bar{i} 's closest contender.

"If $a_{\bar{i}}^{last}(h) > \max(\kappa, a_j^{last}(h))$, then in any equilibrium from this subgame, j stays behinds \bar{i} (A.1. κ) for the remainder of the game.

"If $a_j^{last}(h) \geq \kappa$ and $a_{\bar{i}}^{last}(h) \geq a_j^{last}(h) + 1 - \lfloor \pi^* \rfloor$, then in any equilibrium from this (A.2. κ) subgame, \bar{i} must beat j to the finish line."

The proof is by induction. If $a_{\bar{i}}^{last}(h) > j$, then \bar{i} 's current investment level exceeds j 's maximum possible investment level, so that j must obviously stay behind \bar{i} . Therefore, (A.1. κ) is trivially true for $\kappa = j$. Now fix $1 \leq k \leq j$ and assume that (A.1. κ) is true for $\kappa = k$. We will show that (A.2. κ) is true for $\kappa = k - 1$. First, consider a profile in which j beats \bar{i} , starting from the node $(0, h)$, so that \bar{i} 's payoff is at most $-a_{\bar{i}}^{last}(h)$. We will show that this profile cannot be an equilibrium, because \bar{i} can profitably deviate by upgrading to $a_j^{last}(h) + 1$. (Since $\bar{i} > j$, this upgrade is certainly affordable for \bar{i} .) Also, by assumption, $a_j^{last}(h) \geq k$, so that if \bar{i} deviates in this way, his level will weakly exceed $k + 1$. If he deviates, therefore, the conditions of (A.1. k) will therefore be satisfied, so by assumption he will stay ahead of j . Since j is \bar{i} 's closest contender, this means that \bar{i} can win the race and attain a payoff of $\pi^* - (a_j^{last}(h) + 1) > \lfloor \pi^* \rfloor - (a_j^{last}(h) + 1)$, which by assumption exceeds $-a_{\bar{i}}^{last}(h)$. The deviation is thus profitable. This establishes (A.2. k).

Now assume that (A.2. κ) is true for $\kappa = k$. We will show that (A.1. k) is also true. Fix a decision node $(0, h')$ satisfying the assumptions of (A.1. k) and a profile that has j upgrading to a level weakly exceeding \bar{i} 's at this node. We will show that j would do better to stay put, so that this profile cannot be an equilibrium. Let $(0, h)$ be the decision node that is reached immediately after j 's upgrade. By assumption, $a_j^{last}(h) \geq a_i^{last}(h') \geq k$, so that the first condition of (A.2. κ) is satisfied. Moreover, j 's upgrade cannot be profitable if he jumps by more than $\lfloor \pi^* \rfloor$ steps. We can assume, therefore, that $a_j^{last}(h) \leq a_j^{last}(h') + \lfloor \pi^* \rfloor < a_i^{last}(h') + \lfloor \pi^* \rfloor = a_i^{last}(h) + \lfloor \pi^* \rfloor$. That is, $a_i^{last}(h) \geq a_j^{last}(h) + 1 - \lfloor \pi^* \rfloor$, so that the second condition of (A.2. k) is satisfied. Since (A.2. k) is true, \bar{i} must win the race. Therefore, j would have done better not to jump at all. This establishes that j 's jump could not have been part of an equilibrium profile from $(0, h')$ and establishes (A.1. k).

The inductive argument just given establishes that (A.1. κ) and (A.2. κ) are true for $j < \bar{i}$, and all $\kappa < j$. In particular, if \bar{i} chooses an investment level of one at the beginning of the game, and no other agent invests, then \bar{i} must win the race.

The equilibrium in which $i < \bar{i}$ takes the lead is simpler to analyze. If i invests $\lfloor \pi^* \rfloor$ units at the start of the game, then he cannot be beaten unless some player upgrades to a level exceeding π^* . But no such upgrade can be profitable. By upgrading to $\lfloor \pi^* \rfloor$, a challenger might force a tie, but since $\pi^* > 2$, this would also result in a net loss.

We now show that there can be no other equilibria. First note that in *any* equilibrium that has \bar{i} investing at time zero and subsequently winning, \bar{i} must announce an initial level of one. This is because if \bar{i} is going to win, then all other agents should not invest at all. But if no other agents invest, then \bar{i} should invest at the smallest level necessary to ensure a win. On the other hand, there can be no equilibrium in which some $i < \bar{i}$ invests less than $\lfloor \pi^* \rfloor$ at time zero and wins. To see this, suppose that the highest bidder at the outset bids $k < \lfloor \pi^* \rfloor$. In this case, the conditions of (A.2. k) are satisfied, since $k + 1 - \lfloor \pi^* \rfloor < 0$. Therefore, \bar{i} must win from this subgame. Finally, we show that there cannot be an equilibrium in which the first investment occurs at some positive time. If $\bar{i} > \lfloor \pi^* \rfloor$, there are at least two potential winners of the race; unless the ultimate winner invests at the start of the game, one of the other potential winners can preempt and

win. If $\bar{i} \leq \lfloor \pi^* \rfloor$, then the only potential winner is \bar{i} . If he postpones his initial investment until $t > 0$, then he must invest at a level exceeding 1 to win. He could win at a lower cost by investing at level one from the outset. This completes the discussion of this example.

Proof of Proposition II.

For any history, h , let $M(h) = (M_1(h), M_2(h))$ denote the number of moves that #1 and #2 have left to make. Now fix an arbitrary decision node, (t, h) . If $M(h) = (n, 0)$ and $a_1^{last}(h) \neq a_2^{last}(h)$, then, clearly, any strategy that has #1 failing to match #2 immediately is weakly dominated. We have established, therefore, that:

If $M(h) = (n, 0)$ and $a_1^{last}(h) \neq a_2^{last}(h)$, then after iterated elimination (a.i.e.), #1 must switch to #2's action immediately. (A.3)

Similarly, it is obvious that

If $M_i(h) = 0$ and $a_j^{last}(h) = 'bad'$, then a.i.e., j will continue to play 'bad'. (A.4)

Next, consider a node (t, h) such that $M(h) = (1, 1)$ and $a_1^{last}(h) = a_2^{last}(h)$. If both agents are playing 'good', then for each i , switching at any point beyond t is weakly dominated with respect to j 's reduced strategy set: this is because if i jumps, then j will follow suit immediately, so that a jump by one results in strictly lower payoffs for each. On the other hand, if both are consuming, then player #2 must stop investing immediately. If he does so, then #1 will follow suit immediately, and thus #2 will achieve the maximum possible continuation payoff. Thus, any strategy other than "switch immediately" is now weakly dominated for #1. We have established, therefore:

If $M(h) = (1, 1)$ and $a_1^{last}(h) = a_2^{last}(h)$, then a.i.e., both agents will continue to play 'good' for the remainder of the game. (A.5)

We now proceed to the inductive step. We will prove the following statements, each $\kappa \leq n^*$.

if $M(h) = (\kappa, \kappa - 1)$ and $a_1^{last}(h) \neq a_2^{last}(h)$, then in any iteratively undominated equilibrium, at least one of the players will immediately switch to whatever action the other is playing. (A.6. κ)

If $M(h) = (\kappa, \kappa)$ and $a_1^{last}(h) = a_2^{last}(h) = 'bad'$, a.i.e., #2 will switch immediately to play 'good'. (A.7. κ)

Statement (A.6.1) is implied by (A.4) and (A.7.1) by (A.5). We now show that statement (A.6.2) is

true. Suppose that $M(h) = (2, 1)$ and $a_1^{last}(h) \neq a_2^{last}(h)$. Suppose that there exists an uneliminated strategy

for #2 that has him play $a_2^{last}(h)$ for an interval of time in this state. We will show that in this case, "switching immediately" must dominate "not switching immediately." for #1. First suppose that $a^{last}(h) = (bd\ gd)$. If #1 switches to 'good' at (t, h) and #2 stays put at this node, then each will have one move remaining and (A.5) above will apply; both players will continue to play 'good' for the remainder of the game. This is the best possible outcome for #1 from this node. Had #1 *not* switched immediately and #2 continued to play 'bad' for an interval, this outcome would not be attained. This establishes that "switching" is strictly better than "not switching" against some uneliminated strategy for #1. Now suppose that #2 switches simultaneously with #1 at (t, h) . In this case, the new state will be $(gd\ bd)$; #1 will have one move remaining, and #2 will have none. Condition (A.6.1) above will apply: #1 will switch back to 'bad', and the state will be $(bd\ bd)$ for the rest of the game. Had #1 played 'bad' at this node, then the next state would again be $(bd\ bd)$. Since #2 would now have no more moves, (A.4) would apply and, once again, the state would remain at $(bd\ bd)$ for the rest of the game. We have established, therefore, that #1 cannot do worse, and may do better, if he switches immediately, rather than continues to play 'bad'.

Now suppose that $a^{last}(h) = (gd\ bd)$ and #1 switches to 'bad' at (t, h) . If #2 stays put: each will have one move left; (A.7.1) will apply, #2 will switch back to 'good' and #1 will immediately follow suit. The best possible outcome for #1 will result. On the other hand, if #2 switches to 'good' at the same instant that #1 switches to 'bad', then #2 will have no moves left, and (A.6.1) will apply. #1 will switch back to 'good' and, once again, the best possible outcome for #1 will result. This completes our verification of (A.6.2).

(A.7.2) now follows immediately: If $M(h) = (2, 2)$ and $a^{last}(h) = bd\ bd$, then #2 should switch immediately to 'good'. (A.6.2) would then apply; #1 will switch to 'good'. By (A.4), this state will be maintained until the end of the game.

To complete the inductive step, we need to show that for $k > 2$, if (A.6.k-1) and (A.7.k-1) are true, then (A.6.k) and (A.7.k) are also true. In the argument that established (A.6.k) and (A.7.k) for $k = 2$, we used the particular value of k at only one point: when $a^{last}(h) = bd\ gd$ and both players switched their actions simultaneously. We now return to this particular stage of the argument, for the general case in which Suppose that $M(h) = (k, k-1)$, $k \geq 3$ and both switch simultaneously. In this case, (A.6.k-1) will apply.

Now retrace our earlier argument to establish that agents will eventually (but in zero time!) arrive at $(gd\ gd)$ and stay there for the rest of the game. This completes the proof of the inductive step.

To complete the proof of Proposition II, we need to consider agents' actions at the start of the game. We can conclude from facts (A.5), (A.6. κ) and (A.7. κ) that if #1 and #2 choose the *same* action at the start of the game, then they will "end up" playing $(gd\ gd)$ at time zero. They will then remain in this state for the remainder of the game.⁵³ If agents choose the same actions initially, therefore, we are done. Consider a profile that has them choosing different actions at the outset, and that in the resulting outcome, they choose some state *other* than $(gd\ gd)$ for an interval of time. In this case, #1's payoff will be strictly lower than if he had matched #2 at the outset. But whatever #2 plays at the beginning of the game, #1 *could have* matched his action, and so *could have* attained the maximum possible payoff. Therefore the candidate profile cannot be even a *Nash* equilibrium. This completes the proof of Proposition II. \square

Proof of Theorem III.

Suppose there is a strategy \hat{f}_i , a decision node $(\underline{t}, \underline{h})$ and $\gamma > 0$ such that $P_i(\hat{f}_i, f_{-i}, \underline{t}, \underline{h}) > P_i(f, \underline{t}, \underline{h}) + 2\gamma$. Let $\hat{f} = (\hat{f}_i, f_{-i})$. Let η ($\hat{\eta}$) denote the outcome generated by f (\hat{f}) from $(\underline{t}, \underline{h})$. For each n , let $r^n = \lfloor \underline{t} \rfloor^{R^n}$, and let \underline{h}^n denote the closest R^n -admissible history whose jump times weakly exceed those of \underline{h} . Let h^n (\hat{h}^n) denote the outcome generated by $f_{|R^n}$ ($\hat{f}_{|R^n}$) from (r^n, \underline{h}^n) . Using Theorem I, and the fact that strategies are strongly right continuous (F3), we have $h^n \xrightarrow{n} \eta$ and $\hat{h}^n \xrightarrow{n} \hat{\eta}$. For n sufficiently large, therefore, $V_i(\hat{h}^n) > V_i(h^n) + \gamma$ and $\epsilon^n < \gamma$. But this contradicts the assumption that the R^n -admissible profile derived from $f_{|R^n}$ is an ϵ^n -SGP equilibrium for the game played on R^n . This completes the proof of Theorem III. \square

⁵³ This is not strictly accurate. Agents may "frivolously" cycle through a sequence of consecutive switches, but these cycles must end at the same instant that they started. Cycles of this kind would of course be silly, but strategies that are at worst silly cannot be excluded by weak dominance arguments.

Proof of Theorem IV

To operationalize the restriction on jumps, assume that any history that has i jumping more than n^* jumps yields a payoff of $-\infty$. Once a grid R has been fixed, we will simplify notation by referring to $[r]^R$ simply as $[r]$ and $[r]^R$ as $[r]$. Let $\lceil t \rceil$ denote the smallest $r \in R$ that is *weakly* greater than t . Recall from p. 40 that $h_{\parallel r}$ is the history defined by truncating h after the last jump that occurs strictly before r . Finally, let $k(h, r)$ denotes the last jump of h that occurs strictly before r .

Picking a sufficiently fine grid.

Fix $\varepsilon < 1$. First pick a positive $\delta^a < \varepsilon$ such that condition (*) is satisfied for $\varepsilon/3$, i.e., if two histories h and h' agree before $1 - \delta^a$ and $h_i = h'_i$, then $|V_i(h) - V_i(h')| < \varepsilon/3$. Then pick a positive δ^b such that $\max_i |V_i(h) - V_i(h')| < \varepsilon/3$ whenever $d^H(h, h') < \varepsilon$. Finally, set $\delta^* = \frac{\delta^a \wedge \delta^b}{2(n^* \bar{i})^2}$.⁵² Now pick any δ^* -fine grid R . Our choice of δ^* guarantees that there will be $n^* \bar{i}^2$ grid points between $1 - \delta^b$ and 1. We will now define an R -admissible strategy profile, g that will ε -approximate f and be an ε -SGP equilibrium for the game played on R .

Constructing the ε -approximation of f .

We will simultaneously define a map m from R -admissible decision nodes to continuous-time nodes. $m(r, h)$ will be the node close to (r, h) identified in the statement of the theorem, i.e., the outcomes generated by g and f from, respectively, (r, h) and $m(r, h)$, will be close. The procedure is the formal version of the one described on pp. 38-40.

Set $m(0, \emptyset) = (0, \emptyset)$. Proceed as follows: (a) construct the continuous-time history $\bar{\eta}$ that is generated by f from $(0, \emptyset)$; (b) identify a history \bar{h} that is very close to $\bar{\eta}$; (c) define the map m at each decision node along the path defined by \bar{h} after $(0, \emptyset)$. The map m is defined only for nodes (r, h) such that $r \leq 1 - \delta^b$. (d) specify the profile g at each node for which m has been defined; construct g so that the out-

⁵² ' $x \wedge y$ ' denotes the minimum of x and y . Similarly, ' $x \vee y$ ' denotes the maximum of x and y .

come generated by g from the node is \bar{h} . (e) consider each node $(\underline{r}, \underline{h})$ such that $m(\lfloor \underline{r} \rfloor, \underline{h}_{\parallel \lfloor \underline{r} \rfloor^*})$ was defined in step (c) and define $m(\underline{r}, \underline{h})$. We can now go to the inductive step. After step (e), there are number of nodes for which m has been defined but for which g has *not* been defined. At each such node, repeat steps (a)-(d) above. replacing $(0, \emptyset)$ with the appropriate node. Clearly, after finitely many repetitions of the inductive step, we will have defined m and g on DN^R .

Step (a): Let $\bar{\eta}$ denote the history generated by f from $(0, \emptyset)$.

Step (b): Associate to $\bar{\eta}$ the close R -admissible history \bar{h} , uniquely identified as follows: $\mathbf{a}(\bar{h}) = \mathbf{a}(\bar{\eta})$,

$$t^1(\bar{h}) = t^1(\bar{\eta}) \text{ and for } \kappa > 1, t^\kappa(\bar{h}) = \begin{cases} \max([t^\kappa(\bar{\eta})], [t^{\kappa-1}(\bar{\eta})]) & \text{if } t^\kappa(\bar{\eta}) \leq 1 - \delta^b \\ \max([t^{\kappa-1}(\bar{h})], [(1 - \delta^b)]) & \text{if } t^\kappa(\bar{\eta}) > 1 - \delta^b \end{cases} \cdot \bar{h} \text{ is the}$$

closest R -measurable history to $\bar{\eta}$ whose vector of jump-times exceeds the vector of jump-times of \bar{h} .

Observe that for \bar{h} such that $t^{last}(\bar{h}) < (1 - \delta^b)$, a jump-time of \bar{h} can be separated by at most $n^* \bar{i}$ grid points from the corresponding jump-time of $\bar{\eta}$. Since the maximum number of jumps is $n^* \bar{i}$, the "total" separation between two histories is $n^* \bar{i}^2$ grid points. Since $\delta^* \leq \frac{\delta^b}{(n^* \bar{i})^2}$, it follows that $d^H(\bar{\eta}, \bar{h}) < \frac{1}{2} \delta^b$.

If $t^{last}(\bar{h}) \geq (1 - \delta^b)$, extend the above argument, incorporating the fact that $\delta^* \leq \frac{\delta^a}{(n^* \bar{i})^2}$,

Step (c): For each node $(r, \bar{h}_{\parallel r}) \in DN^R$ such that $r > 0$, we define the node

$$m(r, \bar{h}_{\parallel r}) = \left(m_{[0,1)}(r, \bar{h}_{\parallel r}), m_H(r, \bar{h}_{\parallel r}) \right) \text{ as follows: } m_{[0,1)}(r, \bar{h}_{\parallel r}) = \begin{cases} t^\kappa(\bar{\eta}) & \text{if } r = t^\kappa(\bar{h}), \text{ for some } \kappa \\ r & \text{otherwise} \end{cases},$$

and $m_H(r, \bar{h}_{\parallel r}) = \bar{\eta}_{\lfloor \kappa(\bar{h}, r) \rfloor}$.

Step (d): Define g so that it generates \bar{h} from $(\underline{r}, \underline{h})$, i.e., for $(r, \bar{h}_{\parallel r})$ such that $0 < r < 1 - \delta^b$, set $g(r, \bar{h}_{\parallel r}) = a^{last}(\bar{h})$. Beyond $1 - \delta^b$, pick the history that maximizes i 's payoff, starting from the appropriate subgame, and construct g_i to "follow" that history. (In general, other players' will not be following the same history, but by (*), i 's payoff will still be almost maximized).

Step (e): Consider a κ -length history \tilde{h} and a node (\tilde{r}, \tilde{h}) such that $\tilde{r} \leq 1 - \delta^b$ and $m(\lfloor \tilde{r} \rfloor, \tilde{h}_{\lfloor \kappa-1 \rfloor})$

was defined in step (c). Let $\underline{\eta} = m_H(\lfloor \tilde{r} \rfloor, \tilde{h}_{\lfloor \kappa-1 \rfloor})$. Let

$\bar{\tau}_{-i} = \inf\{s > m_{[0,1]}([\bar{F}], \tilde{h}_{1\kappa-1}) : f_{-i}(s, \underline{\eta}) \neq a_{-i}^{last}(\tilde{h})\}$. We now define the continuous time node $m(\bar{F}, \tilde{h})$. There are three cases to consider. Case (3) requires the choice of a positive γ^* . Since the details of choosing γ^* are tedious but routine, we defer them to the end. Before proceeding, we describe the delicate part of the specification of m . Suppose i deviates unilaterally from g by not jumping at some node at which he is supposed to be the only agent to jump. If this deviation were made at the corresponding node in continuous time, the result would be that some agent other than i would immediately afterwards. m is designed to "catch" this continuous time reaction and replicate it at the appropriate discrete-time decision node. For example, suppose that f has i alone playing 'left' at the continuous-time node $(\frac{1}{2}, \text{nobody has moved})$, and j alone playing 'left' at the node $(t, \text{nobody has moved})$, for every $t \in (\frac{1}{2}, \frac{1}{2} + \epsilon)$. In continuous time, if i didn't move at $\frac{1}{2}$, the resulting outcome would be that j would move at $\frac{1}{2}$. To preserve the strategic flavor of f in discrete-time, we need to "catch" j 's reaction and replicate it at the node discrete-time node $(\lceil \frac{1}{2} \rceil, \text{nobody has moved})$. To accomplish this, we set $m(\lceil \frac{1}{2} \rceil, \text{nobody has moved}) = \frac{1}{2} + \gamma^*$, where γ^* is chosen in advance to be small enough to catch each of the finite number of possible such reactions. We then run f from this node, and use the resulting outcome to define g at $(\lceil \frac{1}{2} \rceil, \text{nobody has moved})$. In this case $g_j(\lceil \frac{1}{2} \rceil, \text{nobody has moved})$ will be equal to 'left'. We now consider the various cases.

(1) At least two agents deviated unilaterally from g at the preceding grid point, i.e., for all i , there exists $j \neq i$ such that $g_j([\bar{F}], \tilde{h}_{1\kappa-1}) \neq a_j^{last}(\tilde{h})$. In this case, set $m(\bar{F}, \tilde{h}) = (\bar{F}, \tilde{h})$. We can be casual about how we define m at such nodes, because they cannot be reached by any unilateral deviation.

(2) Some i deviated unilaterally from g at the preceding grid point and \bar{F} is the first grid point after a jump in \tilde{h} . That is, $i^{last}(\tilde{h}) = [\bar{F}]$ and there exists i such that $g_i([\bar{F}], \tilde{h}_{1\kappa-1}) \neq a_i^{last}(\tilde{h})$ while $g_{-i}([\bar{F}], \tilde{h}_{1\kappa-1}) = a_{-i}^{last}(\tilde{h})$. In this case, $m_{[0,1]}(\bar{F}, \tilde{h})$ is different depending on whether or not agents other than i jump at the preceding grid point. Specifically, set

$$m_{[0,1]}(\bar{F}, \tilde{h}) = \begin{cases} m_{[0,1]}([\bar{F}], \tilde{h}_{1\kappa-1}) & \text{if } a_{-i}^x(\tilde{h}) = a_{-i}^{x-1}(\tilde{h}) \\ \min[m_{[0,1]}([\bar{F}], \tilde{h}_{1\kappa-1}) + \gamma^*, \bar{\tau}_{-i}] & \text{if } a_{-i}^x(\tilde{h}) \neq a_{-i}^{x-1}(\tilde{h}) \end{cases} \quad \text{where}$$

$\bar{\tau}_{-i} = \inf\{s > m_{[0,1]}([\bar{F}], \tilde{h}_{1\kappa-1}) : f_{-i}(s, \underline{\eta}) \neq a_{-i}^{last}(\tilde{h})\}$. Let $m_H(\bar{F}, \tilde{h})$ be the κ -length history η defined by

$(t^{last}(\eta), a^{last}(\eta)) = (m_{[0,1]}(\tilde{r}, \tilde{h}), a^{last}(\tilde{h}))$ and $\eta_{|\kappa-1} = \underline{\eta}$.

(3) Some i deviated unilaterally from g at the preceding grid point and there is at least one grid point between \tilde{r} and the most recent jump in \tilde{h} . That is, $t^{last}(\tilde{h}) < \lfloor \tilde{r} \rfloor$ and there exists i such that $g_i(\lfloor \tilde{r} \rfloor, \tilde{h}_{|\kappa-1}) \neq a_i^{last}(\tilde{h})$ while $g_{-i}(\lfloor \tilde{r} \rfloor, \tilde{h}_{|\kappa-1}) = a_{-i}^{last}(\tilde{h})$. Choose γ^* according to the procedure described

at the end of the proof and define $m_{[0,1]}(\tilde{r}, \tilde{h}) = \begin{cases} \min[m_{[0,1]}(\lfloor \tilde{r} \rfloor, \tilde{h}_{|\kappa-1}) + \gamma^*, \tilde{\tau}_{-i}] & \text{if } \tilde{\tau}_{-i} \leq \tilde{r} \\ \tilde{r} & \text{otherwise} \end{cases}$. Our

choice of γ^* ensures that $f_{-i}(\cdot, \underline{\eta}) = f_{-i}(\tilde{r} + \gamma^*, \underline{\eta})$ on the interval $(\tilde{r}, \tilde{r} + \gamma^*)$. Let $m_H(\tilde{r}, \tilde{h}) = m_H(\lfloor \tilde{r} \rfloor, \tilde{h})$.

We now proceed to the inductive step in the definition of g . Pick a regular node $(\underline{r}, \underline{h})$ such that $m(\underline{r}, \underline{h})$ has been defined by step (e) above, but at which g has not yet been defined. The procedure described below virtually duplicates steps (a)-(e) above, but with slight modifications.

Step (a): Let $\bar{\eta}$ denote the history generated by f from $m(\underline{r}, \underline{h})$.

Step (b): Associate to $\bar{\eta}$ the close R -admissible history \bar{h} , uniquely identified as follows: $a(\bar{h}) = a(\bar{\eta})$;

for $\kappa \leq k(\bar{h})$ such that $t^\kappa(\bar{\eta}) < (1 - \delta^b)$, $t^\kappa(\bar{h}) = t^\kappa(\bar{\eta})$; otherwise,

$$t^\kappa(\bar{h}) = \begin{cases} \max([t^\kappa(\bar{\eta})], [t^{\kappa-1}(\bar{\eta})]) & \text{if } t^\kappa(\bar{\eta}) \leq 1 - \delta^b \\ \max([t^{\kappa-1}(\bar{h})], [1 - \delta^b]) & \text{if } t^\kappa(\bar{\eta}) \geq 1 - \delta^b \end{cases} \text{ . Observe as before that } d^H(\bar{\eta}, \bar{h}) < \frac{1}{2}\delta^b \text{ .}$$

Step (c): For each node $(r, \bar{h}_{\parallel r}) \in DN^R$ such that $r > \underline{r}$, define $m(r, \bar{h}_{\parallel r}) = (m_{[0,1]}(r, \bar{h}_{\parallel r}), m_H(r, \bar{h}_{\parallel r}))$

as follows: $m_{[0,1]}(r, \bar{h}_{\parallel r}) = \begin{cases} t^\kappa(\bar{\eta}) & \text{if } r = t^\kappa(\bar{h}), \text{ for some } \kappa > k(\bar{h}) \\ r & \text{otherwise} \end{cases}$, and $m_H(r, \bar{h}_{\parallel r}) = \bar{\eta}_{|\kappa(\bar{h}, r)}$.

Step (d): Define g so that it generates \bar{h} from $(\underline{r}, \underline{h})$, i.e., for $(r, \bar{h}_{\parallel r})$ such that $\underline{r} \leq r < 1 - \delta^b$, set $g(r, \bar{h}_{\parallel r}) = a^{last}(\bar{h})$. Beyond $1 - \delta^b$, pick the history that maximizes i 's payoff, starting from the appropriate subgame, and construct g , to "follow" that history.

Step (e): Repeat exactly as above.

This completes the inductive definition of g and m . Note that for every $(\underline{r}, \underline{h})$, $m_{[0,1]}(\underline{r}, \underline{h})$ and r are separated by at most $n^* \bar{i}$ grid points. Therefore $|m_{[0,1]}(\underline{r}, \underline{h}) - r| < \frac{1}{2} \delta^b$, so that $|m_{[0,1]}(\underline{r}, \underline{h}) - r| + d^H(\bar{\eta}, \bar{h}) < \frac{1}{2} \delta^b$. Also, from (c), the history $\bar{\eta}$ generated by f from $m(\underline{r}, \underline{h})$ is within δ^b of the history \bar{h} generated by g from $(\underline{r}, \underline{h})$. Therefore g indeed ε -approximates f .

Showing that g is an ε -equilibrium.

Suppose that there exists a strategy \hat{g}_i and a decision-node (r^*, h^*) such that i prefers the history generated by (\hat{g}_i, g_{-i}) from this node to the one generated by g . Let \hat{h} denote this history. We can assume that \hat{h}_i has no more than n^* jumps (otherwise $V_i(\hat{h}) = -\infty$). Also, we will assume that $t^{last}(\hat{h}) < 1 - \delta^b$. (The argument below needs modifying if this is not true. Since the modification is obvious, we will leave it to the reader.) We will construct a corresponding deviation \hat{f}_i such that when (\hat{f}_i, f_{-i}) is played from the close continuous-time decision node $m(r^*, h^*)$, the resulting history is very close to \hat{h} . Since the outcome generated by g from (r^*, h^*) is also very close to the outcome generated by f from $m(r^*, h^*)$, we will be done. (We will assume that (r^*, h^*) is a regular decision node; the argument if $(r^*, h^*) = (0, \emptyset)$ is virtually the same.)

Define \hat{f}_i as follows, for each $k(\hat{h}) \geq \kappa > k(h^*)$, each $s \geq t^{\kappa-1}(\hat{h})$, each η such that $a(\eta) = a(\hat{h}_{1_{\kappa-1}})$

and each $s > t^{last}(\eta)$,

$$\hat{f}_i(s, h) = \begin{cases} a_i^\kappa(\hat{h}) & \text{if } a_i^\kappa(\hat{h}) \neq a_i^{\kappa-1}(\hat{h}), a_{-i}^\kappa(\hat{h}) \neq a_{-i}^{\kappa-1}(\hat{h}) \text{ and } f_{-i}(s, m_H(t^\kappa, \hat{h}_{1_\kappa})) \neq a_{-i}^{\kappa-1}(\hat{h}) \\ a_i^\kappa(\hat{h}) & \text{if } a_i^\kappa(\hat{h}) \neq a_i^{\kappa-1}(\hat{h}), a_{-i}^\kappa(\hat{h}) = a_{-i}^{\kappa-1}(\hat{h}) \text{ and } s \geq m_{[0,1]}(\lceil t^\kappa(\hat{h}) \rceil, \hat{h}_{1_\kappa}) \\ a_i^{\kappa-1}(\hat{h}) & \text{otherwise} \end{cases} \quad \text{Define } \hat{f}_i$$

to be constant on the remaining decision nodes. Since \hat{h}_i has no more than n^* jumps, $\hat{f}_i(\cdot, \eta)$ will instruct i to jump no more than n^* times. Also, for every η , $\hat{f}_i(\cdot, \eta)$ has at most one discontinuity, so that F2 is satisfied. Finally, \hat{f}_i is jump-time independent so that F3 is satisfied.

Let $\hat{f} = (\hat{f}_i, f_{-i})$ and let $\hat{\eta}$ denote the outcome generated by \hat{f} from $m(r^*, h^*)$. We will show that

$$\text{for all } \kappa \geq k(h^*), m_H(r^* \vee \lceil t^\kappa(\hat{h}) \rceil, \hat{h}_{1_\kappa}) \underset{f}{\sim} \hat{\eta}_{1_\kappa}. \quad (\text{A.8.}\kappa)$$

where we write $h \underset{f}{\sim} h'$ if f treats h and h' identically, in the sense of restriction F3. Since $\hat{h}_{1k(\hat{h})} = \hat{h}$, $\hat{\eta}_{1k(\hat{\eta})} = \hat{\eta}$ and the map m separates histories by no more than δ^* , we will have established that $d^H(\hat{\eta}, \hat{h}) < \delta^*$, once we have established (A.8. κ) for $\kappa = k(\hat{h})$. For $\kappa = k(h^*)$, $r^* \geq [t^\kappa(\hat{h})]$, so we are concerned with $m(r^*, \hat{h}_{1\kappa})$. But by definition, $m(r^*, \hat{h}_{1\kappa}) = m(r^*, h^*)$, so that (A.8. $k(h^*)$) is trivially satisfied.

Now assume that for $\kappa > k(h^*)$, (A.8. $\kappa-1$) is true. We will show that (A.8. κ) is true. Set $(t^*, e^*) = m(r^*, h^*)$. Set $\underline{r} = r^* \vee [t^{\kappa-1}(\hat{h})]$ and $\underline{t} = t^* \vee [t^{\kappa-1}(\hat{\eta})]$. Set $r^{+1} = [t^\kappa(\hat{h})]$ and $r^0 = t^\kappa(\hat{h})$. Set $t^0 = m_{[0,1]}(r^0, \hat{h}_{1\kappa-1})$. Let \bar{h} denote the history generated by g from $(\underline{r}, \hat{h}_{1\kappa-1})$ and $\bar{\eta}$ denote one generated by f from $m(\underline{r}, \hat{h}_{1\kappa-1})$.

First suppose that $r^0 = r^*$. If i moves alone at this time, then our construction of \hat{f}_i has i jumping exactly at t^* . Since $g_{-i}(r^*, \hat{h}_{1\kappa-1}) = a_{-i}^{\kappa-1}(\hat{h})$ it must be the case that $f_{-i}(t^*, \hat{\eta}_{1\kappa-1}) = a_{-i}^{\kappa-1}(\hat{h})$. Therefore, $\hat{\eta}$ has i jumping alone to $a_{-i}^{\kappa-1}(\hat{h})$ at t^* . Now suppose some agent other than i jumps at r^* . By our construction of g_{-i} , f_{-i} must have agents other than i moving either at or just after t^* . If i also jumps at r^* , then our construction of \hat{f}_i ensures that i moves exactly when these other agents move, i.e., either at or just after t^* .

Next, suppose that $r^0 > r^*$ and $\hat{h}_{1\kappa} = \bar{h}_{1\kappa}$. This case can happen only if $\hat{g}_i(\cdot, \hat{h}_{1\kappa-1}) = g_i(\cdot, \hat{h}_{1\kappa-1})$ at every grid point in the interval $[\underline{r}, r^0]$. In this case, it follows from the definition of \hat{f}_i that $\hat{f}_i(\cdot, \hat{\eta}_{1\kappa-1}) = f_i(\cdot, \hat{\eta}_{1\kappa-1})$ on the interval $(\underline{t}, t^\kappa(\bar{\eta})) = (\underline{t}, t^0]$. In this case, by step (c) of the algorithm, $m_H([t^\kappa(\hat{h})], \hat{h}_{1\kappa}) = \bar{\eta}_{1k(\bar{h}, [t^\kappa(\hat{h})])} = \hat{\eta}_{1\kappa}$, so that (A.8. κ) is true.

Finally, suppose that $r^0 > r^*$ and that $\hat{h}_{1\kappa} \neq \bar{h}_{1\kappa}$. In this case, $m_H(r^{+1}, \hat{h}_{1\kappa})$ is defined in Step (e) of the algorithm. Since a jump occurred at the previous grid point, case (2) applies. There are two cases to consider: (A) player i jumps alone at $t^\kappa(\hat{h})$; (B) players other than i jump at $t^\kappa(\hat{h})$. In either case, $m_H(r^{+1}, \hat{h}_{1\kappa})$ is the κ -length history η defined by $(t^{last}(\eta), a^{last}(\eta)) = (m_{[0,1]}(r^{+1}, \hat{h}_{1\kappa}), a^\kappa(\hat{h}))$ and $\eta_{1\kappa-1} = \hat{\eta}_{1\kappa-1}$, where

$$m_{[0,1]}(r^{+1}, \hat{h}_{1\kappa}) = \begin{cases} m_{[0,1]}(r^0, \hat{h}_{1\kappa-1}) & \text{if case (A) applies} \\ \min[m_{[0,1]}(r^0, \hat{h}_{1\kappa-1}) + \gamma^*, \hat{t}_{-i}] & \text{if case (B) applies} \end{cases} \quad \text{where}$$

$\hat{\tau}_{-i} = \inf\{s > m_{[0,1]}(r^0, \hat{h}_{1\kappa-1}) : f_{-i}(s, \hat{\eta}_{1\kappa-1}) \neq a_i^{\kappa-1}(\hat{h})\}$. We will establish that $\hat{\eta}_{1\kappa} \stackrel{f}{\sim} m_H(r^{+1}, \hat{h}_{1\kappa})$, by

showing that the difference in the κ 'th jump-times of the two histories is at most γ^* . By construction of \hat{f}_i ,

we have:
$$\hat{f}_i(s, \hat{h}_{1\kappa-1}) = \begin{cases} a_i^{\kappa-1}(\hat{h}) & \text{if case (A) applies and } s < m_{[0,1]}(r^{+1}, \hat{h}_{1\kappa}) \\ a_i^{\kappa-1}(\hat{h}) & \text{if case (B) applies and } f_{-i}(s, \hat{\eta}_{1\kappa-1}) = a_i^{\kappa-1}(\hat{h}) \\ a_i^{\kappa}(\hat{h}) & \text{otherwise} \end{cases}$$
. In case (B), our

choice of γ^* guarantees that g_{-i} jumps at $(r^0, \hat{h}_{1\kappa-1})$ to the same place that f_{-i} jumps to at $m_{[0,1]}(r^0, \hat{h}_{1\kappa-1})$; moreover, our construction of \hat{f}_i ensures that in continuous time, i jumps simultaneously with other agents, just as he does in discrete time. Therefore, by definition of $m_H(r^{+1}, \hat{h}_{1\kappa})$, we will be done once we have shown that in case (A) no agent other than i jumps at or before $m_{[0,1]}(r^0, \hat{h}_{1\kappa-1})$, while in case (B), no agent other than i jumps before $\hat{\tau}_{-i}$.

Define $r^{-1} = \lfloor r^0 \rfloor$ and set $t^{-1} = m_{[0,1]}(r^{-1}, \hat{h}_{1\kappa-1})$. We first establish that in case (A), f_{-i} has no jumps in the interval $(t^{-1}, m_{[0,1]}(r^{+1}, \hat{h}_{1\kappa}))$. Since $m_{[0,1]}(r^{+1}, \hat{h}_{1\kappa})$ was defined at Step (e)(2) and no agent other than i jumped at the last jump, $m_{[0,1]}(r^{+1}, \hat{h}_{1\kappa}) = t^0$. We need therefore consider how t^0 was defined. If i played according to g_i at the node $(r^{-1}, \hat{h}_{1\kappa-1})$, then g called for all players to play $a_i^{\kappa-1}(\hat{h})$ at r^{-1} , then t^0 was defined at step (c) of the algorithm: it was set equal to r^0 , if nobody jumped at the node $(r^0, \hat{h}_{1\kappa-1})$, and otherwise equal to the first time after t^{-1} that some agent jumped. In either case, therefore, f_{-i} cannot have jumps between t^{-1} and t^0 . If i deviated at $(r^{-1}, \hat{h}_{1\kappa-1})$, t^0 was defined at step (e), case (3), of the algorithm. In this case, had f_{-i} had a jump between t^{-1} and r^0 , then t^0 would have "caught" this jump and g_{-i} would have replicated this jump at $(r^0, \hat{h}_{1\kappa-1})$, contradicting the hypothesis that case (A) applies. Now suppose that case (B) applies. In this case, $m_{[0,1]}(r^{+1}, \hat{h}_{1\kappa})$ was set equal to or just beyond $\hat{\tau}_{-i}$, so that agents other than i jump only barely before $m_{[0,1]}(r^{+1}, \hat{h}_{1\kappa})$.

If $\underline{t} = r^{-1}$, then we are done. Assume therefore that $\underline{t} < r^{-1}$. We now show that $f_{-i}(\cdot, \hat{\eta}_{1\kappa-1}) = a_i^{\kappa-1}(\hat{h})$ on $[\underline{t}, t^{-1}]$. Suppose to the contrary that there exists $s^0 \in [\underline{t}, t^{-1}]$ such that $f_{-i}(\cdot, \hat{\eta}_{1\kappa-1}) \neq a_i^{\kappa-1}(\hat{h})$. Clearly, $s^0 \neq \underline{t}$; otherwise f_{-i} would have agents other than i jumping way from

$a_i^{\kappa-1}(\hat{h})$ exactly at \underline{t} and so g_{-i} would similarly have agents jumping away from $a_i^{\kappa-1}(\hat{h})$ at \underline{t} . Assume therefore, that $s^\circ > \underline{t}$ and let r' denote the largest $r \in R$ such that $m_{[0,1]}(r, \hat{h}_{1\kappa-1}) < s^\circ$. Let $t' = m_{[0,1]}(r', \hat{h}_{1\kappa-1})$. Since $t^{-1} = m_{[0,1]}(r^{-1}, \hat{h}_{1\kappa-1}) \geq s^\circ$, we know that $r' < r^{-1}$ so that $\lceil r' \rceil \leq r^{-1}$. If $f_i(\cdot, \hat{\eta}_{1\kappa-1}) = a_i^{\kappa-1}(\hat{\eta})$ on the interval $[t', s^\circ)$, then the outcome generated by f from the continuous time node $(t', \hat{\eta}_{1\kappa-1})$ has some player other than i jumping at s° . But since this outcome is what defines g at $(\lceil r' \rceil, \hat{h}_{1\kappa-1})$, the agents that f_{-i} has jumping at $(t', \hat{\eta}_{1\kappa-1})$ must also be jumping at $(\lceil r' \rceil, \hat{h}_{1\kappa-1})$. But this is a contradiction, since $r' \leq r^{-1} < r^0$ and by assumption, $g_i(\cdot, \hat{h}_{1\kappa-1})$ is constant on (\underline{t}, r^0) . Assume therefore that $f_i(\cdot, \hat{\eta}_{1\kappa-1}) \neq a_i^{\kappa-1}(\hat{\eta})$ on $[t', s^\circ)$. In this case, the node $(\lceil r' \rceil, \hat{h}_{1\kappa-1})$ must be defined at step (e), case (3) of the algorithm. But then $m_{[0,1]}(\lceil r' \rceil, \hat{h}_{1\kappa-1})$ is chosen so that $f_{-i}(t, \hat{\eta}_{1\kappa-1}) \neq a_{-i}^{\kappa-1}(\hat{\eta})$. To determine g at $(\lceil r' \rceil, \hat{h}_{1\kappa-1})$, we run f from the node $(m_{[0,1]}(\lceil r' \rceil, \hat{h}_{1\kappa-1}), \hat{\eta}_{1\kappa-1})$; some agent other than i jumps immediately, so that once again, $g_{-i}(\lceil r' \rceil, \hat{h}_{1\kappa-1}) \neq a_{-i}^{\kappa-1}(\hat{\eta})$. Once again, we reach a contradiction.

Choosing γ^ (this is the deferred substep of (e)(3) above).*

(The reader will note that our choice below depends crucially on the finiteness of the space H^R and the fact that f satisfies F2 and F3.) For any subset B of H , let $J(B)$ denote the set of jump-times of histories in B , i.e., $J(B) = \{t \in [0, 1) : \exists h \in B \text{ and } 1 \leq \kappa \leq k(h) \text{ s.t. } t^\kappa(h) = t\}$. Let $J(h) = J(\{h\})$.

We now describe a method of picking a finite set $S^1(h) \subset H$, for each $h \in H$. Let $h_{1K(h,t)}$ denote the largest jump of h that is weakly before t . Now define $DN(h) = \{(t, h') \in DN : t \in R \cup J(h), h'_{1K(h',t)} = h_{1K(h,t)}, t^{\text{last}}(h') = t\}$. $W(h)$ consists of all t 's in $R \cup J(h)$, paired with possible truncations of h and histories that are "built" by adding chains of consecutive moves at the last jump-times of these truncations. For $B \subset H$, set $W(B) = \bigcup_{h \in B} W(h)$. Let $o^{f,t,h'}$ denote the outcome generated by f from the decision node (t, h') . Now define $S^1(h)$ by $S^1(h) = \{h'' \in H : h'' = o^{f,t,h'} \text{ or } h'' = \lim_{\delta \downarrow 0} o^{f,t+\delta,h'}\}$, for some $(t, h') \in W(h)$. For a subset B of H , set $S^1(B) = \bigcup_{h \in B} S^1(h)$. Now inductively define the k -fold composition of $S^1(B)$ by: $S^k(B) = S^1(S^{k-1}(B))$. Let $S^* = S^{\bar{k}}(H^R)$, where $\bar{k} = 2n^*i\#R$. Note that for all $k \geq \bar{k}$, $S^k(H^R) = S^{\bar{k}}(H^R)$. Define $W^* = W(S^*)$.

Pick $\gamma^* > 0$ to satisfy: (i) $\min\{|s - t| : s, t \in J(S^*) \text{ and } s \neq t\} \leq \gamma^*$; (ii) $\forall (t, h) \in W^*, \forall \gamma < \gamma^*,$
 $d^H(o^{(f, t+\gamma, h)}, \lim_{\delta \downarrow 0} o^{(f, t+\delta, h)}) \leq \gamma^*$; (iii) $\forall (t, h) \in W^*, \forall h'$ such that $t(h') \geq t(h),$
 $f(t^{last}(h), h) = f(t^{last}(h'), h')$ and, for all $s > t^{last}(h'), f(s, h) = f(s, h')$. [(i) uses the fact that $J(S^*)$ is a
finite set; (ii) uses the finiteness of W^* and assumptions F2 and F3; (iii) uses the finiteness of W^* and F3.

REFERENCES

- [1] R. Anderson, "Quick Response Equilibria," University of California Berkeley, C.R.M. I.P. #323, (1984).
- [2] Aumann, R. and S. Sorin, "Cooperation and Bounded Rationality," August, 1986, mimeo.
- [3] Fudenberg, D., R. Gilbert, J. Stiglitz and J. Tirole, "Preemption, Leapfrogging and Competition in Patent Races," *European Economic Review*, 22 (1983) 3-31.
- [4] Fudenberg, D., and D. Levine, "Subgame-Perfect Equilibria of Finite- and Infinite-Horizon Games," *Journal of Economic Theory*, 31, (1984), 251-268.
- [5] Fudenberg, D., and D. Levine, "Limit Games and Limit Equilibria," *Journal of Economic Theory*, 38, (1986), 261-279.
- [6] Fudenberg, D., and E. Maskin, "The Folk Theorem in Repeated Games with Discounting or with Incomplete Information," *Econometrica*, 50, (1982), 97-109.
- [7] Fudenberg, D., and J. Tirole, "Preemption and Rent Equalization in the Adoption of a New Technology," *Review of Economic Studies*, 52 (1985), 383-401.
- [8] Fudenberg, D., and J. Tirole, "Understanding Rent Dissipation: on the Use of Game Theory in Industrial Organization," *American Economic Review*, 77(2), (1987), 176-183.
- [9] Harris, C. and J. Vickers, "Perfect Equilibrium in a Model of a Race," *Review of Economic Studies*, 52 (1985), 193-209.
- [10] Harris, C., "Perfect Equilibrium in Infinite Horizon Games," *Journal of Economic Theory*, 37, (1985), 99-125.
- [11] Ken Hendriks and Charles Wilson, "Discrete versus Continuous Time in Games of Timing," mimeo, New York University, (1985).
- [12] Kreps, D., P. Milgrom, J. Roberts and R. Wilson, "Rational Cooperation in the Finitely-Repeated Prisoner's Dilemma," *Journal of Economic Theory*, 27, (1982), 245-252.
- [13] Vijay Krishna, "Strategic Informational Equilibrium," Columbia University D.P. #198, 1983.
- [14] Kreps, D., and R. Wilson, "Reputation and Imperfect Information," *Journal of Economic Theory*, 27, (1982), 253-279.
- [15] Luce, D. and H. Raiffa, "*Games and Decisions*," New York: John Wiley and Sons, 1957.
- [16] Moulin, H., "Dominance Solvable Voting Schemes," *Econometrica*, 47, (1979), 1337-1351.
- [17] Myerson, R., "Refinements of the Nash Equilibrium Concept," *International Journal of Game Theory*, 4, (1975), 448-460.
- [18] Radner, R., "Collusive Behavior in Oligopolies with Long but Finite Lives," *Journal of Economic Theory*, 22, (1980), 136-156.
- [19] Rubinstein, A., "Perfect Equilibrium in a Bargaining Model," *Econometrica*, 50, (1982), 97-109.
- [20] Selten, R., "Reexamination of the Perfectness Concept for Equilibrium Points in Extensive Games," *International Journal of Game Theory*, 4, (1975), 448-460.
- [21] Simon, L., "Basic Timing Games," University of California, Berkeley, Working Paper, May 1987.
- [22] Simon, L., "A Multi-stage Duel in Continuous Time," mimeo, July 1987.
- [23] Simon, L., "Properness in Games with Large Strategy Spaces," (in preparation).
- [24] Simon, L. and M. Stinchcombe, "Extensive Form Games in Continuous Time: Behavior Strategies," (in preparation).
- [25] Stinchcombe, M., "Game Theory in Continuous Time: a Process Theoretic Approach," Ph. D. dissertation, Berkeley, 1986.

Jul-87

RECENT ISSUES OF THE WORKING PAPER SERIES
OF THE DEPARTMENT OF ECONOMICS
UNIVERSITY OF CALIFORNIA, BERKELEY

Copies may be obtained from the Institute of Business and Economic Research. See the inside cover for further details.

8733 Lloyd Ulman
INDUSTRIAL RELATIONS

Feb-87.

8734 Dale J. Poirier and Paul A. Ruud
PROBIT WITH DEPENDENT OBSERVATIONS

Mar-87.

8735 Pranab Bardhan
ALTERNATIVE APPROACHES TO DEVELOPMENT ECONOMICS: AN EVALUATION

Mar-87.

8736 D. Fudenberg and E. Maskin
NASH AND PERFECT EQUILIBRIA OF DISCOUNTED REPEATED GAMES

Apr-87.

8737 Jeffrey A. Frankel
OBSTACLES TO INTERNATIONAL MACROECONOMIC POLICY COORDINATION

May-87.

8738 Jeffrey A. Frankel and Richard Meese
ARE EXCHANGE RATES EXCESSIVELY VARIABLE?

May-87.

8739 Jeffrey A. Frankel
INTERNATIONAL CAPITAL FLOWS AND DOMESTIC ECONOMIC POLICIES

May-87.

Jul-87

RECENT ISSUES OF THE WORKING PAPER SERIES
OF THE DEPARTMENT OF ECONOMICS
UNIVERSITY OF CALIFORNIA, BERKELEY

Copies may be obtained from the Institute of Business and Economic Research. See the inside cover for further details.

8740 Joseph Farrell and Garth Saloner
COORDINATION THROUGH COMMITTEES AND MARKETS

May-87.

8741 Paul R. Milgrom
EMPLOYMENT CONTRACTS, INFLUENCE ACTIVITIES AND EFFICIENT ORGANIZATION DESIGN

May-87.

8742 Paul Milgrom
ADVERSE SELECTION WITHOUT HIDDEN INFORMATION

Jun-87.

8743 Roger Craine
ARBITRAGE q : AN EQUILIBRIUM THEORY OF INVESTMENT

Jun-87.

8744 Jeffrey A. Frankel and Katharine E. Rockett
INTERNATIONAL MACROECONOMIC POLICY COORDINATION
WHEN POLICY-MAKERS DISAGREE ON THE MODEL
Jul-87.

8745 Leo K. Simon
BASIC TIMING GAMES

Jul-87.

8746 Leo K. Simon and Maxwell B. Stinchcombe
EXTENSIVE FORM GAMES IN CONTINUOUS TIME:
PURE STRATEGIES
Jul-87.

