

UC Berkeley

Other Recent Work

Title

A Multistage Duel in Continuous Time

Permalink

<https://escholarship.org/uc/item/7186g3c4>

Author

Simon, Leo K.

Publication Date

1987-10-01

Peer reviewed

UNIVERSITY OF CALIFORNIA, BERKELEY

Department of Economics

Berkeley, California 94720

Working Paper 8757

A MULTISTAGE DUEL IN CONTINUOUS TIME

Leo K. Simon

October 1, 1987


Key words: Duel, continuous time, timing games, subgame perfect equilibrium, preemption, innovation.

Abstract

Two protagonists face an intertemporal allocation problem. Each has limited, indivisible resources and must decide how best to "spend" them over time. We formulate this problem as a continuous-time game, using techniques recently developed by Simon and Stinchcombe. For the class of parameters that we consider in this paper, the problem has a unique solution, with surprising properties.

*The author has benefited greatly from discussions with Martin Osborne Carolyn Pitchik, Suzanne Scotchmer, Bill Zame and, especially, Maxwell Stinchcombe.

JEL Classification: 022, 026, 611



This paper is part of a series that develops a new approach to modelling continuous time games.¹ In the first paper (Simon-Stinchcombe [10]), we studied a pure strategy model in which agents were allowed to move finitely many times. In the second (Simon [9]), we introduced behavior strategies, but restricted attention to games that end at the instant that some agent moves. We proposed a "calculus for continuous time games," which enabled us to solve a certain class of timing games with a minimum of computation. This paper combines and illustrates the methodologies of the two earlier papers in the context of a specific application.² We exploit our "calculus" to study a problem that would be extremely tedious to solve in discrete time, and impossible to model using conventional continuous-time techniques.³ Consider the following scenario.

I. The Problem.

Two gunfighters stand at either end of a dusty street. Each has a Colt 45 and spare bullets. Apart from their ammunition supplies, they are equally matched. As they begin to walk toward each other, their accuracy increases. Since both bullets and time are expensive, they prefer not to waste either. The guns have no silencers, so that each fighter can keep track of the number of bullets his opponent has left to fire. Who will fire first? In what order will subsequent bullets be fired? Will the bullets be fired sooner or later than the optimal times?⁴ Will agents' aggregate payoffs be higher or lower as the disparity between their initial numbers of bullets is increased?

Our protagonists face an intertemporal allocation problem. Each has limited, indivisible resources, and must decide how best to allocate them over time. Since economic agents frequently face problems of this kind, the answers to the questions raised above are of interest to economists. Indeed, the scenario above can be rephrased as the following problem for industrial organization theorists.

¹ For an alternative approach to modelling continuous time games, see Stinchcombe [12].

² We also draw extensively on ideas that will appear in a third paper (Simon-Stinchcombe [10]) which is still in preparation.

³ Simon [9] discusses in detail the advantages of continuous time relative to discrete time for problems of this kind and points out why conventional continuous time techniques are inappropriate.

⁴ Our benchmark notion of optimality is the usual one: we solve the single person decision problem that defined by maximizing the *sum* of agents' expected payoffs.

There is a large potential market for some medical product. Two firms can produce the product, but neither has obtained government certification. To obtain this, a firm must undertake a costly pilot study, which may or may not succeed. As time passes and the products are refined, the probability of a successful pilot study increases. Each firm has limited financial resources and so can initiate only a finite number of studies; if all of these fail, it must quit the race. If exactly one firm obtains certification, it earns flow profits until a finite time horizon is reached. If both are certified, price competition drives profits to zero. The firms are identical except for their financial resources.

The analysis below is concerned with the economic rather than the Western scenario. We will, however, borrow some convenient duelling terminology. In particular, our protagonists will "shoot bullets" rather than "undertake pilot projects." Our main result addresses the following question. Assuming that agents start out with unequal numbers of bullets, who will shoot the first bullet and when? In what order and at what times will the remaining bullets be fired? Intuition strongly suggests that the agent who starts out with more bullets will fire first, and continue to do so until he has no more bullets than his adversary; once this point is reached, we would expect an alternating pattern thereafter. This intuition is incorrect. In the unique equilibrium for our model, the player who starts out with *fewer* bullets will fire the first shot. If this bullet misses, he will keep on shooting until he either scores a hit or runs out of bullets. The second player will hold his fire until the first has completely exhausted his supply of bullets.

Since some readers may be unfamiliar with the techniques developed in [9] [10] [11], we will study this problem at two levels. We first analyze it informally, highlighting the economic principles involved and using conventional economic reasoning. We then introduce our continuous-time model. We will proceed heuristically, presuming no prior familiarity with our methodology. Even so, less specialized readers may well be content with the first level of the analysis.

We conclude with some caveats about our result. Our purpose in this paper is to develop an interesting point and illustrate continuous-time game theory, rather than prove an abstract theorem. With this in mind, we impose some special restrictions. First, we consider only the most elementary functional form. In particular, we assume that the probability of a hit (or success) increases linearly with time.⁵ Second, our

⁵ The only role this assumption plays is to guarantee that our duopolists always fire *strictly before* the time that a monopolist with the sum of their bullets would fire. This is certainly an intuitive property and, obviously, does not depend on *linearity*, per se.

model has a unique and easily computed solution only for a certain range of parameter values. We have been unable to identify this range by analytic methods. There is, however, a simple computational test we can perform, to see whether the parameters for a given game fall within this range. Our results apply to the class of games that pass this test. By running the test repeatedly, we have verified that this is a rather large class. Finally, we impose two rather arbitrary restrictions on agents' strategies.⁶ Without question, these restrictions are unobjectionable from a behavioral standpoint. They are, however, distasteful because they are ad hoc. On the other hand, they drastically simplify our analysis.

The paper is organized as follows. In section II we analyze the problem at an heuristic level and state the results. Our continuous-time model is introduced in section III and formalized in section IV. In section V we construct equilibrium strategies for the game.

II. An introduction to the analysis

Our aim in this section is to introduce and motivate our result. We will ignore several delicate modelling issues; indeed, we shall not even specify an explicit model! Our duel is played on the interval $[0, 2)$. (The horizon is set equal to '2' to simplify the arithmetic.) There are two players.⁷ Each player can fire at any time he chooses. If he fires at $t \leq 1$, his chance of scoring a hit is t ; if he fires beyond $t = 1$, he scores a hit with certainty. If one or more players fire a bullet, play is momentarily frozen, while "nature" determines whether any bullet will hit its target. If a hit is scored, the game is over. If exactly one player scores, he is declared the winner. If both players score, then nobody wins. (This is consistent with the scenario we outlined on pp. 1-2: we assumed that if both firms' projects succeeded, price competition would drive profits to zero.) If each bullet misses, then play resumes immediately and continues until either some player scores a hit or all bullets have been fired.

Our two players will be called i and j . The 'generic' player will be denoted by ι (iota) and the generic 'other player' by $\neg\iota$ (not iota). A multistage duel is completely described by a triple (β^*, π^*, c^*) .

$\beta^* = (\beta_i^*, \beta_j^*)$ is a pair of integers, denoting the number of bullets with which i and j start out the game.

⁶ Assumptions F3 and F5, specified in section IV below.

⁷ For ease of exposition, we assume these players are male.

This pair will be called the endowment vector for the duel. The scalar $c^* > 0$ denotes the resale value of each bullet. If an player wins the game at t , he earns a flow profit of π^* from t until time 2. Thus, if the game ends at t , and a player has fired n bullets, his *ex post* payoff will be $- \$nc^*$ plus, if he wins, the "prize" of $\$ \pi^*(2 - t)$. We assume that $\pi^* > c^*$, so that playing the game is individually rational.

If i has bullets and j has none, the game reduces to a single person decision problem and is trivial to solve. Set $\Pi_i^{0,0} = 0$. When i has β_i bullets, we shall say that we are in phase $(\beta_i, 0)$ of the duel. Let $L_i^{\beta_i, 0}(t)$ denote i 's payoff if he fires his β_i -th to last bullet at t and, if he misses, fires at the optimal times thereafter. Clearly, for t less than the time he should fire his $\beta_i - 1$ -th to last bullet,

$$L_i^{\beta_i, 0}(t) = t\pi^*(2 - t) - c^* + (1 - t)\Pi_i^{\beta_i - 1, 0} \quad (2.1)$$

Let $t^{\beta_i, 0}$ denote the maximizer of $L_i^{\beta_i, 0}(t)$ and $\Pi_i^{\beta_i, 0}$ denote i 's expected payoff when he fires at this time. It is easy to verify that $(t^{\beta_i, 0})$ is a strictly decreasing sequence, starting at $t^{1,0} = 1$. Set $\Pi_j^{\beta_i, 0} = 0$.

We now consider the duopoly phases of the duel, in which each agent has bullets. For $\beta = (\beta_i, \beta_j) \gg 0$, we will say that we are in phase β of the duel if player i has β_i bullets remaining. We will solve the duopoly phases one at a time, using backward induction. We start with phase $(1,1)$, and then proceed to $(2,1)$, ..., $(\beta_i^*, 1)$, $(2,2)$, $(3,2)$, etc. Once we have determined how players will play, if at any time they enter a phase with fewer bullets than β , we can determine agents' expected payoffs if one or both of them fire at any time in phase β .⁸ We can then reduce this phase to a strategically equivalent "single move" game that ends as soon as one agent fires. Payoffs in this game correspond to the payoffs that would result if agents played in the predetermined way once they left phase β . Since there will not be a unique equilibrium way to continue to play, there will be many single move games that correspond equally well to each phase of the original game. We will establish, however, that these games are all identical "on the equilibrium path."

The single-move games we construct are examples of a classical problem in game theory, known as the "noisy duel."⁹ To distinguish these simple games from our original, multi-stage duel (which is also noisy),

⁸ Of course, we have not yet even defined the game in which these strategies will be played! We will be able to establish, however, the kinds of outcomes that these strategies must generate.

⁹ A noisy duel is so named because each player can hear his opponent's bullet at the instant it is fired. In a silent duel, by contrast, the shots cannot be heard, so that at any given time, i cannot determine whether j has fired and missed or is has been holding his fire.

we will refer to them as classical duels. Since the classical duel plays a fundamental role in our subsequent analysis, we will study it in some detail.

The classical duel.

In a classical duel, each player must choose a time in $[0, 2)$ to fire one bullet. If i fires at time t and j does not, we will say that i leads and j follows at t . Alternatively, both players may fire simultaneously at t . To define a payoff function for the classical duel, we must specify the payoff that player i will receive if either or both players fire at any $t \in [0, 2)$. We call these the "lead," "follow" and "fire simultaneously" functions for player i . In the generic duel, we denote them, respectively, by $L_i(\cdot)$, $F_i(\cdot)$ and $S_i(\cdot)$. To avoid trivialities, we will always assume that if neither player ever fires, each earns a payoff of minus infinity. We will also need to analyze the "subduels" of a given duel. Given $\tau \in [0, 2)$ and a classical duel CD , we denote by $CD(\tau)$ the subduel of CD that begins at τ . This game is defined by restricting the payoff functions for CD to the interval $[\tau, 2)$, and starting play at time τ .

The classical duels corresponding to phase (1,1) are easy to solve. We will show that if players start playing in this phase at time zero, then in any equilibrium, the first bullet will be fired at the time $t^{1,1}$ at which each player is indifferent between leading and following. That is, $t^{1,1}$ must solve the equation

$$L_i^{1,1}(t) = t\pi^*(2 - t) - c^* = (1 - t)(\pi^* - c^*) = F_i^{1,1}(t). \quad (2.2)$$

Call $t^{1,1}$ the firing time for phase (1,1). In any equilibrium, exactly one player fires at this time. There is a unique equilibrium payoff vector, $\Pi^{1,1}$, defined by:

$$\Pi^{1,1} = (\Pi_i^{1,1}, \Pi_j^{1,1}) = (t^{1,1}\pi^*(2 - t^{1,1}) - c^*, (1 - t^{1,1})(\pi^* - c^*)). \quad (2.3)$$

An important fact is that $t^{1,1} < t^{2,0}$, where $t^{2,0}$ is the optimal firing time for an industry with two bullets.

If players enter phase (1,1) at some time $t > t^{1,1}$, the next bullet will be fired immediately. There are three possible outcomes. For each t , there is an equilibrium in which i fires with probability one and $-i$ fires with probability zero. The third outcome is random. Each outcome generates a distinct payoff vector.

Clearly, the two games have very different properties. For other discussions of this game, see Pitchik [7], Hendriks-Wilson [6] and Dixit [1]. Pitchik's paper contains a useful list of references.

Phase (2,1): introducing our main result.

As we have observed, there are many classical duels that correspond to this phase. However, their payoff functions all agree on the interval $[0, t^{1,1})$.¹⁰ Moreover, there is a unique equilibrium outcome that is common to each of them and is completely determined by payoffs on the above interval. With no essential ambiguity, therefore, we can refer to *the* duel corresponding to phase (2,1), denoted by $CD^{2,1}$. For each τ , the lead and follow payoffs will be denoted by $L_i^{2,1}(\cdot)$ and $F_i^{2,1}(\cdot)$. We will specify these payoffs *only* on the interval $[0, t^{1,1})$. If i leads at time $t < t^{1,1}$, he will score a hit with probability t . With the remaining probability, the duel will enter phase (1,1), the next bullet will be fired at $t^{1,1}$ and agents will earn the payoff vector $\Pi^{1,1}$. If j leads and misses at t , then he earns nothing and i earns $\Pi_i^{2,0}$. Summarizing, for $t \in [0, t^{1,1})$,

$$L_i^{2,1}(t) = t\pi^*(2-t) - c^* + (1-t)\Pi_i^{1,1} \quad (2.4.a)$$

$$F_i^{2,1}(t) = (1-t)\Pi_i^{2,0} \quad (2.4.b)$$

$$L_j^{2,1}(t) = t\pi^*(2-t) - c^* \quad (2.4.c)$$

$$F_j^{2,1}(t) = (1-t)\Pi_j^{1,1} \quad (2.4.d)$$

Obviously, $L_j^{2,1}(\cdot)$ is strictly increasing and, for each τ , $F_i^{2,1}(t)$ is strictly decreasing. Also, as one would expect, $\pi^*(2-t)$ exceeds $\Pi_i^{1,1}$ on $[0, t^{1,1})$ so that $L_i^{2,1}(\cdot)$ is also strictly increasing on this interval. An important property of this game is that for each τ , there is a time $t_i < t^{1,1}$ at which $L_i^{2,1}(\cdot)$ intersects $F_i^{2,1}(\cdot)$. Moreover, these times are different for each player.

When modelled as a closed-loop, continuous-time game, $CD^{2,1}$ has a unique solution. The bullet is fired with probability one at $\max(t_i, t_j)$. Denote this time by $t^{2,1}$ and call it the firing-time for phase (2,1). The player who fires the bullet will be the one whose lead function *first* intersects his follow function. It turns out that that this player is j !

This result is at first sight surprising. Since i has a "second chance," while j does not, we might expect i to be the first to shoot. Indeed, this latter conclusion is suggested by the law of diminishing returns: i

¹⁰ The reason is that these payoff functions are determined by the different possible equilibrium strategies for phase (1,1). As we have observed, all of the equilibria for phase (1,1) agree, provided agents enter the phase *before* $t^{1,1}$. If they enter after this time, there are several distinct equilibrium outcomes.

presumably values his second bullet less highly than his first, and so should be more willing than j to risk "wasting" a bullet in the hope of scoring an early hit. As we shall see, these intuitions are incorrect. They fail to take into account the nature of the trade-off faced by i . In addition to the obvious "first-mover advantage," there is a less obvious "second-mover advantage" that more than offsets the former. To see this, suppose that i fires first and, if he misses, fires second. In this case, both of his bullets will be fired "too early." If j fires first and misses, on the other hand, i will become a monopolist and will fire his bullets at the optimal times. It turns out that in equilibrium, the gain to firing at the optimal times rather than suboptimally more than offsets the risk i takes that he will never get to shoot, because j 's bullet finds its target.

We now explain the argument in more detail. To establish that player j is the one to fire at this time, we need only establish that at $t^{2,1}$, j *already* strictly prefers leading to following. Observe from (2.4.a)-(2.4.d) that i 's payoffs both to leading and following exceed the corresponding payoffs for j . We will show that the *difference* between the two players' 'follow' payoffs strictly exceeds the *difference* between their 'lead' payoffs. As a consequence, the intersection of i 's lead and follow functions must occur strictly later than the corresponding intersection for j . From (2.4.b) and (2.4.d) above, the difference between $F_i^{2,1}(t)$ and $F_j^{2,1}(t)$ is $(1-t)(\Pi_i^{2,0} - \Pi_j^{1,1})$; From (2.4.a) and (2.4.c), the difference between $L_i^{2,1}(t)$ and $L_j^{2,1}(t)$ is $(1-t)\Pi_j^{1,1}$. (These differences are illustrated in Figure 1 for $\pi^* = 2$ and $c^* = 1$.) To establish that the former difference exceeds the latter, therefore, we need only show that $\Pi_i^{2,0} > \Pi_i^{1,1} + \Pi_j^{1,1}$, i.e., that the payoff that a monopolist with two bullets receives strictly exceeds the *aggregate* payoff for the two players in phase (1,1). But this inequality must be satisfied, because the monopolist fires at the optimal times for an industry with two bullets, while in the duopoly (1,1), the first bullet is fired too early. More precisely, we have:

$$\begin{aligned}
 \Pi_i^{1,1} + \Pi_j^{1,1} &= L_i^{1,1}(t^{1,1}) + F_i^{1,1}(t^{1,1}) & (2.5) \\
 &= t^{1,1}\pi^*(2 - t^{1,1}) - c^* + (1 - t^{1,1})(\pi^* - c^*) \\
 &\equiv t^{1,1}\pi^*(2 - t^{1,1}) - c^* + (1 - t^{1,1})\Pi_i^{1,0} \\
 &= L_i^{2,0}(t^{1,1}) < L_i^{2,0}(t^{2,0}) \equiv \Pi_i^{2,0}.
 \end{aligned}$$

The first three equalities follow, respectively from equations (2.2), (2.3) and (2.1). The inequality holds because $t^{2,0}$ is the unique maximizer of $L_i^{2,0}(\cdot)$.




FIGURE 1

The generic phase of the problem.

Fix a pair β , where, for convenience, we assume $\beta_i > \beta_j + 1$. Assume that for every pair $(\beta'_i, \beta'_j) \leq \beta$, there is a common, unique solution for each of the classical duels corresponding to phase (β'_i, β'_j) . (We will return to this assumption at the end of the section.) Let $t^{\beta'}$ and $\Pi^{\beta'_i, \beta'_j}$ denote, respectively, the firing time and equilibrium payoff vector for phase (β'_i, β'_j) . An important fact is that $t^{\beta_i-1, \beta_j} < t^{\beta_i, \beta_j-1}$; that is, the bullet is fired sooner in the less asymmetric phase than the more asymmetric one.

Our analysis of this phase exactly parallels our analysis of phase (2,1). All of the classical duels corresponding to this phase have identical payoffs on the interval $[0, t^{\beta_i-1, \beta_j}]$.¹¹ Moreover, there is a common, unique solution for each of these duels, which is completely determined by payoffs on the above interval. Once again, therefore, we will refer to *the* duel corresponding to phase β , and will denote it as CD^β . For each $t < t^{\beta_i-1, \beta_j}$, the lead and follow functions for each player are defined as follows:

$$L_i^\beta(t) = t\pi^*(2-t) - c^* + (1-t)\Pi_i^{\beta_i-1, \beta_j} \quad (2.6.a)$$

$$F_i^\beta(t) = (1-t)\Pi_i^{\beta_i, \beta_j-1} \quad (2.6.b)$$

$$L_j^\beta(t) = t\pi^*(2-t) - c^* + (1-t)\Pi_j^{\beta_i, \beta_j-1} \quad (2.6.c)$$

$$F_j^\beta(t) = (1-t)\Pi_j^{\beta_i-1, \beta_j} \quad (2.6.d)$$

The duels corresponding to *this* phase will all share a common, easily computible solution if:

$$\text{for each } t, \text{ the functions } L_i^\beta(t) \text{ and } F_i^\beta(t) \text{ intersect before } t^{\beta_i, \beta_j-1} \quad (*)$$

For the moment, we will assume that (*) is satisfied and return to this assumption at the end of the section.

To establish that player j fires in this phase, we need to show that throughout $[0, t^{\beta_i-1, \beta_j}]$:

$$F_i^\beta(\cdot) - F_j^\beta(\cdot) > L_i^\beta(\cdot) - L_j^\beta(\cdot). \quad (2.7)$$

In this general case, it is less easy to see why this inequality should hold. Once we have established (2.7),

¹¹ See footnote 10 above. The equilibria for phase (β_i-1, β_j) all agree before t^{β_i-1, β_j} ; similarly, for phases (β_i, β_j-1) and (β_i-1, β_j-1) , the equilibria all agree before, respectively, t^{β_i, β_j-1} and t^{β_i-1, β_j-1} . t^{β_i-1, β_j} is the smallest of these three times.

however, we can conclude as before that i 's intersection occurs later than j 's, so that, as in phase (2,1), j fires the bullet at i 's intersection time.

From (2.6.b)-(2.6.d) above, the difference between $F_i^\beta(t)$ and $F_j^\beta(t)$ is $(1-t)(\Pi_i^{\beta_i, \beta_j-1} - \Pi_j^{\beta_i-1, \beta_j})$; from (2.6.a) and (2.6.c), the difference between $L_i^\beta(t)$ and $L_j^\beta(t)$ is $(1-t)(\Pi_i^{\beta_i-1, \beta_j} - \Pi_j^{\beta_i, \beta_j-1})$; To establish (2.7), therefore, we need only show that $\Pi_i^{\beta_i, \beta_j-1} + \Pi_j^{\beta_i, \beta_j-1} > \Pi_i^{\beta_i-1, \beta_j} + \Pi_j^{\beta_i-1, \beta_j}$. That is, (2.7) will hold if and only if aggregate profits in phase (β_i, β_j-1) exceed those in phase (β_i-1, β_j) . Now since $\beta_i > \beta_j$, the two players are less evenly matched in the former phase than in the latter. It seems quite likely, therefore, that aggregate payoffs will be higher in the former phase. The intuition is that when players are less evenly matched, competition between them is less intense, so that downward pressure on profits is weaker. Indeed, this intuition is consistent with several studies of related problems in the literature.¹²

While this intuition is persuasive, it is by no means straightforward to verify. The proof is deferred to the Appendix. What follows is a very brief outline of the argument. We first establish that in every duopoly phase of the duel, the bullet is fired earlier than in the corresponding monopoly phase. That is, for each $\beta_i' \geq \beta_j' > 0$, the firing time, $t^{\beta_i'}$, in phase (β_i', β_j') occurs strictly earlier than the time, $t^{\beta_i'+\beta_j', 0}$, that would maximize aggregate profits. As we have observed, this seems as though it should be an easy fact to establish. In fact, it is the most difficult step in the proof, and the only one for which we have no economic intuition!

Assume that we have established that player j fires the bullet in phases (β_i, β_j-1) and (β_i-1, β_j) . (Since $\beta_i > \beta_j + 1$, j still has fewer bullets in phase (β_i-1, β_j)). We first argue that the firing time in phase (β_i, β_j-1) will be later than it is in phase (β_i-1, β_j) . By assumption, j fires in each of these phases at the moment that i 's lead and follow functions intersect. We need to establish, therefore, that i 's intersection will occur later in phase (β_i, β_j-1) than in phase (β_i-1, β_j) . The intuition runs as follows. In the latter phase, j is "one miss closer" to running out of bullets than he is in the former. Moreover, i has one more bullet at his disposal in the latter phase. Consequently, following is even more attractive relative to leading in phase (β_i, β_j-1) than it is in phase (β_i-1, β_j) . Therefore, the time at which $F_i^{\beta_i, \beta_j-1}(\cdot)$ intersects $L_i^{\beta_i, \beta_j-1}(\cdot)$ occurs

¹² See, for example, Fudenberg et. al. [3], Harris-Vickers [5], Grossman-Shapiro [4].

strictly later than the time at which $F_i^{\beta_i-1, \beta_j}(\cdot)$ intersects $L_i^{\beta_i-1, \beta_j}(\cdot)$. It now follows that t^{β_i, β_j-1} occurs later than t^{β_i-1, β_j} . The final step of the argument is immediate. The firing time in phase (β_i-1, β_j) occurs sooner than the optimal firing time in the corresponding monopoly phase $(\beta_i+\beta_j-1, 0)$. The monopoly phase that corresponds to phase (β_i, β_j-1) is also $(\beta_i+\beta_j-1, 0)$, so that the optimal firing time for phase (β_i-1, β_j) is also $t^{\beta_i+\beta_j-1, 0}$. That is, t^{β_i-1, β_j} is even earlier than t^{β_i, β_j-1} while t^{β_i, β_j-1} is already too early. Moreover, aggregate profits are a strictly concave function of time. Therefore, they must be lower in phase (β_i-1, β_j) than in phase (β_i, β_j-1) .

A uniqueness test for multi-stage duels.

To conclude this section, we return to the issue of whether or not the classical duels associated with phase β will have unique solutions. The answer depends on whether or not condition (*) defined on p. 8 above will be satisfied for this phase. It is easy to check that for any β_i , this condition will always be satisfied in phase $(\beta_i, 1)$. For β such that $\beta_i > \beta_j > 1$, however, assumption (*) will not always hold. That is, there are parameter values such that when i 's payoffs to leading and following are defined according to (2.6.a)-(2.6.b) above, $F_i^\beta(\cdot)$ will strictly exceed $L_i^\beta(\cdot)$ throughout the interval $[0, t^{\beta_i-1, \beta_j})$. When this happens, we will be unable to solve phase β without taking into account agents' continuation payoffs if the phase continues beyond t^{β_i, β_j-1} . These values will not be uniquely defined, and the our task will be much more complicated. Accordingly, we will restrict our attention to the range of parameters for which (*) is satisfied.

We have been unable to identify this range by analytic methods. There is, however, a simple numerical test that we can perform to check the condition. This test can easily be run on a computer. We have in fact run it many times, and verified that for a wide range of parameter values, condition (*) will indeed be satisfied.¹³ Whenever it is, we can follow the inductive procedure described above and construct the unique solution for our multistage game.

¹³ The condition is more likely to be satisfied for given β , the smaller is c^* relative π^* . Indeed, when $c^* = 0$, the condition was satisfied for all numbers that we tested (up to twenty bullets each).

The easiest way to explain our test is by analogy. Faced with a constrained maximization problem, the first step is to solve the corresponding unconstrained problem. If the latter problem has a solution that lies inside the original constraint set, then this solution will also solve the constrained problem. Our constrained problem that does not involve maximization, but we *are* looking for an interior solution. The constraint in our problem is condition (*) on p. 8. The algorithm below solves the sequence of problems defined by simply ignoring this constraint. If the "test firing times" defined by this algorithm all satisfy constraint (*), then these will also be the unique firing times that solve *our* problem. We now state the test.

Fix a triple (β^*, π^*, c^*) . For each β_i , define the vectors $\bar{\Pi}^{\beta_i, 0}$ and $\bar{\Pi}^{0, \beta_i}$ to be the solutions to the monopoly phases $(\beta_i, 0)$ and $(0, \beta_i)$ of the corresponding multistage duel. That is, set $\bar{\Pi}^{\beta_i, 0}$ equal to the vector $\Pi^{\beta_i, 0}$ defined on p. 4. Now fix $\beta \leq \beta^*$ and assume that $\bar{\Pi}^{\beta_i, -1, \beta_j}$ and $\bar{\Pi}^{\beta_i, \beta_j, -1}$ have been defined.¹⁴ For $i = i, j$, define the test payoff functions $\bar{L}_i^\beta(\cdot)$ and $\bar{F}_i^\beta(\cdot)$ on the interval $[0, 1)$ as follows.

$$\bar{L}_i^\beta(t) = t\pi^*(2-t) - c^* + (1-t)\bar{\Pi}_i^{\beta_i, -1, \beta_j} \quad (2.6.a')$$

$$\bar{F}_i^\beta(t) = (1-t)\bar{\Pi}_i^{\beta_i, \beta_j, -1} \quad (2.6.b')$$

$$\bar{L}_j^\beta(t) = t\pi^*(2-t) - c^* + (1-t)\bar{\Pi}_j^{\beta_i, \beta_j, -1} \quad (2.6.c')$$

$$\bar{F}_j^\beta(t) = (1-t)\bar{\Pi}_j^{\beta_i, -1, \beta_j} \quad (2.6.d')$$

For each t , let \bar{t}_i^β be the solution to $\bar{L}_i^\beta(\cdot) - \bar{F}_i^\beta(\cdot) = 0$. The left hand side is a quadratic with a unique solution on $[0; 1)$. Now set $\bar{t}^\beta = \max(\bar{t}_i^\beta, \bar{t}_j^\beta)$. Finally, for each t , define $\bar{\Pi}_i^\beta = \bar{L}_i^\beta(\bar{t}^\beta)$. Call \bar{t}^β the test firing time for phase β and $\bar{\Pi}^\beta = (\bar{\Pi}_i^\beta, \bar{\Pi}_j^\beta)$ the test payoff vector for phase β . Proceed in this way to construct $\bar{t}^{\beta'}$, for every $\beta' \leq \beta^*$.

Suppose that each of the test firing times defined above satisfy condition (*). Let f be any equilibrium profile for the multistage duel with parameters (β^*, π^*, c^*) . Pick any duopoly phase $\beta \leq \beta^*$. We will show

¹⁴ More specifically, we proceed inductively as follows. We first determine $\bar{\Pi}^{1,1}$. For $\beta_i > 1$, once $\bar{\Pi}^{\beta_i, -1, 1}$ and $\bar{\Pi}^{\beta_i, 0}$ have been determined, we have enough information to determine $\bar{\Pi}^{\beta_i, 1}$. Now fix β_j and assume that we have determined $\bar{\Pi}^{\beta_i, \beta_j, -1}$, for every β_i . In particular, we will have determined $\bar{\Pi}^{\beta_i, \beta_j, -1}$ and $\bar{\Pi}^{\beta_i, -1, \beta_j}$ so we can determine $\bar{\Pi}^{\beta_i, \beta_j}$. Finally, for $\beta_i > \beta_j$, once $\bar{\Pi}^{\beta_i, -1, \beta_j}$ and $\bar{\Pi}^{\beta_i, \beta_j, -1}$ have been determined, we can proceed as follows.

that the payoff functions for the phase β classical duel defined by playing f beyond phase β coincides on $[0, \tilde{t}^{\beta, -1, \beta_j})$ with the "test payoff functions" defined by (2.6.d')-(2.6.d') above. Moreover, we will show that when f is played in this phase starting from $t \leq \tilde{t}^\beta$, the bullet will be fired with probability one at \tilde{t}^β . Since this will be true for every β , we have the following result:

Prop'n I: Consider the triple (β^*, π^*, c^*) , where $\beta^* = \beta^*$ and $\beta_i^* \geq \beta_j^*$. Define the family of scalars $(\tilde{t}^{\beta'})_{(\beta_i, \beta_j) \leq \beta^*}$ according to the algorithm above. If $\tilde{t}^\beta < \tilde{t}^{\beta_i - 1, \beta_j}$, for all $1 \leq \beta_j \leq \beta_j^*$ and all $\beta_j \leq \beta_i \leq \beta_i^*$, then there is a unique equilibrium payoff vector for the multi-stage duel defined by (β^*, π^*, c^*) , which coincides with the test payoff vector $\tilde{\Pi}^{\beta^*}$.

For the remainder of the paper, we will restrict our attention to multistage duels that pass our test.

Results.

Since we have not yet defined our continuous time game, our results are stated in a nontechnical way.

Our main result is:

Th'm II: Consider the multistage duel with parameters (β^*, π^*, c^*) . Assume that these parameters pass the test specified in Proposition I. If $\beta_i^* > \beta_{-i}^*$, the game has a unique equilibrium outcome, generating the payoff vector Π^{β^*} , defined inductively as above. Moreover, τ will fire a bullet only if $-\tau$ has already fired all of his. That is, the phases of the game that will be reached with positive probability are $((\beta_i^*, \beta_{-i}^*), (\beta_i^*, \beta_{-i}^* - 1), \dots, (\beta_i^*, 0), (\beta_i^* - 1, 0), \dots, (1, 0))$. If phase β is reached, the bullet will be fired at t^β , defined inductively as above. If $\beta_i^* = \beta_j^*$, then there are three equilibrium outcomes. For each τ , there is an equilibrium in which τ never fires a bullet unless $-\tau$ has fired all of his. There is a third one in which, for each τ , with probability $1/2$, τ never fires a bullet until $-\tau$ has fired all of his.

We outlined a nontechnical version of the proof of this result in the preceding section. The details are deferred to the Appendix.

In order to prove Theorem II, we established another result. Since we believe it to be of independent interest, we will state it as separate theorem. It compares the efficiency of two duels that are identical except for the initial relative strengths of the two agents. It states that if one bullet is transferred from the weak player to the strong player, then competition will be less intense and *aggregate* efficiency will increase.

Th'm III: Consider two multistage duels with parameters (β^*, π^*, c^*) and $((\beta_i^*+1, \beta_j^*-1), \pi^*, c^*)$, where $\beta_i^* \geq \beta_j^*$. Aggregate expected payoffs are higher in the second duel. Moreover, each bullet in the second duel is fired later than the corresponding bullet in the first. That is, for each $1 \leq \beta_j \leq \beta_j^*$, $t^{\beta_i^*+1, \beta_j-1} > t^{\beta_i^*, \beta_j}$.

In our next two sections, we develop our continuous time model. We introduce our methodology in section III, restricting attention to the classical duel. The presentation in this section is informal. In the following section, we will introduce our multistage model, adopting a much more formal approach.

III. Introducing our continuous-time model: the classical duel

We will restrict attention in this section to the particular family of classical duels that are relevant for the purposes of this paper. The family is identified by conditions (A1)-(A4), defined below. These conditions are rather special. The reason is that they arise "endogenously" in the course of solving our multistage game. The conditions are:

For each τ , $L_\tau(\cdot)$, $F_\tau(\cdot)$ and $S_\tau(\cdot)$ are right continuous, piecewise polynomial functions.¹⁵ (A1)

For each τ , there exists $t_\tau \in [0, 2)$ such that (A2)

before t_τ , $L_\tau(\cdot)$ and $F_\tau(\cdot)$ are continuous and $L_\tau(\cdot)$ lies strictly below $F_\tau(\cdot)$; (A2.a)

beyond t_τ , $F_\tau(\cdot)$ lies strictly below $L_\tau(\cdot)$ (A2.b)

and for all $s > t_\tau$, $\lim_{\delta \downarrow 0} L_\tau(s-\delta) > F_\tau(s)$. (A2.c)

There exists $T > \max(t_i, t_j)$ such that for each τ , $L_\tau(\cdot)$ is strictly increasing on $[0, T)$. (A3)

$S_\tau(\cdot)$ lies strictly below $F_\tau(\cdot)$ on $[0, 2)$; (A4)

Assumption (A1) is purely technical and is trivially satisfied. Assumptions (A3), (A2.a) and (A2.b) are natural in the context of the duel: as time progresses, τ 's accuracy increases, and so he becomes increasingly more likely to score a hit if he fires. Hence, $L_\tau(\cdot)$ is strictly increasing at the start of the game. At the very beginning of the game, both players will be sufficiently inaccurate that each will prefer to be fired at than to fire; in this region of the game, $F_\tau(\cdot)$ will exceed $L_\tau(\cdot)$. Eventually, however, both players will become sufficiently accurate that τ prefers to fire than be fired at. At this point $L_\tau(\cdot)$ will overtake $F_\tau(\cdot)$. It is harder at this stage

¹⁵ A function θ is piecewise polynomial on $[0, 2)$ if there exists a finite subset, $\{0 = t^0 < \dots < t^r < \dots < t^{last} = 2\}$, of this interval such that for each $r < last$, the restriction of θ to (t^r, t^{r+1}) is a polynomial.

to explain why $L_i(\cdot)$ should be continuous and strictly increasing until *after* it is overtaken by $F_i(\cdot)$. Condition (A2.c) is simply a strengthening of the condition that $L_i(\cdot)$ exceeds $F_i(\cdot)$. Since either function may be discontinuous beyond t_i , the condition is not implied by (A2.a). Finally, (A4) is a very delicate assumption. We have to choose our equilibrium strategies very carefully to ensure that it is satisfied.

We first review how a duel in this class would be modelled in discrete time.¹⁶ Let R denote a discrete-time grid, i.e., a finite subset of $[0, 2)$. A discrete time behavior strategy for agent i is a function, ξ_i^R , that assigns a probability weight to each point in R . ξ_i^R has the following interpretation: "for each $r \in R$, if nobody has fired by the time r is reached, I will play 'fire' at this time with probability $\xi_i^R(r)$." A pair of strategies, one for each agent, will be called a strategy profile. When any strategy profile is played, starting from any grid-point in R , it will generate by induction a unique probability distribution over endings, which we shall call an outcome.

The model just described has no direct analog in continuous-time. In particular, it is not obvious that one can sensibly define a continuous-time behavior strategy. The "natural" counterpart of ξ_i^R would be a function ξ_i , mapping each point in $[0, 2)$ to a probability weight in $[0, 1]$. But how should this function be interpreted? The literal interpretation would be: "for each $t \in [0, 2)$, if nobody has fired by the time t is reached, I will play 'fire' at this time with probability $\xi_i(t)$." In order to "play" this strategy, however, an agent would have to be able to perform a continuum of independent randomizations. It is well known that this is not an easy thing to do.

Our continuous-time model is specified in a way that finesses this difficulty.¹⁷ Our agents do choose behavior strategies that have the form just described. That is, each behavior strategy for the classical duel is a function mapping $[0, 2)$ to $[0, 1]$.¹⁸ We will, however, propose a novel interpretation of these strategies, together with a novel way to associate strategy profile to outcomes.¹⁹ We interpret a continuous-time

¹⁶ Dixit [1] contains an analysis of the discrete-time duel.

¹⁷ There are many other difficulties that need to be addressed. These are discussed at length in [9] and [10].

¹⁸ This specification is actually quite restrictive. See Simon [9], pp. 7-8 for a discussion of this point. A more general class of strategies will be proposed in Simon-Stinchcombe [11].

¹⁹ The model we present below was motivated by Fudenberg-Tirole [2], in which an alternative model is proposed.

behavior strategy as a set of instructions about how to play the game on every conceivable discrete-time grid. Specifically, the restriction of any continuous-time profile to a discrete-time grid will be a well-defined discrete-time profile. When played starting from any grid point, this profile will define a unique discrete-time outcome, in exactly the conventional way. When strategies are interpreted in this way, there is a natural candidate procedure for defining outcomes: (a) fix a starting point t in $[0, 2)$ and a continuous-time profile; (b) restrict these strategies to an arbitrary, increasingly fine sequence of discrete-time grids and play the restricted profiles, starting from t ; (c) define the continuous time outcome generated by playing this profile from t to be the limit of these discrete-time outcomes.

In general, this procedure may break down. For some pairs of strategies, the limit of discrete-time outcomes may not exist. For others, it may exist, but depend on the particular sequence of grids on which the strategies are played. Accordingly, for our procedure to be coherent, we need to identify a class of strategies with the following property: when any two members of this class are paired together and played from some starting point on an arbitrary sequence of grids, the resulting sequence of outcomes must converge to a unique limit that is independent of the particular sequence of grids. 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 classical duel.

For a problem as simple as the classical duel, the restrictions on strategies that we need are minimal. It suffices to require that agents choose 'piecewise rational' behavior strategies. (A rational function is the ratio of two polynomial functions.) We will also impose a second restriction, purely because it considerably simplifies the exposition. We will require that if an agent is properly randomizing at t (i.e., he both waits and fires with positive probability), then his strategy is right continuous at t . It will become apparent that in our present context, this restriction is completely innocuous.

Our continuous-time duel has a continuum of 'subgames,' one for each $t \in [0, 2)$. The subgame beginning at t is simply the subduel of the original duel that begins at t . The natural solution concept for our model is subgame perfection. A pair of strategies will be a subgame perfect equilibrium for a

¹⁹ This solution concept was originally proposed in Selten [8].

continuous-time classical duel if the strategies form a Nash equilibrium for every subduel.

Equilibrium strategy profiles for a symmetric classical duel.

In this section we will construct equilibrium strategies for a symmetric classical duel. Specifically, we will consider the duel, $CD^{1,1}$, corresponding to phase (1,1). When $t \in [1, 2)$, an agents who fires scores a hit with probability one. Consequently, the payoffs to leading, following and moving simultaneously are $L_i^{1,1}(t) = \pi^*(2 - t)$, $F_i^{1,1}(t) = 0$ and $S_i^{1,1}(t) = -c^*$. This part of the game is completely uninteresting and we will henceforth ignore it. If i fires and misses at $t < t^{1,0} = 1$, player $-i$ will wait until time 1, fire at point blank range and earn the payoff $(\pi^* - c^*)$. If both players fire simultaneously at $t < 1$, then i wins with the probability $t(1 - t)$ that his bullet hits and $-i$'s misses. Thus, for $t \in [0, 1)$, we have

$$L_i^{1,1}(t) = t\pi^*(2 - t) - c^* \tag{3.1.a}$$

$$F_i^{1,1}(t) = (1 - t)(\pi^* - c^*) \tag{3.1.b}$$

$$S_i^{1,1}(t) = t(1 - t)\pi^*(2 - t) - c^* \tag{3.1.c}$$

It is easy to check that conditions (A1)-(A4) above are indeed satisfied. In particular, as we claimed on p. 5,

$L_i^{1,1}(\cdot)$ intersects $F_i^{1,1}(\cdot)$ at $t^{1,1} \in (0, 1)$. As usual, we call $t^{1,1}$ the firing-time for this duel.

In any equilibrium for this game, player i earns the payoff $L_i^{1,1}(t^{1,1})$. There is a unique *symmetric* equilibrium profile that generates these payoffs. The strategies, denoted by $\xi^0 = (\xi_i^0, \xi_j^0)$, are defined as follows:

$$\text{low: } \xi_i^0(t) = \begin{cases} 0 & \text{if } t \in [0, t^{1,1}) \\ \frac{L_i(t) - F_i(t)}{L_i(t) - S_i(t)} & \text{if } t \in [t^{1,1}, 2) \end{cases} . \text{ Assumptions (A2.b) and (A4) guarantee that } \xi_i^0(\cdot) \in (0, 1)$$

beyond $t^{1,1}$.

We first describe the outcome generated by these strategies, when outcomes are defined as on pp. 14-15. The outcome function will be defined formally in section III below. If these strategies are played starting from any $t < t^{1,1}$, the outcome is that with probability 1, exactly one player fires at exactly $t^{1,1}$. Each agent is equally likely to fire. (We emphasize that players fire simultaneously with probability zero. That is, the limit outcome generated by these strategies is not a product distribution!²⁰) We will show that

²⁰ An almost universal source of confusion is that this outcome *looks* like a "correlated equilibrium." This is misleading. Agents

this is indeed the limit of the outcomes generated by restricting ξ^0 to any increasingly fine sequence of grids. Fix $\varepsilon > 0$, and let $\gamma = \xi_i^0(t^{1,1} + \frac{1}{2}\varepsilon)$. Clearly, γ is strictly positive. Now pick a very fine discrete-time grid and play ξ^0 on this grid. Since ξ_i^0 is strictly increasing beyond $t^{1,1}$, each agent will fire with probability at least γ at each grid point in the interval $(t^{1,1} + \frac{1}{2}\varepsilon, t^{1,1} + \varepsilon)$, provided that the game has not ended by the time this grid point is reached. If this interval contains at least n grid points, then the probability that the end of this interval will be reached without a shot being fired will be no more than $(1 - \gamma)^{2n}$. If n is sufficiently large, this probability will be less than ε . We have verified, therefore, that for every positive ε , if ξ^0 is played on a sufficiently fine grid, the probability that some agent will have moved before $t^{1,1} + \varepsilon$ will exceed $1 - \varepsilon$. It follows that in the limit, the bullet must be fired with probability one at $t^{1,1}$. Next, note that if ε is sufficiently small, the probability that both agents fire simultaneously will be arbitrarily small relative to the probability that only one agent fires. In the limit, therefore, the probability that agents fire simultaneously at $t^{1,1}$ must be zero. Finally, note that by definition of $t^{1,1}$, each agent is indifferent between leading and following at this time. Therefore, the payoff vector that is generated by playing ξ^0 , starting from before $t^{1,1}$, is $(L_i^{1,1}(t^{1,1}), L_i^{1,1}(t^{1,1}))$.

Now suppose that the profile ξ^0 is played starting from $t > t^{1,1}$. With the help of a formula that we provide in the following section, it is easy to compute the outcome in this case. With probability one, the bullet will be fired exactly at t . Player i fires alone with probability $\frac{1 - \xi_i^0(t)}{2 - \xi_i^0(t)}$. With probability $\frac{\xi_i^0(t)}{2 - \xi_i^0(t)}$ both players fire simultaneously. Some tedious algebra verifies that player i 's expected payoff from this outcome is $F_i^{1,1}(t)$.

We next argue informally that these strategies are best responses to each other in every subgame. First fix $\tau \leq t^{1,1}$ and consider the subgame that begins at τ . Since $L_i^{1,1}(\cdot)$ is strictly increasing before $t^{1,1}$, i cannot gain by shooting before this time. If i fires at exactly $t^{1,1}$, his payoff will be $L_i^{1,1}(t^{1,1})$, since $-i$ fires with probability zero at this time. This deviation yields him no more than his original payoff. Now suppose that

are randomizing independently. The limit of a sequence of product distributions need not be a product distribution.

for some ε , τ fires with probability zero between τ and $\tau^{1,1} + \varepsilon$, and then fires with probability one. We will describe what happens when this strategy is played against $\xi_{-\tau}^0$ on an increasingly fine sequence of grids. Just as before, τ will fire with probability γ at each of an increasingly large number of grid points in the interval $(\tau^{1,1} + \frac{1}{2}\varepsilon, \tau^{1,1} + \varepsilon)$. In the limit, τ must fire with probability one at $\tau^{1,1}$. Therefore, this deviation by τ earns him a payoff of $F_i^{1,1}(\tau^{1,1})$, which is exactly what he was getting originally. This completes our argument that ξ^0 is Nash equilibrium for each subgame that begin before $\tau^{1,1}$.

The argument is similar for subgames that begin beyond $\tau^{1,1}$. Fix $\tau > \tau^{1,1}$. Since τ is randomizing at τ , we need to verify that he is indifferent between firing with probability one at τ and waiting until $\tau + \varepsilon$, and then firing with probability one. If τ fires with probability zero on $[\tau, \tau + \varepsilon)$, then by exactly the argument given above, the outcome will be that τ will fire with probability one at τ . In this case, τ 's payoff will be $F_i^{1,1}(\tau)$. Now suppose that τ fires with probability one at τ . In this case, he will earn $S_i^{1,1}(\tau)$ with the probability that τ also fires at this time, and $L_i^{1,1}(\tau)$ with the remaining probability. His expected payoff, therefore, is

$$S_i^{1,1}(\tau) + (1 - \xi_i^0(\tau))L_i^{1,1}(\tau) = \frac{L_i^{1,1}(\tau) - F_i^{1,1}(\tau)}{L_i^{1,1}(\tau) - S_i^{1,1}(\tau)} S_i^{1,1}(\tau) + \frac{F_i^{1,1}(\tau) - S_i^{1,1}(\tau)}{L_i^{1,1}(\tau) - S_i^{1,1}(\tau)} L_i^{1,1}(\tau) = F_i^{1,1}(\tau).$$

Thus, we have verified that τ cannot do better by playing any *pure* strategy other than ξ_i^0 and thus established that ξ^0 is a Nash profile for the subgame beginning at τ .²¹

There are two other equilibrium outcomes for this duel. Each of them generates the same payoff vector as ξ^0 . For each τ , there is an equilibrium in which τ fires at $\tau^{1,1}$ with probability one. A strategy profile that implements one of these is ξ^1 , which is identical to ξ^0 except that $\xi_i^1(\tau^{1,1}) = 1$.²² We have already verified

²¹ If we were working in discrete-time, it would go without saying that we need only check pure-strategy deviations. It is also true in continuous time, but no longer self-evident. The reader is referred to Proposition VII of Simon [9] for formal details.

²² Because ξ_i is discontinuous w.r.t. time, the interpretation of this strategy is a subtle matter. If we simply restricted this profile to an arbitrary grid, the outcome would depend critically on whether or not the grid contained the discontinuity point $\tau^{1,1}$. To avoid this dependence, we "adapt" the profile to each grid, ensuring that if the grid is sufficiently fine, each of the profile's discontinuities will be "captured" by the grid. Specifically, we define an operator called a "graph preserving restriction" (g.p.r.), which "shifts to the right" the discontinuities of a profile in the appropriate way. For example, the g.p.r. of ξ^1 to the grid R has player i terminating with probability one at the first grid point in R greater than or equal to $\tau^{1,1}$. For details see Simon Stinchcombe [10], [11].

that ξ_1^1 is a best response against $\xi_{-1}^1 \equiv \xi_{-1}^0$. To see that ξ_{-1}^1 is a best response against ξ_1^1 , observe that if -1 fires with probability one at $t^{1,1}$, he will earn $S_1^{1,1}(t^{1,1})$, which is strictly less than $F_1^{1,1}(t^{1,1})$.

Not surprisingly, there are countless other equilibrium profiles for this duel. To construct one, take, say, ξ^0 and modify it as follows: at each of a finite number of times beyond $t^{1,1}$, have 1 fire with probability one and -1 fire with probability zero. By rearranging the arguments given above, it is straightforward to verify that the modified profile is an equilibrium. An important fact, however, is that for *any* equilibrium profile, ξ , there must exist $\epsilon > 0$ such that ξ will agree with *one* of the three profiles defined above on the interval $[0, t^{1,1} + \epsilon)$.²³

A striking property of the symmetric duel just described is that all three of its equilibrium outcomes are payoff equivalent and yield each player equal payoffs. Not surprisingly, this result depends on the fact that agents started playing the duel strictly before the bullet was due to be fired. Once this inequality is reversed, these properties no longer hold. For example, pick $\tau > t^{1,1}$ and consider the subduel $CD^{1,1}(\tau)$, of $CD^{1,1}$ that begins at τ . It is straightforward to verify that this duel has three equilibrium outcomes, each generating a distinct payoff vector. Specifically, for each τ , there is an equilibrium in which 1 fires with probability one, and -1 waits with probability one, at τ . Obviously, 1 's payoff in this equilibrium is $L_1^{1,1}(\tau)$, while -1 's is $F_1^{1,1}(\tau)$. In the third kind of equilibrium, strategies must coincide with ξ^0 both at and immediately after τ . The payoff vector corresponding to this equilibrium is $(F_1^{1,1}(\tau), F_1^{1,1}(\tau))$. As we have noted, this multiplicity of solutions can significantly complicate our analysis.

The asymmetric classical duel.

The analysis of an asymmetric duel is similar to the one just given except for two differences. First, as we asserted when we were studying $CD^{2,1}$, if agents start playing in an asymmetric duel *before* its firing time, there will be only one equilibrium outcome, rather than three. If 1 's lead function intersects his follow function while -1 still strictly prefers following to leading, then the unique outcome will be that 1 fires with probability one at -1 's intersection time. The second difference concerns the intensity with which agents

²³ To prove this from scratch is a little tedious. It follows almost immediately, however, from Proposition VII of Simon [9].

randomize beyond the firing time. At each t , i must fire with sufficiently high probability that $-i$ is indifferent between firing and waiting. It is easy to check that if agents are randomizing at t beyond the firing time of an asymmetric duel, then i must fire with probability $\frac{L_{-i}(t) - F_{-i}(t)}{L_{-i}(t) - S_{-i}(t)}$.

Summary: equilibrium strategies for the classical duel.

The results described above will be used repeatedly in the analysis that follows. For future reference, we will tabulate them below in the form of a proposition. Let \overline{CD} denote the classical duel with payoffs $(\overline{L}_i, \overline{F}_i, \overline{S}_i)_{i=j}$. Assume that these functions satisfy assumptions (A1)-(A4) (p. 13 above). Let \overline{t}_i denote the time at which \overline{L}_i and \overline{F}_i intersect. By assumption A2, \overline{t}_i exists. Let \overline{t} denote the firing time for $\overline{CD}(\tau)$, i.e., the maximum of \overline{t}_i and \overline{t}_j . Let $\overline{CD}(\tau)$ denote the subduel of \overline{CD} that begins at τ .

We define three "canonical profiles," $\xi^{\tau,0}$, $\xi^{\tau,i}$ and $\xi^{\tau,j}$, for the subduel $\overline{CD}(\tau)$. All three are identical except for their values at the maximum of \overline{t} and τ : at this time, $\xi^{\tau,i}$ has i firing with probability one, and $-i$ firing with probability zero. $\xi^{\tau,0}$ is symmetric.

First, define $\xi^{\tau,0}$ by for each t , $\xi^{\tau,0}(t) = \begin{cases} 0 & \text{if } t \in [\tau, \overline{t}) \\ \frac{\overline{L}_{-i}(t) - \overline{F}_{-i}(t)}{\overline{L}_{-i}(t) - \overline{S}_{-i}(t)} & \text{if } t \geq \max(\tau, \overline{t}) \end{cases}$. Now, for each i

and each $\chi \in \{i, j\}$, define $\xi^{\tau,\chi}$ by $\xi^{\tau,\chi}(t) = \begin{cases} 0 & \text{if } t \in [\tau, \overline{t}) \\ 1 & \text{if } t = \max(\tau, \overline{t}) \text{ and } \chi = i \\ 0 & \text{if } t = \max(\tau, \overline{t}) \text{ and } \chi = -i \\ \frac{\overline{L}_{-i}(t) - \overline{F}_{-i}(t)}{\overline{L}_{-i}(t) - \overline{S}_{-i}(t)} & \text{if } t > \max(\tau, \overline{t}) \end{cases}$. The

following result describes equilibrium strategies, outcomes and payoff vectors for the subduels of symmetric and asymmetric classical duels.

Prop'n IV: Let \overline{CD} satisfy assumptions (A1)-(A4).

- (i): If $\tau \leq \bar{t}$ and $\bar{t}_i = \bar{t}_j$, then the subduel $\overline{CD}(\tau)$ has exactly three distinct equilibrium outcomes, all of which generate the payoff vector $(\bar{L}_i(\bar{t}), \bar{L}_j(\bar{t}))$. For each $\chi \in \{0, i, j\}$, $\xi^{\tau, \chi}$ is an equilibrium strategy profile.
- (ii): If $\tau < \bar{t}$ and $\bar{t}_j < \bar{t}_i$, then $\overline{CD}(\tau)$ has a unique equilibrium outcome, yielding payoffs $(\bar{L}_i(\bar{t}), \bar{L}_j(\bar{t}))$. $\xi^{\tau, j}$ is an equilibrium profile.
- (iii): For each $\tau \geq \bar{t}$, $\overline{CD}(\tau)$ has exactly three distinct solutions and exactly three distinct equilibrium outcomes. For each $\chi \in \{0, i, j\}$, $\xi^{\tau, \chi}$ is an equilibrium for the duel $\overline{CD}(\tau)$. $\xi^{\tau, 0}$ yields the outcome $(\bar{F}_i(\bar{t}), \bar{F}_j(\bar{t}))$. $\xi^{\tau, i}$ yields the outcome $(\bar{L}_i(\bar{t}), \bar{F}_{-i}(\bar{t}))$.

The existence parts of this result follow immediately from Proposition V, Proposition VII and Theorem VIII in Simon [9]. We do need to prove uniqueness, however. Since the proof is instructive, we will include it in the text. The proof of part (iii) is a straightforward extension of the other proofs, and is omitted.

First note that for any t , if i fires with probability one at t , then $-i$ strictly prefers to follow at this time. (This is a consequence of (A4).) Therefore, since $\bar{L}_i(\cdot)$ is strictly increasing on $[0, T) \supset [0, \bar{t})$, i will choose to fire at some $t < \bar{t}$ only if $-i$ is intending to fire immediately after t . Otherwise, i would prefer to wait a little. On the other hand, if $-i$ is intending to fire immediately after t , then i will fire at t only if he prefers firing to being fired at at this time. Otherwise, he could earn $\bar{F}_i(t) > \bar{L}_i(t)$ by letting $-i$ fire! Together, these facts imply that neither agent will fire unless *both* weakly prefer leading to following. That is, there can be no equilibrium in which the bullet is fired before \bar{t} . We now argue that if an equilibrium exists, the bullet must be fired at \bar{t} . Suppose that it is not fired until $t > \bar{t}$. Now it turns out that in any equilibrium, i 's payoff must be either $\bar{L}_i(t)$ or $\bar{F}_i(t)$. (This follows from Proposition V in Simon [9].) Since i and $-i$ cannot *both* win the duel, at least one of them, say i , must earn $\bar{F}_i(t)$. However, i can preempt, fire just before t , attain essentially $\bar{L}_i(t)$ and thus do better. (If $t < T$, this is true by continuity; if $t > T > t_i$, it follows from (A2.c).) This establishes (informally) that if an equilibrium exists, then the bullet must be fired at \bar{t} . Moreover, exactly one player will fire at \bar{t} . This is true because if both were to fire with positive probability, then i would earn $S_i(\bar{t}) < \bar{L}_i(\bar{t})$ with positive probability, so that i 's expected payoff would be strictly less than $\bar{L}_i(\bar{t})$. In this case, by continuity (A2.c), i could do better by firing just before \bar{t} and attaining essentially $\bar{L}_i(\bar{t})$. Finally, we need to establish that in any equilibrium, player i attains the payoff $\bar{L}_i(\bar{t})$. If

$t_i = t_j = \bar{t}$, then $\bar{L}_i(\bar{t}) = \bar{F}_i(\bar{t})$, for each i . Since exactly one player fires at \bar{t} , i must earn $\bar{L}_i(\bar{t})$, whether he leads or follows. If $t_i < t_{-i}$, then i must fire with probability one; otherwise, he could do better by preempting. In this case, therefore, i earns $\bar{L}_i(\bar{t})$. Moreover, by definition, \bar{t} must equal t_{-i} . That is, $-i$ must be indifferent between leading and following, so that his payoff when i fires is $\bar{F}_{-i}(\bar{t}) = \bar{L}_{-i}(\bar{t})$. This completes our verification of uniqueness.

IV. The Formal Model.

Outline of the model.

In this section, we extend our model to incorporate histories. We will then be able to formalize our analysis of the multistage duel. A history in our model will be a finite list of 'states', paired with the times at which the system changes from one state to the next. We will distinguish between three kinds of states: there will be '*firing states*', '*terminal states*' and an additional state, called '*prepare to fire!*' The system starts out in the state '*prepare to fire!*' We shall represent this state by the symbol ' \emptyset ' (nobody has fired). The moment a bullet is fired, there is a switch to one of the three '*firing states*.' These three states are denoted by $\{i\}$, $\{j\}$ and $\{i,j\}$. They indicate the set of agents who have just fired. If i alone fires, the system switches to $\{i\}$. If both players fire simultaneously, it switches to $\{i,j\}$. Immediately after this switch, nature moves, to determine whether a hit has been scored. If she decides that all bullets will miss their targets, she resets the state to '*prepare to fire!*' Otherwise, the system moves into one of the three '*terminal states*.' These three states are denoted by $\langle i \rangle$, $\langle j \rangle$ and $\langle i,j \rangle$ and indicate which agents have scored hits. The generic state will be denoted by ' a .' Since our states are denoted by sets, we will often use the notation ' $i \in a$.' If a is a firing state, this means that i is one of the agents who has just fired. If, say, state $\langle i \rangle$ is reached, this means that i has scored a hit and won the game. If both agents score hits, then $\langle i,j \rangle$ is reached; in this case, the game is over, but there is no winner.

We can now state more precisely how we model histories. Each history will consist of a string of pairs. Each pair is a time, together with a state. Each history is 'initialized' with the same "zero'th" pair, $(0, \emptyset)$. (Recall that \emptyset denotes the state '*prepare to fire!*') Each *odd* numbered pair will be of the form (t, a) , where a is one of the three '*firing states*'. For example, the history $\{(0, \emptyset), (t, \{i\})\}$ is interpreted as

"the first bullet in the game is fired by i alone at t ; nature's judgment has not yet been announced." Each *even* numbered pair (beyond the zero'th) will record nature's decision, which is made immediately after the last bullet was fired. For example, the history above can be augmented in one of two ways: if i 's bullet hits, then the "next" history will be $[(0, \emptyset), (t, \{i\}), (t, \langle i \rangle)]$ and the game will be over. If i misses, it will be $[(0, \emptyset), (t, \{i\}), (t, \emptyset)]$ and play can continue.

A decision node is a point in time, paired with a past history of the system. (Thus, a continuous time decision node is just like a discrete-time one.) A strategy for an agent is a function that assigns a probability weight to each decision node. If the history part of a decision node ends in either a firing or a terminal state, then the agent has no decision to make. In the former case, he must await nature's decision; in the latter, the game is over. If the last state of the history is '*prepare to fire!*' and the agent still has bullets remaining, then he must decide whether or not to fire at this node. Thus, for each *given* history, the agent's decision problem is just as it was in the classical duel. Formally, therefore, a strategy for our multistage game is just a family of 'classical duel' strategies, one for every history that ends in the state '*prepare to fire!*'

An outcome for our game will be a probability measure on the set of possible histories. Because of the restrictions on strategies that are imposed in this paper, the outcomes that can actually arise in our model have an extremely simple structure: each outcome will concentrate mass on a finite set of histories. We will define an outcome function for the game. This function assigns an outcome to each strategy profile and decision node. This outcome is interpreted as "what happens" when the given strategy profile is played, starting at the given decision node.

Each agent has a valuation function that assigns a value to each history. This function has already been described in section II. The expected payoff function for the game assigns a pair of payoffs to each strategy profile and decision node. The expected payoff an agent receives if a profile is played from a decision node is the expected value of the outcome generated by these strategies, starting from this node. We will now discuss each of the components of the model in more detail.

Histories.

Recall that a history of the game is a string of pairs, recording the various states through which the system has passed, together with the times at which the system moved from one state to the next. Let h denote

the generic history. Each history begins with the 'initial pair,' $(0, \emptyset)$. The history consisting of just this pair will be called the zero-length history and will be denoted by h^\emptyset . A history in which the state has changed k times will be called a " k -length history." Thus, a " k -length history" actually has $k+1$ pairs. We denote the length of a history h by $k(h)$. The pairs of h will be enumerated as $[(t^0(h), a^0(h)), (t^1(h), a^1(h)), \dots, (t^\kappa(h), a^\kappa(h)), \dots, (t^{k(h)}(h), a^{k(h)}(h))]$, where $(t^0(h), a^0(h))$ is the initial pair $(0, \emptyset)$ and $(t^\kappa(h), a^\kappa(h))$ denotes, respectively, the κ 'th transition time and the state to which the system moves at this time.

We will frequently need to refer to both the last and the penultimate pair of a history. We will refer to these, respectively, as $(t^{last}(h), a^{last}(h))$ and $(t^{last-1}(h), a^{last-1}(h))$. Also, we shall refer to $a^{last-1}(h)$ as the last firing state of h . We now summarize the conditions that a history h must satisfy.

- (i) $(t^0(h), a^0(h)) = (0, \emptyset)$;
- (ii) for each odd κ , $a^\kappa(h)$ is a firing state, i.e., either $\{i\}$, $\{j\}$ or $\{i,j\}$;
- (iii) for each odd κ , $t^\kappa(h) \geq t^{\kappa-1}(h)$;
- (iv) for each even $\kappa > 0$, $t^\kappa(h) = t^{\kappa-1}(h)$;
- (v) for each even $\kappa < k(h)$, $a^\kappa(h) = \emptyset$;
- (vi) if $k(h)$ is even, then $a^\kappa(h)$ is either \emptyset or a firing state, i.e., either $\langle i \rangle$, $\langle j \rangle$ or $\langle i,j \rangle$;
- (vii) if $a^{k(h)}(h)$ contains i then $a^{k(h)-1}(h)$ contains i .

Note the weak inequality in condition (iii). It means that a string of bullets can be fired in instantaneous succession. Condition (iv) states that nature always makes her decision immediately after bullets are fired. Condition (v) states that any even state that is *not* the end of the game must be the state '*prepare to fire!*' Condition (vi) describes the various ways that the game can end. Note that a history may end with the state '*prepare to fire!*' This can happen if agents fire all of their bullets and they all miss. Alternatively, agents may simply stop firing before the game is over. Condition (vii) states that an agent cannot score a hit unless he has just fired a bullet. Let H^k denote the set of k -length histories satisfying (i)-(vii). Let $H = \bigcup_k H^k$ denote the universe of all possible histories.

We will denote by $h_{1\kappa}$ the κ -length truncation of h . The history $h_{1\kappa}$ is obtained from h by truncating it after its first κ changes in state. For example, if $h = [(0, \emptyset), (0, \{i\}), (0, \emptyset), (1/2, \{j\}), (1/2, \langle j \rangle)]$,

then h_{11} is the 1-length history $[(0, \emptyset), (0, \{i\})]$ and h_{12} is $[(0, \emptyset), (0, \{i\}), (0, \emptyset)]$. Note that every truncation of h also satisfies conditions (i)-(vii) above and so is a member of H .

It will be important to keep track of the number of bullets that each agent has left to fire. Accordingly, we define a function $b(\cdot) = (b_i(\cdot), b_j(\cdot))$ that counts the number of bullets each agent has left to fire. That is, for each τ , $b_i(h) = \beta_i^* - \#\{\kappa: \tau \in a^\kappa(h)\}$, where β_i^* is the number of bullets that i had at the beginning of the game.²⁴ In conformity with our earlier terminology, we will refer to $b(h)$ as the phase of the history h .

Decision Nodes.

A decision node is a point in time, paired with a complete description of past activity in the system. Decision nodes play the same role in our model that they play in the conventional discrete-time one. In particular, the set of decision nodes is the domain on which a continuous-time strategy will be defined.

Let (t, h) denote the generic decision node, where t is a point in time, and h is a history in H whose last transition time occurs strictly before t . For example, for $t \geq \frac{1}{2}$, the decision node $(t, [(0, \emptyset), (\frac{1}{2}, \{i\}), (\frac{1}{2}, \emptyset)])$ is interpreted as follows: the present time is t ; player i fired his first bullet at time $\frac{1}{2}$ and nature declared that he missed. Since each history keeps track of each bullet that has been fired in the past, agents can condition their choices at each decision node on this information.

Observe that we allow decision nodes of the form $(t^{last}(h), h)$. That is, agents are allowed to move immediately after a change in state has occurred. In this case, 'immediately' should be interpreted as "at the first available opportunity." On every discrete-time grid, this first opportunity is well-defined. It is, therefore, consistent with our conception of continuous time--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., $DN = \bigcup_{h \in H} \{(t, h): t \in [t^{last}(h), 2)\}$. We denote the generic element of DN by (t, h) . We emphasize that DN is *not* the cartesian product of $[0, 2)$ and H . Whenever we refer to $(t, h) \in DN$, we are implicitly asserting that that t weakly exceeds $t^{last}(h)$, i.e., the last time that a change in state occurred.

²⁴ We are, of course, abusing notation, because $b(h)$ depends on β^* as well as h .

Behavior Strategies.

A behavior strategy is a function assigning a point in $[0, 1]$ to each decision node.²⁵ The scalar $f_i(t, h)$ is interpreted as the probability that i will fire a bullet at the decision node (t, h) . We restrict agents to choose strategies that satisfy restrictions (F1)-(F5) below. The first of these enforces the rules of the game described above. The others are imposed for technical reasons.

First, a player can fire with positive probability at (t, h) only if he has not yet used up all his bullets and if the most recent state of h is 'prepare to fire!' That is,

$$\text{for all } (t, h) \in DN, f_i(t, h) > 0 \text{ only if } b_i(h) > 0 \text{ and } a^{last}(h) = \text{'prepare to fire!'} \quad (F1)$$

Our next two assumptions correspond to the ones that we imposed on strategies for the classical duel. The first is a regularity condition that is imposed to ensure that our outcome function is well-defined. We require that for all h ,

$$f_i(\cdot, h) \text{ is a piecewise rational function of time.}^{26} \quad (F2)$$

The next condition is that i 's strategy must be right continuous at t , unless i is either firing or waiting with probability one. In symbols, we require that for all (t, h) ,

$$\text{If } f_i(t, h) \in (0, 1), \text{ then } f_i(\cdot, h) \text{ is right continuous at } t. \quad (F3)$$

Our two remaining restrictions are purely technical in nature. (F4) restricts the way strategies can depend on the past history of the system. Its role is to guarantee that our outcome function is indeed the limit of the discrete-time outcome functions. In the present context, this restriction is again completely innocuous. Let $\mathbf{a}(h)$ denote the vector of states of h . That is, $\mathbf{a}(h) = (a^1(h), \dots, a^x(h), \dots, a^{t(h)}(h))$. We require that agents treat as indistinguishable any two histories that have the same vector of states. In addition, we impose a right continuity condition on agents' immediate reactions to histories. Precisely, the restriction is:

²⁵ In any given multi-stage game, only a small subset of the universe of decision nodes will actually be relevant. For example, if agents start out with only one bullet each, then histories that have more than two different firing states cannot arise. We will simply ignore these redundant nodes from now on.

²⁶ A function θ is piecewise rational on $[0, 2]$ if there exists a finite subset, $\{0 = t^0 < \dots < t^k < \dots < t^{last} = 2\}$, of this interval such that for each $k < last$, there exist two polynomial functions θ' and θ'' , defined on (t^k, t^{k+1}) , such that the restriction of θ to (t^k, t^{k+1}) is equal to $\frac{\theta'}{\theta''}$.

For all h and h' such that $a(h') = a(h)$, and all $s > \max(t^{last}(h), t^{last}(h'))$, we require that (F4)
 $f_i(s, h') = f_i(s, h)$. Moreover, for each h and sequence (h^n) such that for all n ,
 $a(h^n) = a(h)$ and $t^{last}(h^n) \downarrow t^{last}(h)$, we require that $\lim_n f_i(t^{last}(h^n), h^n) = f_i(t^{last}(h), h)$.

Our final condition is completely ad hoc; we impose it because it drastically simplifies the argument that our equilibria is unique. Indeed, without some assumption of this kind, it would be extremely difficult to characterize the set of equilibria. The condition states that if player i alone has just fired, then at the very next instant, he must fire with probability zero:

$$\text{If } a^{last-1}(h) = \{i\}, \text{ then } f_i(t^{last}(h), h) = 0. \quad (\text{F5})$$

The outcome function.²⁷

Our outcome function assigns to each strategy profile and point in time a probability distribution over histories. As explained above, the outcome generated by f from (t, h) is the limit of the sequence of outcomes--i.e., of probability measures over histories--generated by restricting f to an arbitrary, increasingly fine sequence of discrete-time grids, and playing these restrictions from (t, h) . For the class of satisfying (F1)-(F5) above, this limit distribution has a very simple structure. Assumptions F2-F3 together guarantee that with probability one, the state will change at the first instant that some agent fires with positive probability. Precisely, for each f and decision node (t, h) such that $a^{last}(h) = \emptyset$:

$$\begin{aligned} &\text{if } f_i(\cdot, h) > 0, \text{ for some } i, \text{ either at or immediately after } t \quad (4.1) \\ &\text{then with probability one the state will change to some firing state at } t. \end{aligned} \quad ^{28}$$

A consequence of (4.1) is that the outcome generated by any profile of strategies concentrates mass on only a finite number of histories.

We first define the conditional probability that the system will change to state α at τ given that agents start playing f from the decision node (t, h) . We will denote this probability by $TP((\alpha, \tau); f, t, h)$ and call it the transition probability of (α, τ) given (f, t, h) . There are two cases to consider. If $a^{last}(h)$ is

²⁷ Some work is needed to verify that the outcome function we specify here is indeed the limit of the corresponding discrete-time outcomes. See [9], [10] and [11] for details.

²⁸ This was established on pp. 16-17.

a firing state, then $TP((\alpha, \tau); f, t, h)$ is exogenously determined by 'nature.' If $a^{last}(h)$ is the state *prepare to fire!* then this transition probability is determined endogenously.

First, fix h such that $a^{last}(h)$ is a firing state and let α denote *either* a terminal state or the state *'prepare to fire!'* The transition probability, $TP((\alpha, \tau); f, t, h)$, is determined by nature. She always moves immediately, so that $TP((\alpha, \tau); f, t, h) = 0$ whenever $t \neq \tau$. Also, an agent cannot score a hit unless he has just fired, so that $TP((\alpha, \tau); f, t, h) = 0$ unless all agents in $\alpha \subset a^{last}(h)$. For $t \in a^{last}(h)$, the probability that t scores a hit at τ is τ . Therefore, the conditional probability the system will switch to α is: (the probability that exactly those agent in α score hits) *times* (the probability that agents in $a^{last}(h) \sim \alpha$ all miss). That is, for each h such that $a^{last}(h)$ is a firing state, and each α that is either a terminal state or *'prepare to fire!'*

$$TP((\alpha, \tau); f, t, h) = \begin{cases} \tau^{\#\alpha} (1 - \tau)^{\#(a^{last}(h) - \alpha)} & \text{if } \alpha \subset a^{last}(h) \text{ and } \tau = t^{last}(h) \\ 0 & \text{otherwise} \end{cases} \quad (4.2)$$

For example, if $\alpha = \langle j \rangle$ and $a^{last}(h) = \{i, j\}$, then $(a^{last}(h) \sim \alpha)$ is the singleton set containing i . Since this set and α each have one element, it follows that $TP((\alpha, \tau); f, t, h) = t(1 - t)$.

Now assume that the last state of h , $a^{last}(h)$, is *'prepare to fire!'* From (4.1), we know that $TP((\alpha, \tau); f, t, h)$ will be zero unless τ is the infimum of the times after t at which some agent fires with positive probability. If τ is indeed this infimum, then $TP((\alpha, \tau); f, t, h)$ will be the sum of two terms. The first is the probability that exactly those agents in α fire exactly at τ . The second term is the probability that *no* agent fires at τ , multiplied by the limit, as $s \downarrow \tau$, of the conditional probabilities that the set α fires at s , given that *some* agent fires at s . (See Simon-Stinchcombe [11] for a verification that the outcome defined by this formula is indeed the limit of the discrete-time outcomes.)

To make the above statements precise, we need some additional notation. Fix a strategy profile f and a decision node (s, h) . For each subset α of $\{i, j\}$, we define $\phi^{f, s, h}(\alpha)$ to be the probability that exactly those agents in α fire at s . That is, $\phi^{f, s, h}(\alpha) = \prod_{i \in \alpha} f_i(s, h) \times \prod_{i \notin \alpha} (1 - f_i(s, h))$. Note that $\phi^{f, s, h}(\emptyset)$ is the probability that *no* bullet is fired at s . For each h such that $a^{last}(h)$ is *'prepare to fire!'* and each firing state α , we define $TP((\alpha, \tau); f, t, h)$ as follows. If $\tau \neq \inf\{s > t : \phi^{f, s, h}(\emptyset) < 1\}$, then $TP((\alpha, \tau); f, t, h) = 0$. If

$\tau = \inf\{s > t: \phi^{f,s,h}(\emptyset) < 1\}$, then²⁹

$$TP((\alpha, \tau); f, t, h) = \phi^{f,\tau,h}(\alpha) + \phi^{f,\tau,h}(\emptyset) \left(\lim_{s \downarrow \tau} \frac{\phi^{f,s,h}(\alpha)}{(1 - \phi^{f,s,h}(\emptyset))} \right)^{30} \quad (4.3)$$

For the equilibrium profiles that we construct in this paper, only two kinds of transitions actually arise along the equilibrium path. In the first kind, agents' strategies are symmetric. In this case, agents will actually fire with probability zero at $\tau = \inf\{s > t: \phi^{f,s,h}(\emptyset) < 1\}$, and randomize thereafter. In this case,

$$TP(\{(i), \tau\}; f, t, h) = TP(\{(j), \tau\}; f, t, h) = \left(\lim_{s \downarrow \tau} \frac{\phi^{f,s,h}(\{(i)\})}{(1 - \phi^{f,s,h}(\emptyset))} \right) = 1/2. \quad (\text{We verified this case in the}$$

preceding section.) The second kind of outcome generated by our equilibrium strategies has some agent, say i , firing with probability one at τ . In this case $TP(\{(i), \tau\}; f, t, h) = \phi^{f,\tau,h}(\{(i)\}) = 1$. Off the equilibrium path, obviously, things can be more complicated.

We are now ready to define our outcome function. The outcome generated from (t, h) by f will be the finite support measure on histories defined as follows. First, if neither agent fires with positive probability between $t^{last}(h)$ and the end of the game, then the history will not change. In this case, obviously, $\phi^{f,t,h}\{h\} = 1$. Assume therefore that some agent fires with positive probability before the end of the game. Clearly, $\phi^{f,t,h}$ will assign positive probability to some history η only if $\eta_{|k(h)} = h$, i.e., if the appropriate length truncation of η agrees with h . For any history η satisfying this condition, the probability that this history will be realized is the transition probability of $(t^{k(h)+1}(\eta), a^{k(h)+1}(\eta))$ given (f, t, h) times the product, taken over the κ 's between $k(h)+2$ and $k(\eta)$, of the transition probabilities of $(t^\kappa(\eta), a^\kappa(\eta))$ given $(f, t^{\kappa-1}(\eta), \eta_{|k-1})$. In symbols, for each η such that $\eta_{|k(h)} = h$:

$$\phi^{f,t,h}\{\eta\} = TP((t^{k(h)+1}(\eta), a^{k(h)+1}(\eta)); f, t, h) \times \prod_{\kappa=k(h)+2}^{k(\eta)} TP((t^\kappa(\eta), a^\kappa(\eta)); f, t^{\kappa-1}(\eta), \eta_{|k-1})$$

²⁹ If both the numerator and denominator converge to zero, we need to be careful that $\left(\lim_{s \downarrow \tau} \frac{\phi^{f,s,h}(\alpha)}{(1 - \phi^{f,s,h}(\emptyset))} \right)$ is well-defined. In our model, convergence is guaranteed by our restriction that strategies be piecewise rational. We simply invoke the obvious generalization of L'Hospital's rule.

³⁰ When we change symbols in the obvious way, the right hand side becomes the outcome function for the classical duel. Specifically, define $\phi^{f,s}$ analogously to $\phi^{f,s,h}$. If agents play the classical duel profile, ξ , from t , and $\tau = \inf\{s > t: \phi^{f,s}(\emptyset) < 1\}$, then the probability that the set α fires at τ is precisely $\phi^{f,\tau}(\alpha) + \phi^{f,\tau}(\emptyset) \left(\lim_{s \downarrow \tau} \frac{\phi^{f,s}(\alpha)}{(1 - \phi^{f,s}(\emptyset))} \right)$.

To illustrate how our outcome function works, it is instructive to work through a complete example. Each agent starts out with two bullets; that is, $\beta^* = (2, 2)$. We will partially specify a profile f and study the outcome that it generates. The profile will agree along the equilibrium path with our equilibrium strategies, but will not be fully specified off this path. Choose f to have the following properties. For each h such that $b_i(h) = b_j(h)$, set $f_i(\cdot, h)$ equal to zero on the interval $[t^{last}(h), t^{b(h)}]$. Beyond $t^{b(h)}$, have i randomize, ensuring that $\lim_{\delta \downarrow 0} f_i(t^{b(h)} + \delta, h) = 0$.³¹ Now consider h such that $b_i(h) > b_{-i}(h) > 0$. Define $f_i(\cdot, h)$ in exactly the same way as above. Set $f_{-i}(\cdot, h)$ equal to $f_i(\cdot, h)$, except that exactly at $t^{b(h)}$, $-i$ fires with probability one instead of zero. If $b_i(h) > b_{-i}(h) = 0$, then specify $f_i(\cdot, h)$ so that to solve i 's single person decision problem.

This pair of strategies generates the following outcome in our multi-stage game. It is illustrated schematically by Figure 2. With probability one, the first change in state occurs at $t^{2.2}$. The new firing state will be $\{i\}$ or $\{j\}$ with equal probability. From now on, we will assume that state $\{i\}$ is actually realized. Nature now moves immediately. The system moves to the terminal state $\langle i \rangle$ with probability $t^{2.2}$; with the remaining probability, it is reset to 'prepare to fire!' In the latter event, the next change in state occurs at $t^{2.1}$. This time, the state switches to $\{i\}$ again with probability one. Nature now moves again and the system moves to $\langle i \rangle$ with probability $t^{2.1}$; with the remaining probability, it is reset to 'prepare to fire!' In this event, i no longer has any bullets to fire and j becomes as a monopolist. By making the appropriate substitutions, it is straightforward to verify that for each i the expected value of playing f from the beginning of the game is $\Pi_i^{2.2}$, where this term is defined in the usual way from the appropriate set of classical duels.

Payoffs and Equilibrium Notions.

Player i 's valuation function, V_i , assigns a value to each history. If $a^{last}(h) = \langle i \rangle$, then i wins the game. In this case, $V_i(h)$ is $(2 - t^{last}(h))\pi^*$ minus the unit cost, c^* , of each bullet that he fires. Otherwise, the value of a history to i is, simply, minus the cost of each bullet that he has fired. Summarizing,

$$V_i(\eta) = \begin{cases} (2 - t^{last}(\eta))\pi^* - b_i(\eta)c^* & \text{if } a^{last}(h) = \langle i \rangle \\ -b_i(\eta)c^* & \text{otherwise} \end{cases} .$$

The expected payoff function P assigns a payoff vector to each strategy profile and decision node. $P_i(f, t, h)$ is player i 's payoff if agents play f




FIGURE 2.

from the decision node t, h . That is,

$$P_i(f, t, h) = \int V_i(\eta) d\sigma^{f,t,h}(\eta). \quad (4.4)$$

(Recall that $\sigma^{f,t,h}(\eta)$ concentrates mass on only finitely many η 's.)

Like the classical duel, our multistage duel has a continuum of subgames, one corresponding to each decision node. The subgame associated with the node (t, h) is simply the multistage duel played on the interval $[t, 2)$, in which agents start out with an endowment vector of $b(h)$. That is, the endowment vector for this subgame is determined by subtracting from each agent's initial endowment of bullets the number that he has fired in to the history h . We can now define subgame perfection in the usual way. For each decision node, we will say that f is a Nash equilibrium for the subgame beginning at (t, h) if for all i , and all f'_i , $P_i(f, t, h) \geq P_i(f'_i, f_{-i}, t, h)$. A profile f is a subgame perfect equilibrium if f is a Nash equilibrium for every subgame.

V. Construction of equilibrium strategies.

In this section, we construct a pair of equilibrium strategies, f^* , for the multi-stage duel with parameters (β^*, π^*, c^*) . As usual, we assume that this duel passes our numerical test specified in section II. We will specify an inductive procedure for defining a family of classical duels, one corresponding to each phase of our multi-stage game. The equilibrium profiles for each classical duel will be used to define other duels, corresponding to "earlier" phases of the game. The procedure is delicate because we must choose our profiles so that the duels they define will satisfy conditions (A1)-(A4), allowing us to invoke Proposition IV. Once we have defined and solved the classical duels corresponding to each phase of the multi-stage duel, we define f^* simply by patching together the equilibrium profiles for each of the duels.

The conditions that requires particular care are assumptions (A2.b) and (A4): we need to ensure that in the relevant region, following is preferable to firing simultaneously (A4) but not so attractive that it is preferable to leading, in which case (A2.b) would be violated.^{32,33} Fortunately, there is a simple way to resolve

³² The following example illustrates how condition (A4) can fail. Let $\pi^* = 16$, $c^* = 1$, $t = \frac{1}{2}$ and suppose that each player has a large number of bullets remaining. Suppose that at time $\frac{1}{2}$, i intends to fire as many as three bullets in immediate succession, if his earlier bullets miss. If j waits in the hope that all three bullets miss, his payoff must be strictly less than $(\frac{1}{2})^3 \times 16 \times (2 - \frac{1}{2}) - 1 = 2$. On the other hand, if he fires at the same instant that i fires his first shot, j 's payoff will be at least $\frac{1}{4} \times 16 \times (2 - \frac{1}{2}) - 1 = 5$.

this difficulty. For a duopoly phase, β , let $IS(\beta)$ denote the set of immediate successors of β . $IS(\beta)$ consists of the phases (β_i-1, β_j) , (β_i, β_j-1) and (β_i-1, β_j-1) . These are the phases that are directly relevant for the purposes of defining CD^β . From Proposition IV, we know that for any $\beta' \in IS(\beta)$ and any subduel of $CD^{\beta'}$ that begins beyond the firing time $t^{\beta'}$ for $CD^{\beta'}$, there are three kinds of equilibria. For each τ , there is one kind that i strictly prefers. The third kind, involving randomization, is bad for both players. Accordingly, we can use these different equilibria to reward following and punish firing simultaneously, when we define the payoffs for CD^β . For example, if agents in CD^β fire simultaneously at $\tau \geq t^{\beta_i-1, \beta_j-1}$, we will determine their payoffs by playing the equilibrium profile that is bad for both of them in the subduel $CD^{\beta_i-1, \beta_j-1}(\tau)$.³⁴ Similarly i leads and j follows at τ , payoffs will be determined by playing the equilibrium that j strictly prefers in $CD^{\beta_i-1, \beta_j}(\tau)$. It turns out that this construction strikes just the right balance: in classical duels defined by this procedure, following will be preferable to firing simultaneously, but leading will be preferable to following in the regions where these inequalities are important.

We now construct the strategy pair f^* . As usual, we begin with the monopoly phases of the game. Fix $\beta_i > 0$ and let h be a phase- $(\beta_i, 0)$ history. Set $f_j^*(\cdot, h) \equiv 0$ and define $f_i^*(\cdot, h)$ by, for $t \in [0, 2)$,

$$f_i^*(t, h) = \begin{cases} 1 & \text{if } t < t^{\beta_i, 0} \\ 0 & \text{if } t \geq t^{\beta_i, 0} \end{cases}, \text{ where } t^{\beta_i, 0} \text{ is defined on p. 4. Now, for each firing state } \alpha, \text{ define the}$$

continuation payoff functions $\Psi(\cdot; \alpha; (\beta_i, 0)) = (\Psi_i(\cdot; \alpha; (\beta_i, 0)), \Psi_j(\cdot; \alpha; (\beta_i, 0)))$. (The role of α will be explained later.) That is, $\Psi_i(\tau; \alpha; (\beta_i, 0))$ denotes i 's monopoly payoff if he plays f_i^* starting from the decision node (τ, h) . Note that for each β_i and $t \leq t^{\beta_i, 0}$, $(\Psi_i(\tau; \alpha; (\beta_i, 0)), \Psi_j(\tau; \alpha; (\beta_i, 0))) = (\Pi_i^{\beta_i, 0}, 0)$. Beyond this time, $\Psi_i(\cdot; \alpha; (\beta_i, 0))$ will be strictly decreasing.

³³ In an *asymmetric* game, it is difficult to guarantee existence unless (A4) is satisfied. To see this, suppose that for each i , $L_i(\cdot)$ is strictly increasing, and that beyond the firing time \bar{t} , $L_i(\cdot) > S_i(\cdot) > F_i(\cdot)$ while $L_j(\cdot) > F_j(\cdot) > S_j(\cdot)$. For reasons explained above, i will be willing to fire at t only if j fires with positive probability either at or immediately after t . However, if the inequalities are as specified for i , then i never be indifferent between firing and following, and so, in equilibrium, will fire at t with probability either zero or unity. However, if i fires with probability one, j prefers following to moving simultaneously. Thus it is impossible to construct an equilibrium profile in which any agent moves with positive probability at any time.

³⁴ Obviously, if they fire simultaneously before t^{β_i-1, β_j-1} , then their continuation payoffs are uniquely determined.

Once we have defined f^* for all the monopoly phases, we can proceed inductively. We will solve all the phases in which one player has only one bullet, i.e., $(1,1)$, $(2,1)$, ..., $(\beta_i^*, 1)$ then proceed to phase $(2,2)$, $(3,2)$, etc. Fix a duopoly phase, β . Assume that for each $\beta' \in \{(n, \beta_j - 1), 1 \leq n \leq \beta_i^*\} \cup \{(n, \beta_j), 1 \leq n < \beta_i\}$, the classical duel $CD^{\beta'}$ has been defined and its payoffs satisfy conditions (A1)-(A4). [This hypothesis is certainly satisfied when $\beta = (2,1)$: in this case, the only β' in the above set is $(1,1)$, and $CD^{1,1}$, defined on p. 16 certainly satisfies (A1)-(A4).] To define payoffs for CD^β , we will need continuation payoff functions, $\Psi(\cdot; \alpha; \beta')$, for each α and each immediate successor, β' , of β . If β' is a monopoly phase, the $\Psi(\cdot; \alpha; \beta')$'s have already been defined. Otherwise, we will define them by constructing an equilibrium profile, $\xi^{\tau; \alpha; \beta'}(\cdot)$, for the subduel $CD^{\beta'}(\tau)$ and setting $\Psi(\tau; \alpha; \beta')$ equal to the payoffs generated by playing this profile in that subduel.

The profile $\xi^{\tau; \alpha; \beta'}(\cdot)$ is defined as follows. For each $t \in [\tau, 2)$ and each firing state α , we define:³⁵

$$\xi_i^{\tau; \alpha; \beta'}(t) = \begin{cases} 0 & \text{if } t \in [\tau, t^{\beta'}) \\ 1 & \text{if } t = t^{\beta'} > \tau \text{ and } \beta_i' < \beta_i \\ 0 & \text{if } t = t^{\beta'} > \tau \text{ and } \beta_i' \geq \beta_i \\ 1 & \text{if } t = \tau \geq t^{\beta'} \text{ and } \alpha = \{j\} \\ 0 & \text{if } t = \tau \geq t^{\beta'} \text{ and } \alpha = \{i\} \\ \frac{L_j^{\beta'}(t) - F_j^{\beta'}(t)}{L_j^{\beta'}(t) - S_j^{\beta'}(t)} & \text{if } t = \tau \geq t^{\beta'} \text{ and } \alpha = \{i, j\} \\ \frac{L_j^{\beta'}(t) - F_j^{\beta'}(t)}{L_j^{\beta'}(t) - S_j^{\beta'}(t)} & \text{if } t > \max(\tau, t^{\beta'}) \end{cases} \quad (5.1)$$

Now for each α , define the continuation payoff $\Psi(\tau; \alpha; \beta')$ to be the outcome generated by playing the profile $\xi^{\tau; \alpha; \beta'}(\cdot)$ in the subduel $CD^{\beta'}(\tau)$. It is straightforward to verify that $\Psi(\tau; \alpha; \beta')$ has the following values:

$$\Psi(\tau; \alpha; \beta') = \begin{cases} (L_i^{\beta'}(t^{\beta'}), L_j^{\beta'}(t^{\beta'})) & \text{if } \tau \in [\tau, t^{\beta'}) \\ (F_i^{\beta'}(\tau), L_j^{\beta'}(\tau)) & \text{if } \tau \geq t^{\beta'} \text{ and } \alpha = \{i\} \\ (L_i^{\beta'}(\tau), F_j^{\beta'}(\tau)) & \text{if } \tau \geq t^{\beta'} \text{ and } \alpha = \{j\} \\ (F_i^{\beta'}(\tau), F_j^{\beta'}(\tau)) & \text{if } \tau \geq t^{\beta'} \text{ and } \alpha = \{i, j\} \end{cases} \quad (5.2)$$

³⁵ There is, here, a potential source of confusion. In section II, we constructed the strategy profile $\xi^{\tau, \alpha}$; Now, we define the profile

We can now specify the payoffs that define CD^β . For each $t \in \{0, 2\}$.

$$L_i^\beta(t) = t\pi^*(2-t) - c^* + (1-t)\Psi_i(t; \{i\}; (\beta_i-1, \beta_j)) \quad (5.3.a)$$

$$F_i^\beta(t) = (1-t)\Psi_i(t; \{j\}; (\beta_i, \beta_j-1)) \quad (5.3.b)$$

$$S_i^\beta(t) = (1-t) \left[t\pi^*(2-t) - c^* + (1-t)\Psi_i(t; \{i, j\}; (\beta_i-1, \beta_j-1)) \right] - tc^* \quad (5.3.c)$$

$$L_j^\beta(t) = t\pi^*(2-t) - c^* + (1-t)\Psi_j(t; \{j\}; (\beta_i, \beta_j-1)) \quad (5.3.d)$$

$$F_j^\beta(t) = (1-t)\Psi_j(t; \{i\}; (\beta_i-1, \beta_j)) \quad (5.3.e)$$

$$S_j^\beta(t) = (1-t) \left[t\pi^*(2-t) - c^* + (1-t)\Psi_j(t; \{i, j\}; (\beta_i-1, \beta_j-1)) \right] - tc^* \quad (5.3.f)$$

To complete the inductive step, we need to prove that these payoffs satisfy conditions (A1)-(A4). Once this is established, we can define duels corresponding to phase (β_i', β_j) , for each $\beta_i^* \geq \beta_i' > \beta_i$. Finally, we can use the equilibrium profiles obtained in prior steps to define the duel corresponding to (β_j+1, β_j+1) . In this way, we can define classical duels corresponding to each duopoly phase of the game.

We can now define our profile f^* for the monopoly phases of the game. For each duopoly phase history $h \neq h^\emptyset$, we define $f^*(\cdot, h) = \xi^{(t^{h^{\alpha-1}(h)}; \alpha^{h^{\alpha-1}(h)}; b(h))}$. For the zero-length history h^\emptyset , choose α arbitrarily from among the firing-states, and define $f^*(\cdot, h^\emptyset) = \xi^{(0; \alpha; \beta^*)}$. Our final result is:

Th'm V: Fix a triple (β^*, π^*, c^*) . If this triple satisfies the test specified in Proposition IV, then the strategy profile f^* is a subgame perfect equilibrium profile for the multistage duel with these parameters.

As we have indicated, the delicate step in the proof is the verification that the classical duels defined above satisfy conditions (A1)-(A4). Once this is established, the theorem follows easily from Proposition IV. We will outline the proof that condition (A4) is satisfied for player j . The remainder of the proof is relegated to the appendix.

First suppose that i fires in CD^β at $\tau < t^{\beta_i-1, \beta_j}$. If he misses, the bullet in phase (β_i-1, β_j) will not be fired until t^{β_i-1, β_j} . Suppose that j follows at τ . If i misses, j has the option of leading at τ in phase

$\xi^{(\tau; \{i\}; \beta)}(\cdot)$. The superscripts τ and $\{i\}$ indicate very different things: τ signals that i will fire immediately at τ ; $\{i\}$ indicates that i has just fired, and now it is $-i$'s turn to fire.

(β_i-1, β_j) . The probability that i will miss and j will then hit at τ is $\tau(1-\tau)$. If i and j fire simultaneously at τ , the probability that j alone will score is, once again, $\tau(1-\tau)$. On the other hand, the cost of firing simultaneously exceeds the expected cost of following: if j fires simultaneously, he incurs the cost of the bullet whether i hits or misses; if j follows, he incurs this cost only if i misses. Finally, if i 's bullet and j 's both miss, then phase (β_i-1, β_j-1) will be reached. The bullet in this phase will not be fired until $t^{\beta_i-1, \beta_j-1} > t^{\beta_i-1, \beta_j}$; Thus, j 's payoff conditional on reaching phase (β_i-1, β_j-1) will be the same, i.e., $L_j^{\beta_i-1, \beta_j-1}(t^{\beta_i-1, \beta_j-1})$, whether he follows or fires simultaneously at τ in phase β . Therefore, j can do strictly better than firing simultaneously at τ in phase β by following at τ and then firing immediately in phase (β_i-1, β_j) if i misses. This establishes that $F_j^\beta(\tau) > S_j^\beta(\tau)$.³⁶

Now suppose that i fires at $\tau \in [t^{\beta_i-1, \beta_j}, t^{\beta_i-1, \beta_j-1})$. If j follows and i misses, then f^* has j leading immediately in phase (β_i-1, β_j) and, if he misses, firing again at $\max(\tau, t^{\beta_i-1, \beta_j-1})$. Repeat the argument just given to establish that j does strictly better by following. Finally, suppose that i fires at $\tau \geq t^{\beta_i-1, \beta_j-1}$.³⁷ Once again, f^* has j leading immediately if phase (β_i-1, β_j) is reached. The difference here is that if j also misses, the bullet in phase (β_i-1, β_j-1) will be fired immediately. If j follows in phase β and phase (β_i-1, β_j-1) is reached, f^* has i leading. If j fires simultaneously with i in phase β and phase (β_i-1, β_j-1) is reached, f^* has both players randomizing in phase (β_i-1, β_j-1) . In either case, j 's payoff conditional on phase (β_i-1, β_j-1) being reached is $F_j^{\beta_i-1, \beta_j-1}(t^{\beta_i-1, \beta_j-1})$. Therefore, once again, he strictly prefers following to firing simultaneously. This completes our verification that condition (A4) is satisfied.

³⁶ Of course, j chooses *not* to fire immediately at τ , but waits until t^{β_i-1, β_j} before firing. This fact does not our argument however: it only means that the advantage to following is even greater than we indicated.

REFERENCES

- [1] Avinash Dixit, "Perfect Equilibrium in the Duel," mimeo Princeton University (1983).
- [2] Fudenberg, D., and J. Tirole, "Preemption and Rent Equalization in the Adoption of a New Technology," *Review of Economic Studies*, 52 (1985), 383-401.
- [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] Grossman, G. and C. Shapiro, "Dynamic R&D Competition," *Economic Journal*, 97 (forthcoming, 1987)
- [5] Harris, C. and J. Vickers, "Perfect Equilibrium in a Model of a Race," *Review of Economic Studies*, 52 (1985), 193-209.
- [6] Ken Hendriks and Charles Wilson, "Discrete versus Continuous Time in Games of Timing," mimeo, New York University, (1985).
- [7] Carolyn Pitchik, "Equilibria of a Two-Person Game of Non-Zerosum Noisy Game of Timing," *International Journal of Game Theory*, 10, (1982), pp. 207-221.
- [8] Selten, R., "Spieltheoretische Behandlung eines Oligopolmodells mit Nachfrageträgheit," *Zeitschrift für die gesamte Staatswissenschaft*, 12, (1965), pp. 301-324.
- [9] Simon, L., "Basic Timing Games," U.C. Berkeley Working Paper #8745, July 1987
- [10] Simon, L. and M. Stinchcombe, "Extensive Form Games in Continuous Time: Pure Strategies,," U.C. Berkeley Working Paper #8746, July 1987
- [11] Simon, L. and M. Stinchcombe, "Extensive Form Games in Continuous Time: Behavior Strategies," (in preparation).
- [12] Stinchcombe, M., "Game Theory in Continuous Time: a Process Theoretic Approach," Ph. D. dissertation, Berkeley, 1986.

APPENDIX.

We first prove that the conclusions of Theorems II and III hold for the "test firing times" $(\tilde{t}^\beta)_{\beta \leq \beta^*}$ and "test payoff vectors" $(\tilde{\Pi}^\beta)_{\beta \leq \beta^*}$, defined by the algorithm on pp. -12. We then prove existence (Theorem V), in the following way. We fix a triple (β^*, π^*, c^*) that satisfies our test. We show that for every duopoly phase β , the phase β classical duel defined by playing f^* beyond phase β coincides on $[0, \tilde{t}^{\beta, -1, \beta})$ with the "test payoff functions" defined by (2.6.d')-(2.6.d') above. By assumption, the lead and follow functions for this phase intersect in the above interval. Moreover, when the definition of the payoff functions is completed by playing f^* starting from beyond $\tilde{t}^{\beta, -1, \beta}$, the classical duel will satisfy conditions (A1)-(A4). We can conclude from Proposition IV, therefore, that for any phase β history h , and any t , f^* is an equilibrium for subgame that begins at (t, h) . It follows immediately that the conclusions of Theorems II and III hold for the particular equilibrium profile f^* . Finally, we prove the uniqueness parts of Theorems II and III, and simultaneously prove Proposition I, in the following way. We establish that for any equilibrium profile f , and any duopoly phase, β , the lead and follow payoffs for the phase β classical duel defined by playing f beyond phase β coincide with the lead and follow payoffs defined by playing f^* . (This is an immediate consequence of our assumption F5 on strategies.) The payoffs to firing simultaneously may differ. We show, however, that this difference does not matter, i.e., f and f^* are effectively equivalent. The uniqueness results now follow immediately.

VI. Proof of Theorems II and III: the first step.

In this subsection, we prove that the conclusions of Theorems II and III hold for the "test firing times" $(\tilde{t}^\beta)_{\beta \leq \beta^*}$ and "test payoff vectors" $(\tilde{\Pi}^\beta)_{\beta \leq \beta^*}$, defined by the algorithm on pp. -12. As usual, we will always assume that i has at least as many bullets remaining as j . We shall sometimes refer to j as the "weak" and i as the "strong" player.

The Monopolist's problem

Recall from expression (2.1) that the monopolist with β , bullets has to maximize:

$$\bar{L}_i^{\beta, 0}(\cdot) = i\pi^*(2 - t) - c^* + (1 - t)\tilde{\Pi}_i^{\beta, -1, 0} \quad (\text{A.1})$$

Now $\frac{\partial \tilde{L}_i^{\beta,0}(\cdot)}{\partial t} = 2\pi^*(1-t) - \tilde{\Pi}_i^{\beta,-1,0}$, so that $\tilde{L}_i^{\beta,0}(\cdot)$ is maximized at $\tilde{t}^{\beta,0} = 1 - \frac{\tilde{\Pi}_i^{\beta,-1,0}}{2\pi^*}$. Let

$\tilde{\Pi}_i^{2,0} = \tilde{L}_i^{\beta,0}(\tilde{t}^{\beta,0})$. Substituting for $\tilde{t}^{\beta,0}$ in (A.1), we have

$$\begin{aligned} \tilde{\Pi}_i^{2,0} &= 2\pi^*(2 - \tilde{t}^{\beta,0}) + (1 - \tilde{t}^{\beta,0})\tilde{\Pi}_i^{\beta,-1,0} - c^* & (A.2) \\ &= \left(1 - \frac{(\tilde{\Pi}_i^{\beta,-1,0})^2}{4\pi^*}\right)\pi^* + \frac{\tilde{\Pi}_i^{\beta,-1,0}}{2\pi^*}\tilde{\Pi}_i^{\beta,-1,0} - c^* \\ &= (\pi^* - c^*) + \frac{(\tilde{\Pi}_i^{\beta,-1,0})^2}{4\pi^*} \\ &= \tilde{\Pi}_i^{1,0} + \frac{(\tilde{\Pi}_i^{\beta,-1,0})^2}{4\pi^*}. \end{aligned}$$

We will show that i 's expected profit increases with the number of bullets he has remaining. We have

$$\tilde{\Pi}_i^{2,0} - \tilde{\Pi}_i^{\beta,-1,0} = \tilde{\Pi}_i^{1,0} + \frac{(\tilde{\Pi}_i^{\beta,-1,0})^2}{4\pi^*} - \tilde{\Pi}_i^{\beta,-1,0}.$$

It is straightforward to check that the solutions to this system are $2\sqrt{\pi^*}(\sqrt{\pi^*} \pm \sqrt{c^*})$.³⁸

Since $\tilde{\Pi}_i^{2,0}$ is a convex function of $\tilde{\Pi}_i^{\beta,-1,0}$, the difference $(\tilde{\Pi}_i^{2,0} - \tilde{\Pi}_i^{\beta,-1,0})$ must be positive on the interval $[0, 2\sqrt{\pi^*}(\sqrt{\pi^*} - \sqrt{c^*})]$. Moreover, since $\tilde{\Pi}_i^{1,0} = (\sqrt{\pi^*} + \sqrt{c^*})(\sqrt{\pi^*} - \sqrt{c^*})$, we have:

$$2\sqrt{\pi^*}(\sqrt{\pi^*} - \sqrt{c^*}) - \tilde{\Pi}_i^{1,0} = \left[2\sqrt{\pi^*} - (\sqrt{\pi^*} + \sqrt{c^*})\right](\sqrt{\pi^*} - \sqrt{c^*}) = (\sqrt{\pi^*} - \sqrt{c^*})^2,$$

so that $\tilde{\Pi}_i^{2,0} \in [0, 2\sqrt{\pi^*}(\sqrt{\pi^*} - \sqrt{c^*})]$. This establishes that $\tilde{\Pi}_i^{2,0}$ is increasing in β_i over the region that concerns us.

³⁸ For example, when $\tilde{\Pi}_i^{\beta,-1,0} = 2\sqrt{\pi^*}(\sqrt{\pi^*} - \sqrt{c^*})$, we have:

$$\begin{aligned} \tilde{\Pi}_i^{1,0} + \frac{(\tilde{\Pi}_i^{\beta,-1,0})^2}{4\pi^*} - \tilde{\Pi}_i^{\beta,-1,0} &= (\pi^* - c^*) + \frac{2\sqrt{\pi^*}(\sqrt{\pi^*} - \sqrt{c^*})^2}{4\pi^*} - 2\sqrt{\pi^*}(\sqrt{\pi^*} - \sqrt{c^*}) \\ &= (\pi^* - c^*) + (\sqrt{\pi^*} - \sqrt{c^*})^2 4\pi^* - 2\pi^* + 2\sqrt{\pi^*}c^* \\ &= (\pi^* - c^*) + \pi^* + c^* - 2\sqrt{\pi^*}c^* 4\pi^* - 2\pi^* + 2\sqrt{\pi^*}c^* = 0. \end{aligned}$$

The following property of the monopolists' decision problem plays an important role in our subsequent proof that the monopolist always fires later than the duopolists in the corresponding phase. The result is a bound on the rate at which the unopposed player's expected profits increase with the number of bullets he has remaining. The result is stated here for convenience:

$$\left[2\tilde{\pi}_i^{1,0} - \tilde{\pi}_i^{\beta_i+1,0} \right] - (1 - \tilde{r}^{\beta_i+1,0}) \left[2\tilde{\pi}_i^{1,0} - \tilde{\pi}_i^{2,0} \right] > 0. \quad (\text{A.3})$$

To verify (A.3), first observe that:

$$(1 - \tilde{r}^{\beta_i,0})\tilde{\pi}_i^{\beta_i-1,0} = \frac{(\tilde{\pi}_i^{\beta_i-1,0})^2}{2\pi^*} = 2(\tilde{\pi}_i^{2,0} - \tilde{\pi}_i^{1,0}). \quad (\text{A.4})$$

Therefore,

$$\begin{aligned} & \left[2\tilde{\pi}_i^{1,0} - \tilde{\pi}_i^{\beta_i+1,0} \right] - (1 - \tilde{r}^{\beta_i+1,0}) \left[2\tilde{\pi}_i^{1,0} - \tilde{\pi}_i^{2,0} \right] \\ &= \left[2\tilde{\pi}_i^{1,0} - \tilde{\pi}_i^{\beta_i+1,0} \right] + 2 \left[\tilde{\pi}_i^{\beta_i+1,0} - \tilde{\pi}_i^{1,0} \right] - 2(1 - \tilde{r}^{\beta_i+1,0})\tilde{\pi}_i^{1,0} \\ &\equiv \tilde{\pi}_i^{\beta_i+1,0} - 2\tilde{\pi}_i^{1,0} \frac{\tilde{\pi}_i^{2,0}}{2\pi^*} \\ &\equiv \tilde{\pi}_i^{\beta_i+1,0} - 2(\pi^* - c^*) \frac{\tilde{\pi}_i^{2,0}}{2\pi^*} \\ &> \tilde{\pi}_i^{\beta_i+1,0} - \tilde{\pi}_i^{2,0} > 0. \end{aligned}$$

The first equality follows from (A.4); The first inequality holds because $\pi^* > c^*$; the second because $\tilde{\pi}_i^{\beta_i+1,0} > \tilde{\pi}_i^{2,0}$.

The duopoly phases.

We will establish that the following statements are true, for all $1 \leq \beta_j \leq \beta_j^*$ and all $\beta_j' \leq \beta_i' \leq \beta_i^*$.

The proof is induction on β_j' , i.e., the number of bullets that j has available.

$$\bar{t}^{\beta_i'+\beta_j',0} > \bar{t}^{\beta_j'}; \quad [\text{App.5a.}(\beta_i',\beta_j')]$$

$$\bar{\pi}_i^{\beta_i'+1,\beta_j'-1} + \bar{\pi}_j^{\beta_i'+1,\beta_j'-1} > \bar{\pi}_i^{\beta_j'} + \bar{\pi}_j^{\beta_j'}; \quad [\text{App.5b.}(\beta_i',\beta_j')]$$

player j fires the bullet in phase $(\beta_i'+1,\beta_j')$; [App.5c.}(\beta_i',\beta_j')

$$\bar{t}^{\beta_i'+2,\beta_j'} > \bar{t}^{\beta_i'+1,\beta_j'+1}. \quad [\text{App.5d.}(\beta_i',\beta_j')]$$

We first prove that [App.5a.}(\beta_i',1)]-[App.5d.}(\beta_i',1)] are true, for every $\beta_i' \geq 1$. We then proceed inductively.

Fix $\beta_j > 1$ and assume that for every $v \geq \beta_j - 1$, statements [App.5a.}(\beta_i',\beta_j-1)]-[App.5d.}(\beta_i',\beta_j-1)] are true.

For every $\beta_j' < \beta_j$ and every $\beta_i' > \beta_j'$, [App.5a.}(\beta_i',\beta_j')]-[App.5d.}(\beta_i',\beta_j')] are true. We will prove

[App.5a.}(\beta_i',\beta_j)]-[App.5d.}(\beta_i',\beta_j)], for every $\beta_i' \geq \beta_j$.

Proof of [App.5a.(1,1)]-[App.5d.(1,1)].

Recall that the payoffs $\bar{L}_i^{1,1}(\cdot)$ and $\bar{F}_i^{1,1}(\cdot)$ are defined as follows:

$$\bar{L}_i^{1,1}(\cdot) = t\pi^*(2-t) - c^*$$

$$\bar{F}_i^{1,1}(\cdot) = (1-t)\bar{\pi}_i^{1,0}.$$

Since $\bar{L}_i^{1,1}(\cdot) - \bar{F}_i^{1,1}(\cdot)$ is strictly increasing in t , we need only show that $\bar{L}_i^{1,1}(\bar{t}^{2,0}) - \bar{F}_i^{\beta_i,\beta_j+1}(\bar{t}^{\beta_i,\beta_j+1}) > 0$.

We have:

$$\begin{aligned} \bar{L}_i^{1,1}(\bar{t}^{2,0}) - \bar{F}_i^{\beta_i,\beta_j+1}(\bar{t}^{\beta_i,\beta_j+1}) &= \bar{t}^{2,0}\pi^*(2-\bar{t}^{2,0}) - c^* - (1-\bar{t}^{2,0})\bar{\pi}_i^{1,0} \\ &= \bar{\pi}_i^{2,0} - (1-\bar{t}^{2,0})\bar{\pi}_i^{1,0} - (1-\bar{t}^{2,0})\bar{\pi}_i^{1,0} \\ &= \bar{\pi}_i^{2,0} - 4\left[\bar{\pi}_i^{2,0} - \bar{\pi}_i^{1,0}\right] \\ &= 4\bar{\pi}_i^{1,0} - 3\bar{\pi}_i^{2,0} \\ &= 4(\pi^* - c^*) - 3\left[\pi^* - c^* + \frac{(\pi^* - c^*)^2}{4\pi^*}\right] \\ &> (\pi^* - c^*) - \frac{3}{4}(\pi^* - c^*) > 0. \end{aligned}$$

The first identity follows from (App.1) and the second from (App.2); the two strict inequalities follow from

$\pi^* > c^*$.

We next consider [A.5b.(1,1)]. Since $\bar{\Pi}_j^{2,0} = 0$, we need to prove that $\bar{\Pi}_i^{2,0} > 2\bar{\Pi}_i^{1,1}$. Now $\tilde{t}^{2,0}$ maximizes $\tilde{L}_i^{2,0}(\cdot) = t\pi^*(2-t) - c^* + (1-t)\bar{\Pi}_i^{1,0} = \tilde{L}_i^{1,1}(\cdot) + \bar{F}_i^{1,1}(\cdot)$. Necessarily, therefore, $\bar{\Pi}_i^{1,0} \geq \tilde{L}_i^{1,1}(\tilde{t}^{1,1}) + \bar{F}_i^{1,1}(\tilde{t}^{1,1}) = 2\bar{\Pi}_i^{1,1}$. Since $\tilde{L}_i^{2,0}(\cdot)$ is strictly concave in t and $\tilde{t}^{2,0} \neq \tilde{t}^{1,1}$, it follows that the above inequality is strict.

Statement [A.5c.(1,1)]-- j fires the bullet in phase (2,1)--was proved in section II above.

Now fix $\beta_i > 1$ and assume that statements [A.5a.($\beta_i-1,1$)]-[A.5c.($\beta_i-1,1$)] are true. We will show that [A.5a.($\beta_i,1$)]-[A.5c.($\beta_i,1$)] are true. We first show that $\tilde{t}^{\beta_i+1,0} > \tilde{t}^{\beta_i,1}$. From [A.5c.($\beta_i-1,1$)], we know that j fires the bullet in phase ($\beta_i,1$). Since $\tilde{L}_i^{\beta_i,1}(\cdot) - \bar{F}_i^{\beta_i,1}(\cdot)$ is strictly increasing, we need only check that $\tilde{L}_i^{\beta_i,1}(\tilde{t}^{\beta_i+1,0}) - \bar{F}_i^{\beta_i,1}(\tilde{t}^{\beta_i+1,0})$ is positive. We have:

$$\begin{aligned} \tilde{L}_i^{\beta_i,1}(\tilde{t}^{\beta_i+1,0}) - \bar{F}_i^{\beta_i,1}(\tilde{t}^{\beta_i+1,0}) &\equiv \tilde{t}^{\beta_i+1,0}\pi^*(2 - \tilde{t}^{\beta_i+1,0}) - c^* + (1 - \tilde{t}^{\beta_i+1,0})(\bar{\Pi}_i^{\beta_i-1,1} - \bar{\Pi}_i^{2,0}) \\ &\equiv \tilde{t}^{\beta_i+1,0}\pi^*(2 - \tilde{t}^{\beta_i+1,0}) - c^* + (1 - \tilde{t}^{\beta_i+1,0}) \left[(1 - \tilde{t}^{\beta_i-1,1})\bar{\Pi}_i^{\beta_i-1,0} - \bar{\Pi}_i^{2,0} \right] \\ &> \tilde{t}^{\beta_i+1,0}\pi^*(2 - \tilde{t}^{\beta_i+1,0}) - c^* + (1 - \tilde{t}^{\beta_i+1,0}) \left[(1 - \tilde{t}^{\beta_i,0})\bar{\Pi}_i^{\beta_i-1,0} - \bar{\Pi}_i^{2,0} \right] \\ &\equiv \left[\bar{\Pi}_i^{\beta_i+1,0} - (1 - \tilde{t}^{\beta_i+1,0})\bar{\Pi}_i^{2,0} \right] - (1 - \tilde{t}^{\beta_i+1,0}) \left[\bar{\Pi}_i^{2,0} - (1 - \tilde{t}^{\beta_i,0})\bar{\Pi}_i^{\beta_i-1,0} \right] \\ &\equiv \bar{\Pi}_i^{\beta_i+1,0} + 2 \left[\bar{\Pi}_i^{1,0} - \bar{\Pi}_i^{\beta_i+1,0} \right] - (1 - \tilde{t}^{\beta_i+1,0}) \left\{ \bar{\Pi}_i^{2,0} + 2 \left[\bar{\Pi}_i^{2,0} - \bar{\Pi}_i^{1,0} \right] \right\} \\ &\equiv \left[2\bar{\Pi}_i^{1,0} - \bar{\Pi}_i^{\beta_i+1,0} \right] - (1 - \tilde{t}^{\beta_i+1,0}) \left[2\bar{\Pi}_i^{1,0} - \bar{\Pi}_i^{2,0} \right] \\ &> 0. \end{aligned}$$

The first inequality follows because $\tilde{t}^{\beta_i,0} > \tilde{t}^{\beta_i-1,1}$. The next three identities follow from (A.1), (A.4) and (A.3). The second inequality follows from (A.3).

Statement [A.5b.($\beta_i,1$)] now follows easily. Since $\bar{\Pi}_j^{\beta_i+1,0} = 0$, we need to prove that

$$\bar{\Pi}_i^{\beta_i+1,0} > \bar{\Pi}_i^{\beta_i,1} + \bar{\Pi}_j^{\beta_i,1}. \quad \text{As before, } \tilde{t}^{\beta_i+1,0} \text{ maximizes}$$

$$\tilde{L}_i^{\beta_i+1,0}(\cdot) = t\pi^*(2-t) - c^* + (1-t)\bar{\Pi}_i^{2,0} = \tilde{L}_j^{\beta_i,1}(t) + \bar{F}_i^{\beta_i,1}(\cdot). \quad \text{Necessarily, therefore,}$$

$\bar{\Pi}_i^{\beta_i+1,0} \geq \bar{L}_j^{\beta_i,1}(\bar{t}^{\beta_i,1}) + \bar{F}_i^{\beta_i,1}(\bar{t}^{\beta_i,1})$. Since $\bar{L}_i^{\beta_i+1,0}(\cdot)$ is strictly concave in t and $\bar{t}^{\beta_i+1,0} \neq \bar{t}^{\beta_i,1}$, it follows that the above inequality is strict.

We now prove statement [A.5c.($\beta_i,1$)]. We have:

$$\bar{L}_i^{\beta_i+1,1}(t) = t\pi^*(2-t) - c^* + (1-t)\bar{\Pi}_i^{\beta_i,1}$$

$$\bar{F}_i^{\beta_i+1,1}(t) = (1-t)\bar{\Pi}_i^{\beta_i+1,0}$$

$$\bar{L}_j^{\beta_i+1,1}(t) = t\pi^*(2-t) - c^*$$

$$\bar{F}_j^{\beta_i+1,1}(t) = (1-t)\bar{\Pi}_j^{\beta_i,1}$$

To show that j fires, we need only check that $\bar{\Pi}_j^{\beta_i-1,1} + \bar{\Pi}_i^{\beta_i-1,1} < \bar{\Pi}_i^{2,0}$. As before, since $\bar{t}^{\beta_i,0}$ was chosen to maximize $\bar{L}_i^{\beta_i-1,1}(\cdot) + \bar{F}_j^{\beta_i-1,1}(\cdot)$, we have $\bar{\Pi}_j^{\beta_i-1,1} + \bar{\Pi}_i^{\beta_i-1,1} \leq \bar{\Pi}_i^{2,0}$. Repeating our earlier argument, the inequality must hold strictly.

Now fix $\beta_i \geq 1$. If $\beta_i > 1$, assume [A.5d.($\beta_i-1,1$)] is true. We will show [A.5d.($\beta_i,1$)] is true. We have

$$\bar{L}_i^{\beta_i+2,1}(t) = t\pi^*(2-t) - c^* + (1-t)\bar{\Pi}_i^{\beta_i+1,1}$$

$$\bar{F}_i^{\beta_i+2,1}(t) = (1-t)\bar{\Pi}_i^{\beta_i+2,0}$$

$$\bar{L}_i^{\beta_i+1,2}(t) = t\pi^*(2-t) - c^* + (1-t)\bar{\Pi}_i^{\beta_i,2}$$

$$\bar{F}_i^{\beta_i+1,2}(t) = (1-t)\bar{\Pi}_i^{\beta_i+1,1}$$

To establish $\bar{t}^{\beta_i+2,1} > \bar{t}^{\beta_i+1,2}$, it is sufficient to show $\bar{\Pi}_i^{\beta_i+1,1} - \bar{\Pi}_i^{\beta_i+2,0} < \bar{\Pi}_i^{\beta_i,2} - \bar{\Pi}_i^{\beta_i+1,1}$, i.e.,

$$2\bar{\Pi}_i^{\beta_i+1,1} - \bar{\Pi}_i^{\beta_i,2} < \bar{\Pi}_i^{\beta_i+2,0}. \quad (\text{A.6})$$

By [A.5c.($\beta_i,1$)], j fires in phase ($\beta_i+1,1$), so i is indifferent between firing and following. Therefore,

$\bar{\Pi}_i^{\beta_i+1,1} = \bar{L}_i^{\beta_i+1,1}(\bar{t}^{\beta_i+1,1}) = \bar{F}_i^{\beta_i+1,1}(\bar{t}^{\beta_i+1,1})$. If $\beta_i = 1$, then from [A.5c.($\beta_i,1$)], $\bar{\Pi}_j^{2,1} = \bar{L}_j^{2,1}(\bar{t}^{2,1})$, so that

$$\begin{aligned}
 2\bar{\Pi}_i^{2,1} - \bar{\Pi}_j^{2,1} &= 2\tilde{t}^{2,1}\pi^*(2 - \tilde{t}^{2,1}) - c^* + 2(1 - \tilde{t}^{2,1})\bar{\Pi}_i^{1,1} - \tilde{t}^{2,1}\pi^*(2 - \tilde{t}^{2,1}) - c^* \\
 &= \tilde{t}^{2,1}\pi^*(2 - \tilde{t}^{2,1}) - c^* + 2(1 - \tilde{t}^{2,1})\bar{\Pi}_i^{1,1} \\
 &< \tilde{t}^{2,1}\pi^*(2 - \tilde{t}^{2,1}) - c^* + (1 - \tilde{t}^{2,1})\bar{\Pi}_i^{2,0} \\
 &< \tilde{t}^{3,0}\pi^*(2 - \tilde{t}^{3,0}) - c^* + (1 - \tilde{t}^{3,0})\bar{\Pi}_i^{2,0} \\
 &\equiv \bar{\Pi}_i^{3,0}.
 \end{aligned}$$

The first inequality follows from [A.5b.(1,1)], the second from [A.5a.(2,1)]. If $\beta_i > 1$, we have:

$$\begin{aligned}
 2\bar{\Pi}_i^{\beta_i+1,1} - \bar{\Pi}_i^{\beta_i,2} &\leq \bar{L}_i^{\beta_i+1,1}(\tilde{t}^{\beta_i+1,1}) - \bar{F}_i^{\beta_i,2}(\tilde{t}^{\beta_i,2}) + \bar{F}_i^{\beta_i+1,1}(\tilde{t}^{\beta_i+1,1}) \\
 &\equiv \tilde{t}^{\beta_i+1,1}\pi^*(2 - \tilde{t}^{\beta_i+1,1}) - c^* + (1 - \tilde{t}^{\beta_i+1,1})\bar{\Pi}_i^{\beta_i,1} \\
 &\quad - (1 - \tilde{t}^{\beta_i,2})\bar{\Pi}_i^{\beta_i,1} + (1 - \tilde{t}^{\beta_i+1,1})\bar{\Pi}_i^{\beta_i+1,0} \\
 &< \tilde{t}^{\beta_i+1,1}\pi^*(2 - \tilde{t}^{\beta_i+1,1}) - c^* + (1 - \tilde{t}^{\beta_i+1,1})\bar{\Pi}_i^{\beta_i+1,0} \\
 &< \tilde{t}^{\beta_i+2,0}\pi^*(2 - \tilde{t}^{\beta_i+2,0}) - c^* + (1 - \tilde{t}^{\beta_i+2,0})\bar{\Pi}_i^{\beta_i+1,0} \\
 &\equiv \bar{\Pi}_i^{\beta_i+2,0}.
 \end{aligned}$$

The weak inequality holds because by [A.5c.($\beta_i, 1$)], j fires in phase $(\beta_i+1, 1)$, so that $\bar{\Pi}_i^{\beta_i+1,1} = \bar{L}_i^{\beta_i+1,1}(\tilde{t}^{\beta_i+1,1}) = \bar{F}_i^{\beta_i+1,1}(\tilde{t}^{\beta_i+1,1})$, while $\bar{\Pi}_i^{\beta_i,2} \geq \bar{F}_i^{\beta_i,2}(\tilde{t}^{\beta_i,2})$. The first strict inequality holds because by [A.5d.($\beta_i-1, 1$)], $\tilde{t}^{\beta_i+1,1} > \tilde{t}^{\beta_i,2}$, so that $(1 - \tilde{t}^{\beta_i+1,1})\bar{\Pi}_i^{\beta_i,1} - (1 - \tilde{t}^{\beta_i,2})\bar{\Pi}_i^{\beta_i,1} < 0$; the second strict inequality holds because $\tilde{t}^{\beta_i+2,0}$ maximizes $t\pi^*(2 - t) - c^* + (1 - t)\bar{\Pi}_i^{\beta_i+1,0}$ and by [A.5d.($\beta_i-1, 1$)], $\tilde{t}^{\beta_i+1,1} \neq \tilde{t}^{\beta_i+2,0}$.

The Inductive Step.

Now fix $\beta_j > 1$ and assume that for every $\beta_i \geq \beta_j - 1$, [A.5a.($\beta_i, \beta_j - 1$)]-[A.5d.($\beta_i, \beta_j - 1$)] are true. We will show that for every $\beta_i' \geq \beta_j$, [A.5a.(β_i, β_j)]-[A.5d.(β_i, β_j)] are true. To prove [A.5a.(β_i, β_j)], we need to verify that $\bar{L}_i^\beta(\cdot)$ intersects $\bar{F}_i^\beta(\cdot)$ at a time strictly before $\tilde{t}^{\beta_i+\beta_j, 0}$. Since $\bar{L}_i^\beta(\cdot) - \bar{F}_i^\beta(\cdot)$ is strictly increasing, we need only show that $\bar{L}_i^\beta(\tilde{t}^{\beta_i+\beta_j, 0}) - \bar{F}_i^{\beta_i, \beta_j}(\tilde{t}^{\beta_i+\beta_j, 0}) > 0$. We have:

$$\bar{L}_i^{\beta}(\bar{t}^{\beta_i+\beta_j,0}) - \bar{F}_i^{\beta_i,\beta_j}(\bar{t}^{\beta_i+\beta_j,0}) \quad (\text{A.7})$$

$$\begin{aligned} &\equiv \bar{t}^{\beta_i+\beta_j,0} \pi^*(2 - \bar{t}^{\beta_i+\beta_j,0}) - c^* + (1 - \bar{t}^{\beta_i+\beta_j,0})(\bar{\pi}_i^{\beta_i-1,\beta_j} - \bar{\pi}_i^{\beta_i,\beta_j-1}) \\ &\equiv \bar{t}^{\beta_i+\beta_j,0} \pi^*(2 - \bar{t}^{\beta_i+\beta_j,0}) - c^* + (1 - \bar{t}^{\beta_i+\beta_j,0}) \left[\prod_{\beta'_j=2}^{\beta_j} (1 - \bar{t}^{\beta_i-1,\beta'_j})(1 - \bar{t}^{\beta_i-1,1})\bar{\pi}_i^{\beta_i-1,0} - \prod_{\beta'_j=2}^{\beta_j} (1 - \bar{t}^{\beta_i,\beta'_j-1})\bar{\pi}_i^{2,0} \right] \\ &\geq \bar{t}^{\beta_i+\beta_j,0} \pi^*(2 - \bar{t}^{\beta_i+\beta_j,0}) - c^* + (1 - \bar{t}^{\beta_i+\beta_j,0}) \left[\prod_{\beta'_j=2}^{\beta_j} (1 - \bar{t}^{\beta_i-1,\beta'_j})(1 - \bar{t}^{\beta_i-1,1})\bar{\pi}_i^{\beta_i-1,0} - \prod_{\beta'_j=2}^{\beta_j} (1 - \bar{t}^{\beta_i-1,\beta'_j})\bar{\pi}_i^{2,0} \right] \\ &\equiv \bar{t}^{\beta_i+\beta_j,0} \pi^*(2 - \bar{t}^{\beta_i+\beta_j,0}) - c^* + (1 - \bar{t}^{\beta_i+\beta_j,0}) \prod_{\beta'_j=2}^{\beta_j} (1 - \bar{t}^{\beta_i-1,\beta'_j}) \left[(1 - \bar{t}^{\beta_i-1,1})\bar{\pi}_i^{\beta_i-1,0} - \bar{\pi}_i^{2,0} \right] \\ &> \bar{t}^{\beta_i+\beta_j,0} \pi^*(2 - \bar{t}^{\beta_i+\beta_j,0}) - c^* + (1 - \bar{t}^{\beta_i+\beta_j,0}) \prod_{\beta'_j=2}^{\beta_j} (1 - \bar{t}^{\beta_i+\beta'_j-1,0}) \left[(1 - \bar{t}^{\beta_i,0})\bar{\pi}_i^{\beta_i-1,0} - \bar{\pi}_i^{2,0} \right] \\ &\equiv \left[2\bar{\pi}_i^{1,0} - \bar{\pi}_i^{\beta_i+\beta_j,0} \right] - \prod_{\beta'_j=1}^{\beta_j} (1 - \bar{t}^{\beta_i+\beta'_j,0}) \left[2\bar{\pi}_i^{1,0} - \bar{\pi}_i^{2,0} \right] \\ &> \left[2\bar{\pi}_i^{1,0} - \bar{\pi}_i^{\beta_i+\beta_j,0} \right] - (1 - \bar{t}^{\beta_i+\beta_j,0}) \prod_{\beta'_j=2}^{\beta_j} (1 - \bar{t}^{\beta_i+\beta'_j,0}) \left[2\bar{\pi}_i^{1,0} - \bar{\pi}_i^{\beta_i+1,0} \right] \\ &\vdots \\ &> \left[2\bar{\pi}_i^{1,0} - \bar{\pi}_i^{\beta_i+\beta_j,0} \right] - (1 - \bar{t}^{\beta_i+\beta_j,0})(1 - \bar{t}^{\beta_i+\beta_j-1,0}) \left[2\bar{\pi}_i^{1,0} - \bar{\pi}_i^{\beta_i+\beta_j-1,0} \right] \\ &> 0. \end{aligned}$$

The first inequality will be true if $\prod_{\beta'_j=2}^{\beta_j} (1 - \bar{t}^{\beta_i-1,\beta'_j}) \geq \prod_{\beta'_j=2}^{\beta_j} (1 - \bar{t}^{\beta_i,\beta'_j-1})$. But this inequality is satisfied be-

cause if $\beta'_j = \beta_i$, then $\bar{t}^{\beta_i,\beta'_j-1} \equiv \bar{t}^{\beta_i-1,\beta'_j}$, while if $\beta'_j < \beta_i$, then by [A.5d.(β_i-2,β'_j-1)], the equality holds strictly. The second inequality holds because by [A.5d.(β_i,β'_j-1)], $\bar{t}^{\beta_i+\beta'_j-1,0} > \bar{t}^{\beta_i-1,\beta'_j}$ while by [A.5a.($\beta_i-1,1$)], $\bar{t}^{\beta_i,0} > \bar{t}^{\beta_i-1,1}$. The next identity is obtained by substitution, using (A.1) and (A.4). The last several inequalities follow from repeated application of (A.3).

We now prove [A.5b.(β_i, β_j)] and [A.5c.(β_i, β_j)], for $\beta_i \geq \beta_j$. If $\beta_i > \beta_j$, assume that [A.5b.($\beta_i - 1, \beta_j$)] and [A.5c.($\beta_i - 1, \beta_j$)] have already been proved. We first show that $\tilde{\Pi}_j^\beta + \tilde{\Pi}_i^\beta < \tilde{\Pi}_i^{\beta_i+1, \beta_j-1} + \tilde{\Pi}_j^{\beta_i+1, \beta_j-1}$. From [A.5d.($\beta_i, \beta_j - 1$)], j fires in phase $(\beta_i + 1, \beta_j - 1)$, so that:

$$\begin{aligned}\tilde{\Pi}_i^{\beta_i+1, \beta_j-1} &= \tilde{F}_i^{\beta_i+1, \beta_j-1}(\tilde{t}^{\beta_i+1, \beta_j-1}) = (1 - \tilde{t}^{\beta_i+1, \beta_j-1})\tilde{\Pi}_i^{\beta_i+1, \beta_j-2} \\ \tilde{\Pi}_j^{\beta_i+1, \beta_j-1} &= \tilde{L}_j^{\beta_i+1, \beta_j-1}(\tilde{t}^{\beta_i+1, \beta_j-1}) = \tilde{t}^{\beta_i+1, \beta_j-1} \pi^*(2 - \tilde{t}^{\beta_i+1, \beta_j-1}) - c^* + (1 - \tilde{t}^{\beta_i+1, \beta_j-1})\tilde{\Pi}_j^{\beta_i+1, \beta_j-2}\end{aligned}$$

Also, the payoffs to leading and following in phase (β_j, β_j) are equal, while if $\beta_i > \beta_j$, we know from [A.5b.($\beta_i - 1, \beta_j$)] that j leads in phase β . Therefore

$$\begin{aligned}\tilde{\Pi}_i^\beta &= \tilde{F}_i^\beta(\tilde{t}^\beta) = (1 - \tilde{t}^\beta)\tilde{\Pi}_i^{\beta, \beta_j-1} \\ \tilde{\Pi}_j^\beta &= \tilde{L}_j^\beta(\tilde{t}^\beta) = \tilde{t}^\beta \pi^*(2 - \tilde{t}^\beta) - c^* + (1 - \tilde{t}^\beta)\tilde{\Pi}_j^{\beta, \beta_j-1}\end{aligned}$$

Combining these expressions yields

$$\begin{aligned}\tilde{\Pi}_j^\beta + \tilde{\Pi}_i^\beta &\equiv \tilde{t}^\beta \pi^*(2 - \tilde{t}^\beta) - c^* + (1 - \tilde{t}^\beta)(\tilde{\Pi}_i^{\beta, \beta_j-1} + \tilde{\Pi}_j^{\beta, \beta_j-1}) \\ &< \tilde{t}^\beta \pi^*(2 - \tilde{t}^\beta) - c^* + (1 - \tilde{t}^\beta)(\tilde{\Pi}_i^{\beta_i+1, \beta_j-2} + \tilde{\Pi}_j^{\beta_i+1, \beta_j-2}) \\ &\equiv \tilde{F}_i^{\beta_i+1, \beta_j-1}(\tilde{t}^\beta) + \tilde{L}_j^{\beta_i+1, \beta_j-1}(\tilde{t}^\beta) \\ &< \tilde{F}_i^{\beta_i+1, \beta_j-1}(\tilde{t}^{\beta_i+1, \beta_j-1}) + \tilde{L}_j^{\beta_i+1, \beta_j-1}(\tilde{t}^{\beta_i+1, \beta_j-1}) \\ &\equiv \tilde{\Pi}_i^{\beta_i+1, \beta_j-1} + \tilde{\Pi}_j^{\beta_i+1, \beta_j-1}\end{aligned}$$

The first inequality follows from [A.5b.($\beta_i, \beta_j - 1$)]; the following identity holds because by [A.5c.($\beta_i, \beta_j - 1$)], j fires in phase $(\beta_i + 1, \beta_j - 1)$; the next inequality follows because from [A.5d.($\beta_i - 1, \beta_j - 1$)], $\tilde{t}^{\beta_i+1, \beta_j-1} > \tilde{t}^\beta$, while from [A.5a.($\beta_i + 1, \beta_j - 1$)], $\tilde{t}^{\beta_i+\beta_j, 0} \equiv \tilde{t}^{\beta_i+1+\beta_j-1, 0} > \tilde{t}^{\beta_i+1, \beta_j-1}$, and because aggregate payoffs are a strictly concave function of time. profits are higher when the firing time is closer to the optimum.

We now prove [A.5c.(β_i, β_j)], i.e., that j fires in phase $(\beta_i + 1, \beta_j)$. By the usual manipulations, we need to show that

$$\tilde{\Pi}_j^\beta + \tilde{\Pi}_i^\beta < \tilde{\Pi}_j^{\beta_i+1, \beta_j-1} + \tilde{\Pi}_i^{\beta_i+1, \beta_j-1}.$$

But this inequality follows immediately from step [A.5b.(β_i, β_j)] above.

Now assume that for every $\beta_i' \geq \beta_j$, [A.5b.(\(\beta_i', \beta_j\))] and [A.5c.(\(\beta_i', \beta_j\))] are true. Fix $\beta_i \geq \beta_j$ and, if $\beta_i > \beta_j$, assume that [A.5d.(\(\beta_i - 1, \beta_j\))] is true. We will prove [A.5d.(\(\beta_i, \beta_j\))]. We need to show that whenever i weakly prefers following to leading in phase $(\beta_i + 1, \beta_j + 1)$, he strictly prefers following to leading at this time in phase $(\beta_i + 2, \beta_j)$.

As usual, i 's payoffs to leading and following in the relevant phases are:

$$\begin{aligned}\bar{L}_i^{\beta_i+2, \beta_j}(t) &= t\pi^*(2-t) - c^* + (1-t)\bar{\Pi}_i^{\beta_i+1, \beta_j} \\ \bar{F}_i^{\beta_i+2, \beta_j}(t) &= (1-t)\bar{\Pi}_i^{\beta_i+2, \beta_j-1} \\ \bar{L}_i^{\beta_i+1, \beta_j+1}(t) &= t\pi^*(2-t) - c^* + (1-t)\bar{\Pi}_i^{\beta_i, \beta_j+1} \\ \bar{F}_i^{\beta_i+1, \beta_j+1}(t) &= (1-t)\bar{\Pi}_i^{\beta_i+1, \beta_j}\end{aligned}$$

To show that $\bar{L}_i^{\beta_i+2, \beta_j} > \bar{L}_i^{\beta_i+1, \beta_j+1}$, it is sufficient to show that

$$\bar{\Pi}_i^{\beta_i+1, \beta_j} - \bar{\Pi}_i^{\beta_i+2, \beta_j-1} < \bar{\Pi}_i^{\beta_i, \beta_j+1} - \bar{\Pi}_i^{\beta_i+1, \beta_j},$$

i.e., that

$$\bar{\Pi}_i^{\beta_i+2, \beta_j-1} > 2\bar{\Pi}_i^{\beta_i+1, \beta_j} - \bar{\Pi}_i^{\beta_i, \beta_j+1}. \quad (\text{A.6})$$

We have:

$$\begin{aligned}& \bar{\Pi}_i^{\beta_i+1, \beta_j} - \bar{\Pi}_i^{\beta_i, \beta_j+1} + \bar{\Pi}_i^{\beta_i+1, \beta_j} \\ & \leq \bar{L}_i^{\beta_i+1, \beta_j}(\bar{t}^{\beta_i+1, \beta_j}) - \bar{F}_i^{\beta_i, \beta_j+1}(\bar{t}^{\beta_i, \beta_j+1}) + \bar{F}_i^{\beta_i+1, \beta_j}(\bar{t}^{\beta_i+1, \beta_j}) \\ & = \bar{t}^{\beta_i+1, \beta_j}\pi^*(2 - \bar{t}^{\beta_i+1, \beta_j}) - c^* + (1 - \bar{t}^{\beta_i+1, \beta_j})\bar{\Pi}_i^{\beta_i} \\ & \quad - (1 - \bar{t}^{\beta_i, \beta_j+1})\bar{\Pi}_i^{\beta_i} + (1 - \bar{t}^{\beta_i+1, \beta_j})\bar{\Pi}_i^{\beta_i+1, \beta_j-1} \\ & \leq \bar{t}^{\beta_i+1, \beta_j}\pi^*(2 - \bar{t}^{\beta_i+1, \beta_j}) - c^* + (1 - \bar{t}^{\beta_i+1, \beta_j})\bar{\Pi}_i^{\beta_i+1, \beta_j-1} \\ & < \bar{t}^{\beta_i+2, \beta_j-1}\pi^*(2 - \bar{t}^{\beta_i+2, \beta_j-1}) - c^* + (1 - \bar{t}^{\beta_i+2, \beta_j-1})\bar{\Pi}_i^{\beta_i+1, \beta_j-1} \\ & = \bar{\Pi}_i^{\beta_i+2, \beta_j-1}\end{aligned}$$

The first weak inequality holds because by [A.5c.(\(\beta_i, \beta_j\))], j fires in phases $(\beta_i + 1, \beta_j)$ and $(\beta_i, \beta_j + 1)$, so that

$\bar{\Pi}_i^{\beta_i+1, \beta_j} = \bar{L}_i^{\beta_i+1, \beta_j}(\bar{t}^{\beta_i+1, \beta_j}) = \bar{F}_i^{\beta_i+1, \beta_j}(\bar{t}^{\beta_i+1, \beta_j})$, while $\bar{\Pi}_i^{\beta_i, \beta_j+1} \geq \bar{F}_i^{\beta_i, \beta_j+1}(\bar{t}^{\beta_i, \beta_j+1})$. The second weak inequality

holds because if $\beta_i = \beta_j$, then $\tilde{t}^{\beta_i+1,\beta_j} \equiv \tilde{t}^{\beta_i,\beta_j+1}$, so that $(1 - \tilde{t}^{\beta_i+1,\beta_j})\tilde{\Pi}_i^\beta = (1 - \tilde{t}^{\beta_i,\beta_j+1})\tilde{\Pi}_i^\beta$; if $\beta_i > \beta_j$, then by [A.5d.(\(\beta_i-1,\beta_j\))], $\tilde{t}^{\beta_i+1,\beta_j} > \tilde{t}^{\beta_i,\beta_j+1}$, so that $(1 - \tilde{t}^{\beta_i+1,\beta_j})\tilde{\Pi}_i^\beta < (1 - \tilde{t}^{\beta_i,\beta_j+1})\tilde{\Pi}_i^\beta$; The strict inequality holds because by [A.5d.(\(\beta_i,\beta_j-1\))] and [A.5a.(\(\beta_i+2,\beta_j-1\))], $\tilde{t}^{\beta_i+1,\beta_j} < \tilde{t}^{\beta_i+2,\beta_j-1} < \tilde{t}^{\beta_i+2+\beta_j-1,0}$, while for $t < \tilde{t}^{\beta_i+2+\beta_j-1,0}$, $t\pi^*(2-t) - c^* + (1-t)\tilde{\Pi}_i^{\beta_i+1,\beta_j-1}$ is strictly increasing in t . This completes our verification of [A.5d.(\(\beta_i,\beta_j\))] and completes the proof.

Proof of Theorem V.

Fix a triple (β^*, π^*, c^*) and assume that these parameters pass the test specified in Proposition I. The following facts are either obvious or have been checked in the text: (i) for each monopoly phase history h , the profile $f^*(\cdot, h)$ is the unique equilibrium profile for every subgame beginning at (t, h) , $t \geq t^{last}(h)$; (ii) for each β_i , $\Pi^{\beta_i,0} = \tilde{\Pi}^{\beta_i,0}$ and $\tilde{t}^{\beta_i,0} = t^{\beta_i,0}$; (iii) $t^{1,1} = \tilde{t}^{1,1}$; (iv) for any phase (1,1) history h such that $t^{last}(h) < t^{1,1}$, we have $P(f^*, t^{last}(h), h) = \Pi^{1,1} = \tilde{\Pi}^{1,1}$, where $\Pi^{1,1}$ and $\tilde{\Pi}^{1,1}$ were defined, respectively on p. and p. .

Now fix a duopoly phase β , and a phase β history, h , such that $\beta_i \geq \beta_j > 0$ and $t^{last}(h) < t^\beta$. Assume that for each $\beta' \in \{(n,\beta_j-1), 0 \leq n \leq \beta_i^*\} \cup \{(n,\beta_j), 0 \leq n < \beta_i\}$, the following facts have been established: (a) $t^{\beta'} = \tilde{t}^{\beta'}$; (b) for any phase β' history h such that $t^{last}(h) < \tilde{t}^{\beta'}$, $P(f^*, t^{last}(h), h) = \Pi^{\beta'} = \tilde{\Pi}^{\beta'}$; (c) if β' is a duopoly phase, then the payoff functions for the duel $CD^{\beta'}$ satisfy conditions (A1)-(A4). We will establish that (a), (b) and (c) are true for β . It will then follow immediately from Proposition IV that for each α and τ , the profile $\xi^{(\tau; \alpha; \beta^*)}(\cdot)$ is a subgame perfect equilibrium for the subduel $CD^\beta(\tau)$. In turn, it will follow immediately from expressions (4.2) and (4.3) and the comment in footnote 30 that if in the multi-stage duel, either agent or both fires at some decision node (t, h) , where $t < \tilde{t}^{\beta_i-1,\beta_j}$ and h is a phase β history, agents' expected payoffs will be exactly the same as if the same set of agents had fired at t in CD^β . Finally, we can conclude that from any decision node (τ, h) such that h is a phase β history, f^* is a Nash equilibrium for the subgame beginning at this node.

Since (β^*, π^*, c^*) passes our test, $\tilde{t}^\beta < \tilde{t}^{\beta_i-1,\beta_j}$, so that CD^β has an equilibrium in which the bullet is

fired at $\tilde{t}^\beta = t^\beta$, yielding payoffs $\tilde{\Pi}^\beta$. Since by assumption, the functions defined by expressions (2.6.a')-(2.6.d') and those defined by (2.6.a)-(2.6.d), it follows immediately that $\tilde{\Pi}^\beta = \Pi^\beta$.

We now check that the payoffs in CD^β satisfy conditions (A1)-(A4). It follows from expression (5.2) that for each $\beta' \in IS(\beta)$, the $\Psi_i(\cdot; \alpha; \beta')$'s will be right continuous, piecewise polynomials, provided that $L_i^{\beta'}(t)$ and $F_i^{\beta'}(t)$ are also functions of this form. But this is true by assumption. Now set $T = t^{\beta_i-1, \beta_j}$ in assumption (A3) above. We will argue that $L_i^\beta(\cdot)$ is strictly increasing and continuous on $[0, T)$. The argument for j is similar. For each $t \in [0, t^{\beta_i-1, \beta_j})$, $\Psi_i(t; \{i\}; (\beta_i-1, \beta_j)) = \Psi_i(t^{\beta_i-1, \beta_j}; \{i\}; (\beta_i-1, \beta_j))$. Therefore, $L_i^\beta(t) = \pi^*(2-t) - c^* + (1-t)\Psi_i(t^{\beta_i-1, \beta_j}; \{i\}; (\beta_i-1, \beta_j))$ on $[0, T)$. Next, observe that, obviously, $\Psi_i(t^{\beta_i-1, \beta_j}; \{i\}; (\beta_i-1, \beta_j)) < \Psi_i(0; \{i\}; (\beta_i+\beta_j, 0))$, where the right hand side is i 's payoff when he is the monopolist in the monopoly phase that corresponds to phase (β_i-1, β_j) . Now the monopolist chooses $t^{\beta_i+\beta_j, 0}$ to maximize $\pi^*(2-t) - c^* + (1-t)\Psi_i(0; \{i\}; (\beta_i+\beta_j, 0))$. From [A.5a], $t^\beta < t^{\beta_i+\beta_j, 0}$, so that $\frac{\partial}{\partial t}(\pi^*(2-t) - c^* + (1-t)\Psi_i(0; \{i\}; (\beta_i+\beta_j, 0))) = (\frac{\partial(\pi^*(2-t) - c^*)}{\partial t} - \Psi_i(0; \{i\}; (\beta_i+\beta_j, 0))) > 0$ on $[0, t^\beta]$. Therefore $\frac{\partial L_i^\beta(t)}{\partial t} = (\frac{\partial(\pi^*(2-t) - c^*)}{\partial t} - \Psi_i(t^{\beta_i-1, \beta_j}; \{i\}; (\beta_i-1, \beta_j)))$ must also be positive on the same interval. Next, note that because (β^*, π^*, c^*) passes our test, the payoff functions defined by (5.3.a)-(5.3.f) coincide with the "test functions" defined in expression (2.6.a'), which, by assumption, intersect before T . Moreover, these functions are certainly continuous on $[0, T)$. We have now verified, therefore, that these payoffs satisfy conditions (A1), (A2a) and (A3).

We now check condition (A2.b), i.e., that beyond t^β , each agent strictly prefers leading to following. To prove this, we must consider the explicit values of the "continuation payoff functions," $\Psi_i(\cdot; \{i\}; (\beta_i-1, \beta_j))$ and $\Psi_i(\cdot; \{j\}; (\beta_i, \beta_j-1))$, on this interval. There are three cases to consider, depending on whether the game is symmetric and, if it is not, whether i is the weak or the strong player.

First assume that i is the weak player, i.e., $\beta_i < \beta_j$. Note that from [A.5a], we have $t^{\beta_i, \beta_j-1} < t^{\beta_i-1, \beta_j}$, i.e., the firing time in the "less asymmetric" successor phase, (β_i, β_j-1) , occurs earlier than in the "more asymmetric" phase, (β_i-1, β_j) . If i fires and misses at $t \in [t^\beta, t^{\beta_i, \beta_j-1})$, he will the next bullet with probabili-

ty one at t^{β_i-1, β_j} . Therefore, for each such t in this interval, $\Psi_i(t; \{i\}; (\beta_i-1, \beta_j)) = L_i^{\beta_i-1, \beta_j}(t^{\beta_i-1, \beta_j})$. If j fires and misses in this interval, then either i fires at t^{β_i, β_j-1} or, if $\beta_i = \beta_j-1$ they both randomize immediately after this time. In either case, $\Psi_i(t; \{j\}; (\beta_i, \beta_j-1)) = L_i^{\beta_i, \beta_j-1}(t^{\beta_i, \beta_j-1})$.

If i fires and misses at $t \in [t^{\beta_i, \beta_j-1}, t^{\beta_i-1, \beta_j})$, then he will, once again, fire again at t^{β_i-1, β_j} . Therefore, once again, $\Psi_i(t; \{i\}; (\beta_i-1, \beta_j)) = L_i^{\beta_i-1, \beta_j}(t^{\beta_i-1, \beta_j})$. If j fires and misses in this interval, then i will fire the next bullet immediately. In this case, $\Psi_i(t; \{j\}; (\beta_i, \beta_j-1)) = L_i^{\beta_i, \beta_j-1}(t)$.

If i fires and misses at $t \in [t^{\beta_i-1, \beta_j}, t^{\beta_i-1, \beta_j-1})$, then j will fire immediately afterwards; if j 's bullet misses, i gets to fire again at t^{β_i-1, β_j-1} . Therefore, on this interval, $\Psi_i(t; \{i\}; (\beta_i-1, \beta_j)) = F_i^{\beta_i-1, \beta_j}(t) = (1-t)L_i^{\beta_i-1, \beta_j-1}(t^{\beta_i-1, \beta_j-1})$. If j fires, then i fires immediately afterwards; if i 's bullet misses, j gets to fire again at t^{β_i-1, β_j-1} . Therefore, $\Psi_i(t; \{j\}; (\beta_i, \beta_j-1)) = L_i^{\beta_i, \beta_j-1}(t) = \pi\pi^*(2-t) - c^* + (1-t)F_i^{\beta_i-1, \beta_j-1}(t^{\beta_i-1, \beta_j-1})$.

Finally, if i fires and misses at $t > t^{\beta_i-1, \beta_j-1}$, j fires immediately, and, if j misses, i fires again immediately. Similarly, if j fires and misses in this interval, i fires immediately, and, if i misses, j fires again immediately. Therefore, for $t \in (t^{\beta_i-1, \beta_j-1}, 2)$, $\Psi_i(t; \{i\}; (\beta_i-1, \beta_j)) = F_i^{\beta_i-1, \beta_j}(t) = (1-t)L_i^{\beta_i-1, \beta_j-1}(t)$, while $\Psi_i(t; \{j\}; (\beta_i, \beta_j-1)) = L_i^{\beta_i, \beta_j-1}(t) = \pi\pi^*(2-t) - c^* + (1-t)F_i^{\beta_i-1, \beta_j-1}(t)$.

Substituting for $\Psi_i(t; \{i\}; (\beta_i-1, \beta_j))$ and $\Psi_i(t; \{j\}; (\beta_i, \beta_j-1))$ in (A.3.a) and (A.3.b), we have

$$\text{if } t \in [t^\beta, t^{\beta-1, \beta_j}), \quad \begin{pmatrix} L_i^\beta(t) \\ F_i^\beta(t) \end{pmatrix} = \begin{pmatrix} t\pi^*(2-t) - c^* + (1-t)L_i^{\beta-1, \beta_j}(t^{\beta-1, \beta_j}) \\ (1-t)L_i^{\beta, \beta_j-1}(t^{\beta, \beta_j-1}) \end{pmatrix} \quad (\text{A.8.a})$$

$$\text{if } t \in [t^{\beta-1, \beta_j}, t^{\beta, \beta_j-1}), \quad \begin{pmatrix} L_i^\beta(t) \\ F_i^\beta(t) \end{pmatrix} = \begin{pmatrix} t\pi^*(2-t) - c^* + (1-t)L_i^{\beta-1, \beta_j}(t^{\beta-1, \beta_j}) \\ (1-t)L_i^{\beta, \beta_j-1}(t) \end{pmatrix} \quad (\text{A.8.b})$$

$$\text{if } t \in [t^{\beta-1, \beta_j}, t^{\beta-1, \beta_j-1}), \quad \begin{pmatrix} L_i^\beta(t) \\ F_i^\beta(t) \end{pmatrix} = \begin{pmatrix} t\pi^*(2-t) - c^* + (1-t)^2 L_i^{\beta-1, \beta_j-1}(t^{\beta-1, \beta_j-1}) \\ (1-t)t\pi^*(2-t) - c^* + (1-t)^2 F_i^{\beta-1, \beta_j-1}(t^{\beta-1, \beta_j-1}) \end{pmatrix} \quad (\text{A.8.c})$$

$$\text{if } t \in (t^{\beta-1, \beta_j-1}, 2), \quad \begin{pmatrix} L_i^\beta(t) \\ F_i^\beta(t) \end{pmatrix} = \begin{pmatrix} t\pi^*(2-t) - c^* + (1-t)^2 L_i^{\beta-1, \beta_j-1}(t) \\ (1-t)t\pi^*(2-t) - c^* + (1-t)^2 F_i^{\beta-1, \beta_j-1}(t) \end{pmatrix} \quad (\text{A.8.d})$$

We will show that $L_i^\beta(\cdot) > F_i^\beta(\cdot)$ beyond t^β .

First, consider $t \in [t^\beta, t^{\beta, \beta_j-1})$. In this interval,

$$L_i^\beta(t) - F_i^\beta(t) = t\pi^*(2-t) - c^* + (1-t)[L_i^{\beta-1, \beta_j}(t^{\beta-1, \beta_j}) - L_i^{\beta, \beta_j-1}(t^{\beta, \beta_j-1})]$$

There are two cases to consider. If $L_i^{\beta-1, \beta_j}(t^{\beta-1, \beta_j}) \geq L_i^{\beta, \beta_j-1}(t^{\beta, \beta_j-1})$, it is immediate that $L_i^\beta(\cdot) > F_i^\beta(\cdot)$ on $[t^\beta, t^{\beta, \beta_j-1})$. Suppose therefore that $L_i^{\beta-1, \beta_j}(t^{\beta-1, \beta_j}) < L_i^{\beta, \beta_j-1}(t^{\beta, \beta_j-1})$. In this case,

$$\frac{\partial(L_i^\beta(\cdot) - F_i^\beta(\cdot))}{\partial t} = \frac{\partial(t\pi^*(2-t) - c^*)}{\partial t} - [L_i^{\beta-1, \beta_j}(t^{\beta-1, \beta_j}) - L_i^{\beta, \beta_j-1}(t^{\beta, \beta_j-1})] > 0 \text{ on } [t^\beta, t^{\beta, \beta_j-1}). \text{ By}$$

[A.5c], we know that $L_i^\beta(t^\beta) > F_i^\beta(t^\beta)$; it therefore follows that $L_i^\beta(\cdot) > F_i^\beta(\cdot)$ on $[t^\beta, t^{\beta, \beta_j-1})$.

Next, consider $t \in [t^{\beta, \beta_j-1}, t^{\beta-1, \beta_j})$. On this interval,

$$L_i^\beta(t) - F_i^\beta(t) = t\pi^*(2-t) - c^* + (1-t)[L_i^{\beta-1, \beta_j}(t^{\beta-1, \beta_j}) - L_i^{\beta, \beta_j-1}(t)] \quad (\text{A.9})$$

Now $L_i^{\beta, \beta_j-1}(t) = t\pi^*(2-t) - c^* + (1-t)L_i^{\beta-1, \beta_j-1}(t^{\beta-1, \beta_j-1})$. Moreover, the second term on the right hand side is just $F_i^{\beta-1, \beta_j}(t)$. Substituting in (A.9) and rearranging terms, we have

$$L_i^\beta(t) - F_i^\beta(t) = t\pi^*(2-t) - c^* + (1-t)[L_i^{\beta-1, \beta_j}(t^{\beta-1, \beta_j}) - F_i^{\beta-1, \beta_j}(t)]$$

Once again, there are two cases to consider. First suppose that for some $s \in [t^{\beta, \beta_j-1}, t^{\beta-1, \beta_j})$,

$L_i^{\beta-1, \beta_j}(t^{\beta-1, \beta_j}) < F_i^{\beta-1, \beta_j}(s)$. For this s , we have

$$\frac{\partial(L_i^\beta(s) - F_i^\beta(s))}{\partial t} = \frac{\partial(s\pi^*(2-s) - c^*)}{\partial t} + (1-s) \frac{-\partial F_i^{\beta_i-1, \beta_j}(s)}{\partial t} + [F_i^{\beta_i-1, \beta_j}(s) - L_i^{\beta_i-1, \beta_j}(t^{\beta_i-1, \beta_j})]$$

Since all three terms are positive, it follows that $(L_i^\beta(\cdot) - F_i^\beta(\cdot))$ is strictly increasing at every $s \in [t^{\beta_i-1, \beta_j}, t^{\beta_i-1, \beta_j})$ such that $L_i^{\beta_i-1, \beta_j}(t^{\beta_i-1, \beta_j}) < F_i^{\beta_i-1, \beta_j}(s)$. Moreover, we have already established that $L_i^\beta(t^{\beta_i, \beta_j-1}) > F_i^\beta(t^{\beta_i, \beta_j-1})$. It follows that $L_i^\beta(s) > F_i^\beta(s)$ for each such s . Now assume that there exists $s \in [t^{\beta_i-1, \beta_j}, t^{\beta_i-1, \beta_j})$ such that $L_i^{\beta_i-1, \beta_j}(t^{\beta_i-1, \beta_j}) \geq F_i^{\beta_i-1, \beta_j}(s)$. Since $F_i^{\beta_i-1, \beta_j}(\cdot)$ is strictly decreasing before t^{β_i-1, β_j} , the expression $[\cdot]$ must be positive on the interval $[s, t^{\beta_i-1, \beta_j})$. so that $L_i^\beta(\cdot) > F_i^\beta(\cdot)$ on this interval also.

For $t \geq t^{\beta_i-1, \beta_j-1}$, the inequality follows immediately from our hypothesis that the payoffs for $CD^{\beta_i-1, \beta_j-1}$ satisfy (A2.b), so that $L_i^{\beta_i-1, \beta_j-1}(\cdot) > F_i^{\beta_i-1, \beta_j-1}(\cdot)$ beyond t^{β_i-1, β_j-1} . This completes our verification of assumption (A2.b) when i is the weak player.

Condition (A2.c) is easy to check. This is the requirement that at any discontinuity point t of either $L_i^\beta(\cdot)$ or $F_i^\beta(\cdot)$ beyond t^β , i can do at least as well by leading just before t as he can by following at t . From (A.8.a)-(A.8.c), it is clear that $F_i^\beta(\cdot)$ is continuous on $[t^\beta, t^{\beta_i-1, \beta_j-1})$ (because $L_i^{\beta_i, \beta_j-1}(\cdot)$ is continuous on $[0, t^{\beta_i-1, \beta_j-1}]$). On other other hand, $L_i^\beta(\cdot)$ is continuous on this interval except at t^{β_i-1, β_j} . Since $L_i^\beta(\cdot)$ exceeds $F_i^\beta(\cdot)$ at t^{β_i-1, β_j} and $L_i^\beta(\cdot)$ jumps *down* at this point, condition (A2.c) is certainly satisfied. Beyond t^{β_i-1, β_j-1} , (A2.c) follows from our hypothesis that the same condition is satisfied for $L_i^{\beta_i-1, \beta_j-1}(\cdot)$.

We now assume that i is the strong player, i.e., that $\beta_i > \beta_j$. In this case, the ordering of the firing times is reversed, i.e., $t^{\beta_i-1, \beta_j} < t^{\beta_i, \beta_j-1}$. If i fires and misses at $t \in [t^\beta, t^{\beta_i-1, \beta_j})$, then either j fires at t^{β_i-1, β_j} or, if $\beta_i-1 = \beta_j$, both players randomize immediately after this time. In either case, $\Psi_i(t; \{j\}; (\beta_i-1, \beta_j)) = L_i^{\beta_i-1, \beta_j}(t^{\beta_i-1, \beta_j})$. If j fires and misses at t in this interval then he will not fire again until t^{β_i, β_j-1} so that $\Psi_i(t; \{j\}; (\beta_i, \beta_j-1)) = F_i^{\beta_i, \beta_j-1}(t^{\beta_i, \beta_j-1})$.

Now suppose that i fires and misses at $t \in [t^{\beta_i-1, \beta_j}, t^{\beta_i, \beta_j-1})$. In this case, j will fire immediately, so that $\Psi_i(t; \{j\}; (\beta_i-1, \beta_j)) = F_i^{\beta_i-1, \beta_j}(t)$. If j fires and misses in this interval, then he again hold his fire until

$t^{\beta_i, \beta_j - 1}$ so that, again, $\Psi_i(t; \{j\}; (\beta_i, \beta_j - 1)) = F_i^{\beta_i, \beta_j - 1}(t^{\beta_i, \beta_j - 1})$.

Beyond $t^{\beta_i - 1, \beta_j - 1}$, the definitions of $\Psi_i(t; \{i\}; (\beta_i - 1, \beta_j))$ and $\Psi_i(t; \{j\}; (\beta_i, \beta_j - 1))$ are independent of whether i is strong or weak, so we can ignore this case.

Once again, we substitute for $\Psi_i(t; \{i\}; (\beta_i - 1, \beta_j))$ and $\Psi_i(t; \{j\}; (\beta_i, \beta_j - 1))$ in (App.3.a) and (App.3.b) to obtain:

$$\begin{aligned} \text{if } t \in [t^\beta, t^{\beta_i - 1, \beta_j}), & \quad \begin{pmatrix} L_i^\beta(t) \\ F_i^\beta(t) \end{pmatrix} = \begin{pmatrix} t\pi^*(2-t) - c^* + (1-t)F_i^{\beta_i - 1, \beta_j}(t^{\beta_i - 1, \beta_j}) \\ (1-t)F_i^{\beta_i, \beta_j - 1}(t^{\beta_i, \beta_j - 1}) \end{pmatrix} \\ \text{if } t \in [t^{\beta_i - 1, \beta_j}, t^{\beta_i, \beta_j - 1}), & \quad \begin{pmatrix} L_i^\beta(t) \\ F_i^\beta(t) \end{pmatrix} = \begin{pmatrix} t\pi^*(2-t) - c^* + (1-t)F_i^{\beta_i - 1, \beta_j}(t) \\ (1-t)F_i^{\beta_i, \beta_j - 1}(t^{\beta_i, \beta_j - 1}) \end{pmatrix} \end{aligned}$$

We now show that $L_i^\beta(\cdot) > F_i^\beta(\cdot)$ on the interval $[t^\beta, t^{\beta_i, \beta_j - 1})$. Observe that $L_i^\beta(\cdot)$ is concave on this interval while $F_i^\beta(\cdot)$ is affine. Therefore, $(L_i^\beta(\cdot) - F_i^\beta(\cdot))$ is concave. Also, $L_i^\beta(t^\beta) = F_i^\beta(t^\beta)$ while

$$\begin{aligned} L_i^\beta(t^{\beta_i, \beta_j - 1}) - F_i^\beta(t^{\beta_i, \beta_j - 1}) &= t^{\beta_i, \beta_j - 1} \pi^*(2 - t^{\beta_i, \beta_j - 1}) - c^* + (1 - t^{\beta_i, \beta_j - 1}) [F_i^{\beta_i - 1, \beta_j}(t^{\beta_i, \beta_j - 1}) - L_i^{\beta_i, \beta_j - 1}(t^{\beta_i - 1, \beta_j})] \\ &= t^{\beta_i, \beta_j - 1} t^{\beta_i, \beta_j - 1} \pi^*(2 - t^{\beta_i, \beta_j - 1}) - c^* + (1 - t^{\beta_i, \beta_j - 1})^2 [F_i^{\beta_i - 1, \beta_j - 1}(t^{\beta_i - 1, \beta_j - 1}) - F_i^{\beta_i - 1, \beta_j - 1}(t^{\beta_i - 1, \beta_j - 1})] \\ &= t^{\beta_i, \beta_j - 1} t^{\beta_i, \beta_j - 1} \pi^*(2 - t^{\beta_i, \beta_j - 1}) - c^* > 0 \end{aligned}$$

Therefore, $(L_i^\beta(\cdot) - F_i^\beta(\cdot))$ must be strictly positive throughout $(t^\beta, t^{\beta_i, \beta_j - 1})$. This completes our verification of assumption (A2.b) when i is the strong player.

Condition (A2.c) is immediate in this case, because both $L_i^\beta(\cdot)$ and $F_i^\beta(\cdot)$ are continuous on $[t^\beta, t^{\beta_i - 1, \beta_j - 1})$ (because $F_i^{\beta_i - 1, \beta_j}(\cdot)$ is continuous on $[0, t^{\beta_i - 2, \beta_j}] \supset [0, t^{\beta_i - 1, \beta_j - 1})$). Beyond $t^{\beta_i - 1, \beta_j - 1}$, (A2.c) again follows from our hypothesis that the same condition is satisfied for $L_i^{\beta_i - 1, \beta_j - 1}(\cdot)$.

Finally, we need to check the symmetric case in which $\beta_i = \beta_j$. To verify this case, we need only patch together the easier parts of the arguments given above. First note in this case, $t^{\beta_i - 1, \beta_j}$ is identically equal to $t^{\beta_i, \beta_j - 1}$. Beyond $t^{\beta_i - 1, \beta_j}$, $(L_i^\beta(\cdot) - F_i^\beta(\cdot))$ is strictly positive for the reasons that we gave before. Between t^β and $t^{\beta_i - 1, \beta_j}$, $(L_i^\beta(\cdot) - F_i^\beta(\cdot))$ is a concave function which is zero at one end point and positive at the other. It must therefore be positive throughout the interval. This completes our verification that for each

β , $L_i^\beta(\cdot) - F_i^\beta(\cdot)$ is strictly positive on the interval $(t^\beta, 2)$ and so completes our verification of assumption (A2.b). Condition (A2.c) is again immediate in this case.

We now check condition (A4), i.e., that $F_i^\beta(\cdot)$ exceeds $S_i^\beta(\cdot)$ beyond t^β . Substituting for $\Psi_i(t; \{j\}; (\beta_i, \beta_j - 1))$ and $\Psi_i(t; \{i, j\}; (\beta_i - 1, \beta_j - 1))$ in (A.3.b) and (A.3.c), we have

$$F_i^\beta(t) - S_i^\beta(t) = (1 - t) \left[\Psi_i(t; \{j\}; (\beta_i, \beta_j - 1)) - \left(t\pi^*(2 - t) - c^* + (1 - t)\Psi_i(t; \{i, j\}; (\beta_i - 1, \beta_j - 1)) \right) \right] - tc^*.$$

As usual, there are three regions to consider. If $t \in [t^\beta, t^{\beta_i, \beta_j - 1})$, the term between the square brackets can be written as $(L_i^{\beta_i, \beta_j - 1}(t^{\beta_i, \beta_j - 1}) - L_i^{\beta_i, \beta_j - 1}(t))$, which is strictly positive for all $t < t^{\beta_i, \beta_j - 1}$. If $t \in [t^{\beta_i, \beta_j - 1}, t^{\beta_i - 1, \beta_j - 1})$, then this term becomes $(L_i^{\beta_i, \beta_j - 1}(t^{\beta_i, \beta_j - 1}) - L_i^{\beta_i, \beta_j - 1}(t^{\beta_i, \beta_j - 1})) = 0$. In either case, therefore, the whole expression is positive. Finally, if $t > t^{\beta_i - 1, \beta_j - 1}$, then $[\cdot]$ expands to

$$\left(t\pi^*(2 - t) - c^* + (1 - t)\Psi_i(t; \{i\}; (\beta_i - 1, \beta_j - 1)) \right) - \left(t\pi^*(2 - t) - c^* + (1 - t)\Psi_i(t; \{i, j\}; (\beta_i - 1, \beta_j - 1)) - tc^* \right).$$

In this case, $\Psi_i(\cdot; \{j\}; (\beta_i - 1, \beta_j - 1))$ equals $\Psi_i(\cdot; \{i, j\}; (\beta_i - 1, \beta_j - 1))$, so that $[\cdot] = tc^*$ beyond $t^{\beta_i - 1, \beta_j - 1}$.

This is because if $\alpha = \{i, j\}$, then $\xi^{(t; \alpha; (\beta_i, \beta_j - 1))}(\cdot)$ has each agent randomizing both at and immediately after t . As we have seen, the expected value for i of the random outcome will be $F_i^{\beta_i - 1, \beta_j - 1}(t)$. On the other hand, if $\alpha = \{i\}$, then $\xi^{(t; \alpha; (\beta_i, \beta_j - 1))}(\cdot)$ has j firing with probability one at t , so that once again, i 's payoff is $F_i^{\beta_i - 1, \beta_j - 1}(t)$.

Proof of Proposition I.

We have established that if (β^*, π^*, c^*) passes our test, then there exists an equilibrium for the multi-stage game that has the properties described in Theorem II. To prove Proposition I, we need to establish that this equilibrium is unique. To establish this, we first show that the payoffs to leading and following in CD^β must agree with the payoffs in the classical duel induced by f^* . We can then prove that there is a unique equilibrium to this phase.

By assumption, if agent i leads at any $t^{\beta_i - 1, \beta_j}$, the next bullet must be fired immediately. Similarly, if agent j fires beyond $\bar{t}^{\beta_i, \beta_j - 1}$, the next bullet must be fired immediately. But by restriction (F5), if i fires at t

in phase β , then he cannot fire at the first instant that the next phase is entered. From Proposition IV, therefore, it *must* be the case that player i fires with probability one at this instant. That is, i 's continuation payoff if he leads at t must be $L_i^{\beta, -1, \beta_j}(t^{\beta, -1, \beta_j})$, for $t < t^\beta$, and $F_i^{\beta, -1, \beta_j}(t)$, for $t \geq t^\beta$. But these are exactly the payoffs defined generated by playing f^* beyond phase β . Similar comparisons can be made for the remaining payoffs. Conclude that beyond t^β , conditions (A1)-(A3) are satisfied. It follows from the arguments that we used to prove Proposition IV (in the text) that if there is an equilibrium in which agents wait until some $t > t^\beta$ before firing, they both must fire with probability one, earning payoffs $(S_i^\beta(t), S_j^\beta(t))$. To complete the proof, we must show that such an equilibrium cannot exist. It is, obviously, sufficient to show that at every t , at least one agent must strictly prefer following to firing simultaneously.

This condition is certainly satisfied in phase (1,1). Now assume that this is true for every phase $\beta' \leq \beta$. From arguments given in the text, if i fires at t , $-i$ will necessarily strictly prefer following at t to firing simultaneously, whenever the next firing time (if i misses) is strictly greater than t . Assume therefore that the next bullet is fired immediately. By assumption, in the equilibrium for the next phase, at least one agent fires with probability less than one, and is weakly prefers following to firing. By inspection of (A.8.a)-(A.8.d), this agent strictly prefers following to firing in phase β .

Oct-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.

- 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.
- 8747 Joseph Farrell
INFORMATION AND THE COASE THEOREM

Jul-87.
- 8748 Philippe Aghion, Patrick Bolton and Bruno Jul
LEARNING THROUGH PRICE EXPERIMENTATION BY A MONOPOLIST FACING UNKNOWN DEMAND

Aug-87.
- 8749 Patrick Bolton
THE PRINCIPLE OF MAXIMUM DETERRENCE REVISITED

Aug-87.

Oct-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.

- 8750 Jeffrey A. Frankel, Kenneth A. Froot
CREDIBILITY, THE OPTIMAL SPEED OF TRADE LIBERALIZATION, REAL INTEREST RATES,
AND THE LATIN AMERICAN DEBT
Aug-87.
- 8751 Barry Eichengreen
DID INTERNATIONAL ECONOMIC FORCES CAUSE THE GREAT DEPRESSION?

Sep-87.
- 8752 Jeffrey A. Frankel
MONETARY AND PORTFOLIO-BALANCE MODELS OF EXCHANGE RATE DETERMINATION

Sep-87.
- 8753 Albert Fishlow
MARKET FORCES OR GROUP INTERESTS: INCONVERTIBLE CURRENCY
IN PRE-1914 LATIN AMERICA
Sep-87.
- 8754 Albert Fishlow
SOME REFLECTIONS ON COMPARATIVE LATIN AMERICAN
ECONOMIC PERFORMANCE AND POLICY
Sep-87.
- 8755 Barry Eichengreen
REAL EXCHANGE RATE BEHAVIOR UNDER ALTERNATIVE
INTERNATIONAL MONETARY REGIMES: INTERWAR EVIDENCE
Sep-87.
- 8756 Leo K. Simon and William R. Zame
DISCONTINUOUS GAMES AND ENDOGENOUS SHARING RULES

Oct-87.
- 8757 Leo K. Simon
A MULTISTAGE DUEL IN CONTINUOUS TIME

Oct-87.