A Theory of Endogenous Commitment*

Guillermo Caruana
CEMFI
caruana@cemfi.es

Liran Einav
Stanford University and NBER
leinav@stanford.edu

June 24, 2005

Abstract

Commitment is typically modeled by assigning to one of the players the ability to take an initial binding action. The weakness of this approach is that the fundamental question of who has the opportunity to commit cannot be addressed, as it is assumed. This paper presents a framework in which commitment power arises endogenously from the fundamentals of the model. We construct a finite dynamic game in which players are given the option to change their minds as often as they wish, but pay a switching cost if they do so. We show that for games with two players and two actions there is a unique subgame perfect equilibrium with a simple structure. This equilibrium is independent of the order and timing of moves and robust to other protocol specifications. Moreover, despite the perfect information nature of the model and the costly switches, strategic delays may arise in equilibrium. The flexibility of the model allows us to apply it to various environments. In particular, we study an entry-deterrence situation. Its equilibrium is intuitive and illustrative of how commitment power is endogenously determined.

KEYWORDS: Commitment, entry, switching costs.

JEL classification: C7, C73.

*This paper is a revised version of a chapter in our 2002 dissertations at Boston University and Harvard University, respectively. We would like to thank Drew Fudenberg and Bob Rosenthal for invaluable advice, guidance, and support. To our deep sorrow, Bob passed away while we were working on this paper. We will miss him very much. We have also benefited from discussions with Estelle Cantillon, Eddie Dekel, Hsueh-Ling Huynh, Shachar Kariv, Asim Khwaja, David Laibson, John Leahy, Bart Lipman, Markus Möbius, Dilip Mookherjee, Al Roth, Balazs Szentes, Steve Tadelis, Elena Zoido, participants at the 2002 Annual Congress of the EEA at Venice, the 2003 Harvard Strategy Conference, and seminar participants at Alicante, Autonoma, UC Berkeley, BU, Carlos III, CEMFI, Duke, Harvard, Illinois, Kellogg, NYU Stern, Pompeu Fabra, Rutgers, Stanford, and Tel-Aviv. The usual disclaimer applies. Correspond to Liran Einav, Department of Economics, Stanford University, Stanford, CA 94305-6072; Tel: (650) 723-3704.

1 Introduction

Ever since Schelling (1960), commitment has been a central and widely used concept in economics. Parties interacting dynamically can often benefit from the opportunity to credibly bind themselves to certain actions, or, alternatively, to remain flexible longer than their opponents. Commitment is typically modeled through dynamic games in which one of the players is given the opportunity to take an initial binding action, allowing him to commit first. This approach has the drawback that the fundamental question of who has the opportunity to commit is driven by a modeling decision. The main goal of this paper is to provide a game-theoretic framework in which the set of commitment possibilities is not imposed, but arises naturally from the fundamentals of the model. Thus, issues such as preemption, bargaining power, credibility, and leadership can be addressed.

Consider Schelling's original example of an army burning its bridges, but imagine that both armies have the opportunity to burn their own retreating routes. Which army is more likely to use this commitment opportunity? How would the answer depend on the importance of the disputed land for each army's chances of winning the war? The framework we develop will provide a way to address these questions.

Another illustration is provided by the large literature on entry decisions of firms. A potential entrant considers entering a market. The incumbent has the opportunity to create a tougher environment for the entrant by using some costly device, e.g. by investing in additional capacity. If the entrant can commit to an action before the incumbent decides, he will enter the market, forcing the incumbent to accommodate. If, alternatively, the incumbent can credibly commit to fight before the entrant makes his final decision, entry can be deterred. A simple way to capture these two stories is to consider a game in which each player makes a decision only once. The order of play gives the opportunity to commit to the player who moves first. But which player should move first? This stylized model cannot answer the question of who has the opportunity to commit earlier, as it is assumed. In contrast, the framework we develop does not rely on an exogenously specified choice of the order or timing of moves.

To offer a more specific and recent example, consider the competition between Boeing and Airbus over the launching of the superjumbo. Both firms had initially committed resources to launching a very large aircraft. Ultimately, Boeing backed off and Airbus prevailed. As convincingly argued by Esty and Ghemawat (2002), since both firms are likely to share similar abilities in taking initial binding actions, the ultimate outcome is likely to be driven by the asymmetric effect of competition at the superjumbo segment on Boeing's and Airbus' profits. Since Boeing's existing jumbo (the 747) is the closest substitute to the future superjumbo, Boeing had a stronger incentive to soften superjumbo competition, and therefore greater incentive not to launch. Esty and Ghemawat (2002) make this argument using a simple two-stage game of entry and exit. The premise of our model is that these stages (and their timing) are not imposed; they will endogenously emerge as the key binding entry and exit decisions out of a much larger set of decision opportunities. One of the main advantages of the framework is its wide applicability; it provides a unified way to think about the role of commitment in a broad range of strategic interactions.

The framework we propose considers a fixed and known date in the future at which a final decision has to be made. Prior to that date, players announce the actions they intend to take in this final date. They can change their announced actions as often as they want. But, for the announcements to be credible rather than cheap talk, we assume that if a player changes his previously announced action he incurs a switching cost. In this manner, the announcements play the role of an imperfect commitment device. We assume that as the final deadline approaches, the cost of switching increases, and that just before the deadline these costs are so high that players are fully committed to their announced actions. Generically, the model has a unique subgame perfect equilibrium (henceforth spe). In Section 3 we show that in games with two players and two actions the spe strategies have a simple structure. They can be described by a finite and small number of stages. Within a stage, a player's decision only depends on the most recent announcements made, but not on the exact point in time within the stage. This implies that, although players could potentially vary their decisions often, in equilibrium they seldom do so. In particular, on the equilibrium path at most one player switches, and when he does so he does it only once. This equilibrium feature points to another (endogenous) commitment strategy. Commitment may be achieved by partially committing to a strategy that is eventually abandoned. Such inefficient delays are often explained by uncertainty or asymmetric information. Our model can rationalize such costly delays even in a world of perfect information.

The main result of the paper (Theorem 1) is that the equilibrium of games with two players and two actions is independent of the order and the timing of the moves. As long as both players can revise their announcements frequently enough, the exact order and timing of their moves have no impact. This accomplishes the main task of laying out a setting in which commitment is not driven by order and timing assumptions. Throughout the paper we assume that players move sequentially in a pre-specified order. This restriction simplifies the proofs but does not drive the results. In Section 5.1 we claim that some asynchronicity, as in Lagunoff and Matsui (1997), is sufficient to preserve the qualitative results of the paper. In Section 5.3 we show that while the results should also generalize to bigger action spaces, the order independence result does not extend for general N-player games. Nevertheless, we suggest interesting families of games for which it does: Caruana, Einav, and Quint (forthcoming) study an N-player bargaining setting using this paper's framework and obtain a unique order-independent outcome; Quint and Einav (2005) analyze an N-player entry game with a similar structure which also preserves order independence.

The three main assumptions of the model – a fixed deadline, increasing switching costs, and payoffs (net of switching costs) that only depend on the final decisions taken – cannot fit all possible scenarios. This framework, however, is flexible enough to accommodate a wide range of interesting economic situations. Section 4 analyzes the example of entry deterrence. Consider any new market which is to be opened at a pre-specified date (e.g. as a result of a patent expiration, the introduction of a new technology, or deregulation). All the potential competitors have to decide whether or not to enter. In order to be ready and operative on the opening day they need to take certain actions (build infrastructure, hire labor, etc.). The increasing cost assumption fits well, as one can assume that the later these actions are taken, the more costly they are.[1] Other economic problems that can be analyzed using this setting are those that involve strategic competition in time. Consider, for example, the decision of when to release a motion picture. Movies' distributors compete for high demand weekends, but do not want to end up releasing all at the same time. This raises the question of what determines the final configuration of release dates.[2] Finally, the model may be also applied to elections and other political conflicts in which a deadline is present.[3]

Rosenthal (1991) and Van Damme and Hurkens (1996) also address the issue of commitment and its timing. In their models, however, once an action is taken it cannot be changed later on, while in our model players can reverse their actions. Indeed, in our framework players eventually get locked into their actions, but this happens only gradually. This reversibility aspect also distinguishes our paper from Saloner (1987), Admati and Perry (1991), and Gale (1995, 2001). They all allow for changes in one's actions, but only in one direction. Allowing for the possibility to reverse one's actions is important. This goes back to Judd's (1985) critique of the product proliferation literature (Schmalensee, 1978; Fudenberg and Tirole, 1985); as he points out, high exit costs and not only high entry costs are crucial in making product proliferation credible. Finally, Henkel (2002) has a similar motivation to ours and some of his results are related (e.g. the potential for strategic delays). In his work, however, the players' roles (leader and follower) are exogenously imposed.

The paper most similar to ours is Lipman and Wang (2000) that, like us, analyzes finite games with switching costs. They use similar techniques and share similarities in the structure of their equilibria. Their objective, however, is to study the consequences of introducing switching costs in a repeated game environment. Thus, their framework involves small switching costs and a flow of payoffs, while we have switching costs that become very high and payoffs that only depend on the final decisions. As a consequence, the intermediate decisions taken by the players do not directly impact the final payoffs in our framework, as they do in theirs. In general, this leads to different predictions for the two models. We discuss this further in Section 5.2.

2 The Model

Consider N players, each with a finite action space. The game starts at $t = 0$. There is a deadline T by which each player will have to take a final action. These final actions determine the final

[1] Exiting may seem to be free. But if the assets bought for entering become more specific to the firm as time goes by, this involves an implicit cost in holding them: their scrap value diminishes. See Section 4 for more details.

[2] See Einav (2003) for an empirical analysis of the release date timing game.

[3] Consider for example the strategic decision by presidential candidates regarding the allocation of their resources in the last weeks of the campaign (Stromberg, 2002).

payoffs for each player. Between $t = 0$ and $t = T$ players will have to decide about their final actions, taking into account that any time they change their decisions they have to pay a switching cost. Formally, a game is described by (Π, C, g), where Π stands for the payoff matrix, C for the switching cost technology, and g for the grid of points at which players get to play. We specify each below.

Time is discrete. Each player i takes decisions at a large but finite set of points in time. We refer to this set as the *grid* of player i and denote it by g_i. Formally $g_i \in G$, where G is the set of all finite sets of points in $[0, T]$. For most of the paper, we assume that players play sequentially, so that $g_i \cap g_j = \emptyset$ for any $i \neq j$. Given a grid $g_i = \{t_1^i, t_2^i, ..., t_{L_i}^i\}$ where $t_l^i < t_m^i$ if $l < m$, we define the *fineness* of the grid as $\varphi(g_i) = Max\{t_1^i, t_2^i - t_1^i, t_3^i - t_2^i, ..., T - t_{L_i}^i\}$. Finally, denote the *game grid* by $g = \{g_i\}_{i=1}^N$, and its fineness by $\varphi(g) = \underset{i}{Max}\{\varphi(g_i)\}$. Throughout the paper, $\varphi(g)$ is considered to be "small" (more precisely specified later). The idea is that players have many opportunities to switch their decisions.[4] This is also a convenient point to introduce some additional notation. Given a point in time t, denote the next and previous points on the grid at which player i gets to play by $next_i(t) = Min\{t' \in g_i | t' > t\}$ and $prev_i(t) = Max\{t' \in g_i | t' < t\}$, respectively. Similarly, let $next(t) = Min\{next_i(t) | i \leq N\}$ and $prev(t) = Max\{prev_i(t) | i \leq N\}$.

When player i gets to play at $t \in g_i$, he has to take an action from his action space A_i. At every point in time all previous actions are common knowledge. The very first move by player i, taken at t_1^i, is costless. However, if he later changes his decision, say from a_i to a_i', he has to pay a switching cost $C_i(a_i \to a_i', t)$.[5] We impose the following assumptions on this function:

1. $C_i(a_i \to a_i', t)$ is a continuous and strictly increasing function in t on $[0, T]$, $\forall a_i, a_i' \in A_i$, $a_i \neq a_i'$, $\forall i \in \{1, ..., N\}$.

2. $C_i(a_i \to a_i, t) = 0 \; \forall a_i \in A_i$, $\forall i \in \{1, ..., N\}$ $\forall t \in [0, T]$.

3. $C_i(a_i \to a_i', 0) = 0 \; \forall a_i, a_i' \in A_i$, $\forall i \in \{1, ..., N\}$.

4. $C_i(a_i \to a_i', T) = \infty \; \forall a_i, a_i' \neq a_i \; \forall i \in \{1, ..., N\}$.

5. $C_i(a_i \to a_i', t) + C_i(a_i' \to a_i'', t) \geq C_i(a_i \to a_i'', t) \; \forall a_i, a_i', a_i'' \in A_i$, $\forall i \in \{1, ..., N\}$, $\forall t \in [0, T]$

Finally, all that remains to be specified are the payoffs for the different players. Let $\overline{a}_i = (a_i(t))_{t \in g_i}$ be player i's sequence of decisions over his grid, let $\overline{a} = (\overline{a}_i)$ be the profile of all sequences, and let $a^* = (a_i(t_{L_i}^i))$ be the final actions of all players. Player i's payoffs are:

$$U_i(\overline{a}) = \Pi_i(a^*) - \sum_{t \in g_i - \{t_1^i\}} C_i(a_i(prev_i(t)) \to a_i(t), t) \tag{1}$$

[4] To gain intuition, the reader could imagine the model in continuous time. Our model is constructed in discrete time to avoid the usual problems of existence of equilibria in continuous models.

[5] Notice that the cost does not depend on the opponents' actions. This simplifies the analysis and allowing it would not qualitatively change any of the main results of the paper.

where $\Pi = (\Pi_i)_{i=1}^N$ is the payoff function for the normal-form game with strategy space $A = \prod_{i=1}^{N} A_i$. Thus, the payoffs for player i are the payoffs he collects at the end, which depend on the final play by all players, less the switching costs he incurred in the process, which depend only on player i's own sequence of actions.

The equilibrium concept that we use is subgame perfect equilibrium (spe). Notice that, by construction, for a generic (Π, C, g) there is a unique spe. This is a finite game of perfect information. Hence, one can solve for the equilibrium by applying backward induction. The only possibility for multiplicity arises when at a specific node a player is indifferent between two or more actions. If this happens, any perturbation of the final payoffs Π or the grid g eliminates the indifference. More precisely, given a cost function C, the set of games that have multiple equilibria has measure zero.[6] For this reason and to simplify the analysis we abstract from these cases. We will discuss the non-generic cases as we proceed with the analysis.

We make three additional remarks. First, note that the switching cost function does not literally need to approach infinity as $t \to T$. All we require is that switching late in the game is more costly than any possible extra benefit achieved in the final payoffs, namely

$$C_i(a_i \to a_i', T) > \underset{a_{-i}}{Max} \left(\Pi_i(a_i', a_{-i}) - \Pi_i(a_i, a_{-i})\right) \quad \forall i, a_i, a_i' \tag{2}$$

It is easy to see that in equilibrium no player will ever switch after

$$\bar{t} = \underset{i,a_i,a_i',a_{-i}}{Max} \; C_i^{-1}(a_i \to a_i', \Pi_i(a_i', a_{-i}) - \Pi_i(a_i, a_{-i})) \tag{3}$$

Second, switching costs are sunk. This and the absence of indifference nodes make history irrelevant. If a player has to take an action at t and the last decisions taken by all players are a, when or how often he or other players had changed their minds before this point has no impact on their future payoffs. Thus, we can define the relevant state space by $\{(a,t) \mid a \in A, t \in g\}$ and denote the spe strategy for player i by $s_i(a,t) \in A_i$ $\forall a \in A$, $\forall t \in g_i$.[7]

The third remark regards the way we model the cost technology. In our setting the cost technology is a primitive. It is given exogenously and cannot be changed by the players. Nevertheless, one can think of situations in which commitment is achieved by changing one's switching costs. This possibility can be handled within the model. One just needs to expand the action space to permit players to change their cost function. For example, if a player has an action space A_i and can choose either high or low switching costs (H or L), we just need to consider a new action space, $\{H, L\} \times A_i$. Accordingly, the switching cost function would be higher if the switch is done under the H regime and lower under L, and switching between regimes would be costly.

[6] In the paper we will use the following measures: (i) for the space of g_i's: $\mu(B) = \sum_{n=1}^{\infty} \mu_n(B \cap G_n)$, where G_n is the set of all grids on $[0,T]$ that contain exactly n elements and μ_n is the Lebesgue measure on $[0,T]^n$; (ii) for the space of g's the product of the g_i's measures; (iii) for the space of Π's the usual Lebesgue measure on $\mathbb{R}^{N \cdot K^N}$; and (iv) for the space of (Π, g)'s the product measure of the two.

[7] If at t a player has not played yet, clearly the state does not depend on his action space.

3 Analysis of Two-By-Two Games

3.1 An Example

To illustrate the typical structure of the equilibrium, consider the following entry deterrence game:

	Entry	*No Entry*
Fight	$2, -10$	$10, 0$
No Fight	$5, 3$	$12, 0$

Assume that the switching costs are equal across different actions and for both players; that is, $C_i(a_i \rightarrow a_i', t) = c(t)\ \forall a_i \in A_i\ \forall i \in \{1, 2\}$. For simplicity, let $c(t) = t$.[8] Suppose that T is big ($T > 10$) and that players alternate and decide every 0.01 increment in the following way: the entrant plays at $\varepsilon \in (0, 0.01)$, the incumbent at $0.01 + \varepsilon$, the entrant at $0.02 + \varepsilon$, and so on.

In Section 4 we will analyze in detail a generalized version of this game. We defer to that point the economic interpretation of the equilibrium. Our focus now is on the equilibrium structure. The game is solved using backward induction. Equation (3) specifies the latest date at which a player would switch. This happens at $9.98 + \varepsilon$, the first node before $\bar{t} = c^{-1}(10) = 10$ at which the entrant plays. At this point, if the action profile is [*Fight, Entry*] (that is, if the most recently announced actions were *Fight* by the incumbent and *Entry* by the entrant), he would switch and not enter the market. Consider now a decision node in the interval $(5, 10]$ at action profile [*Fight, Entry*]. For the incumbent it is still too costly to make a change. If it is the entrant's turn, he will play *No Entry* immediately to save on switching costs. For any profile different from [*Fight, Entry*] it is still too costly to consider any change of actions.

Next, consider the profile [*No Fight, Entry*] at $t = 4.99 + \varepsilon$, the last node before $t = 5$ at which the incumbent plays. If he plays *No Fight* now, he will keep on playing it until the end and get a final payoff of 5. By switching to *Fight*, however, the entrant would react by not entering, guaranteeing the incumbent a final payoff of 10. Given that the switching cost is less than 5, the incumbent finds it profitable to switch to *Fight*. We can now move one step backwards and analyze the entrant's decision at [*No Fight, Entry*] at $t = 4.98 + \varepsilon$. He anticipates that if he plays *Entry*, the incumbent will respond by fighting, which will force the entrant out of the market. Thus, the entrant prefers to play *No Entry* immediately in order to save on switching costs. From this point backwards, the entrant always plays *No Entry*. As a consequence, the players' initial decisions are [*No Fight, No Entry*] and on the equilibrium path the players do not switch. Notice that this outcome is not an equilibrium of any of the one-shot sequential games. The table below presents the complete equilibrium strategies.

[8] Strictly speaking, this specification violates the requirement that $c(t) \rightarrow \infty$ as $t \rightarrow T$. As already mentioned, we only need $c(T)$ to be big enough (greater than 10 in this particular case).

Time	**Switches**
Initial actions	NE by the entrant; NF by the incumbent
$[0.02+\varepsilon, 1.99+\varepsilon]$	$[F,E] \to [NF,E]$ and $[F,NE] \to [NF,NE]$ by the incumbent $[F,E] \to [F,NE]$ and $[NF,E] \to [NF,NE]$ by the entrant
$[2+\varepsilon, 4.98+\varepsilon]$	$[F,E] \to [F,NE]$ and $[NF,E] \to [NF,NE]$ by the entrant
$4.99+\varepsilon$	$[NF,E] \to [F,E]$ by the incumbent
$[5+\varepsilon, 9.98+\varepsilon]$	$[F,E] \to [F,NE]$ by the entrant
$[9.99+\varepsilon, T]$	None

The second column indicates the profiles at which a player decides to change his previous action. If a profile is not on the list it is because the player's action is to continue playing the same action as before.

Notice that despite the fact that players have many opportunities to play, the structure of the equilibrium is quite simple: there are long periods of time in which players' incentives remain constant, and there are only a few instances at which they change. In the analysis we refer to these instances as the *critical points* of the game and to the intervals in which the strategies remain constant as *stages*. As we will see, this structure is common to any game (Π, C, g).

3.2 Structure of the Equilibrium Strategies

This section shows how commitment is achieved in this model. Given that the switching costs are low early in the game and only increase as the game advances, real commitment to an action is only attained at some point in the game. This endogenously creates a "commitment ladder," such that over time each player is able to commit better to certain actions. Each step in the ladder corresponds to a stage and each new critical point introduces a new commitment possibility.

For any two-player game (Π, C, g) and the corresponding spe strategies $s_i(a,t)$ for both players, we formally introduce the above mentioned concepts.

Definition 1 *For any i, $t^* \in g_i$ is a **critical point** if there exists an action profile $a = (a_i, a_j)$ such that $s_i((a_i,a_j),t^*) = a'_i$ $(a'_i \neq a_i)$ and $s_i((a_i,a_j), next_i(t^*)) = a_i$.*

Definition 2 *Let $\{t_1^*, t_2^*, ..., t_k^*\}$ be the set of critical points, such that $t_i^* < t_j^*$ if $i < j$. The corresponding k+1 **stages** are the following intervals:* $[0, t_1^*], (t_1^*, t_2^*], (t_2^*, t_3^*], ..., (t_{k-1}^*, t_k^*], (t_k^*, T]$.

Each critical point t^* is associated with a specific action profile a and a specific player i. Player i's response at profile a is changed just after t^*. This happens for one of two reasons. First, it can be due to a pure time phenomenon. It is the last point at which it is still profitable for player i to switch away from a. After this point, such a switch would be too costly, so the player may be thought of as committed to this action. Second, it can be a consequence of a change by the opponent: player i anticipates that immediately afterwards player j will do something new, which in turn changes player i's incentives. In the example of Section 3.1, $t = 9.98 + \varepsilon$ is a critical point of the first type and $t = 4.98 + \varepsilon$ of the second.

Given the definition of a stage, we first establish that the strategies for both players are held constant throughout a stage. This fact is not directly implied by Definition 2.

Proposition 1 *$\forall i \in \{1, 2\}$ $\forall t, t' \in g_i$, if t, t' are in the same stage then $s_i(a, t) = s_i(a, t')$ $\forall a \in A$.*

The proof, as well as all other proofs, is relegated to the appendix. Note that an important consequence of Proposition 1 is that on the equilibrium path of any subgame, switches occur only at the beginning of a stage. The last stage of a game is given by equation (3), which provides a point $\bar{t}$ after which no player switches. Note, however, that before this point players could in principle build up very complicated strategies. As a result, there could potentially be as many stages in the game as points in the grid between 0 and $\bar{t}$. The next proposition shows that this is not the case. The equilibrium has a simple structure and the number of stages is quite limited.

Proposition 2 *Given a cost structure C, generically for every (Π, g), the unique spe of (Π, C, g) is completely characterized by $\overline{m} \leq 7$ critical points $\{t^*_m\}_{m=1}^{\overline{m}}$ and the corresponding stage strategies.*

The number of stages for any game is at most eight. The number eight is of no particular interest, but reflects the fact that the complexity of the equilibrium is limited. The proposition's statement is generic because the argument assumes no indifference at all $t \in g$. The proof uses an algorithm (see Appendix B) that computes the stages and the corresponding strategies. The algorithm finds the equilibrium strategies without the need to apply backward induction at every decision node; it computes continuation values only after each critical point.

An alternative approach to describe the spe strategies is to use the notion of strategic delays. Given that switching is more costly as time goes by, one could think that whenever there is a profitable switch, it would be carried out earlier rather than later in order to save on costs. Nevertheless, we show that delays may occur in equilibrium for strategic reasons.

Definition 3 *Consider a decision node (a, t) for $t \in g_i$ at which player i switches, i.e. $s_i(a, t) = a'_i$ ($a'_i \neq a_i$). This switch is a **delayed switch** if there exist $\widetilde{a}$ and $t' < t$ such that $t' \in g_i$ and (a, t) is on the equilibrium path of the subgame $(\widetilde{a}, t')$.*

Note that a delayed switch may never materialize. It is defined with respect to a subgame, which may be on or off the equilibrium path of the game. The next proposition argues that on the equilibrium path of any subgame (a, t), there can be at most one delayed switch.

Proposition 3 *Given a cost structure C, generically for every (Π, g), the unique spe strategies of (Π, C, g) are such that the equilibrium path of any subgame contains at most one delayed switch.*

In the proof we proceed in two steps. First, we show that for a switch to be delayed it has to be credible. If player i delays a switch, and then reverses this switch later on, then player j will ignore the original delay, making it wasteful – it could have been done earlier at a lower cost. Second, we show that for a switch to be delayed, it has to be beneficial, in the sense that it has to make player j do something different than what he would have done without the delay. For two-action games, this means that a player delays a move until the point at which the other player is committed to an action. Hence, for a delayed switch to be credible and beneficial it must be the last switch on the path. A delayed switch by the row player can be viewed as credible (irreversible) if it eliminates a row of the payoff matrix from further consideration, and as beneficial if it eliminates a column of the payoff matrix from further consideration. For two-by-two games, these eliminations result in a unique outcome, so there are no further switches. Finally, we apply Proposition 3 to the equilibrium path of the full game to obtain:

Corollary 1 *On the equilibrium path, one of two patterns are observed: (a) both players play immediately the final profile and never switch thereafter; or (b) one player immediately plays the final action and the other starts by playing one action and switches to the other later on.*[9]

3.3 Grid Invariance

We want to compare the equilibria of a given game for different grids. Clearly, the exact position of the critical points depends on the grid chosen. We show, however, that as long as the grid is fine enough, the number of stages and the corresponding strategies are invariant to the grid. This allows us to define a notion of equilibrium for a given (Π, C) without making any reference to the specific grid. To do so formally, we define a notion of equivalence between two equilibria.

Definition 4 *Consider two games* (Π, C, g) *and* (Π, C, g'). *The unique spe equilibria of both games are* ***essentially the same*** *if the number of stages in both coincide and the strategies at each stage are the same.*

It is according to this definition of equivalence that we state the grid-invariance property:

Theorem 1 *Given* C*, generically for every* Π *there exists* $\alpha > 0$ *such that for almost every* $g \in G$, $\varphi(g) < \alpha$ *the spe equilibria of* (Π, C, g) *are essentially the same.*

This result is obtained by making extensive use of the limit version of the model, that is, taking the fineness of the grid to zero. Generically, the limit of the equilibria exists. This implies that the order of the stages in the limit is also the order of the stages of the finite game, as long as the grid for that game is fine enough. In other words, in the limit the critical points converge. Therefore, as long as the critical points for different players are separated in the limit, for fine grids players get the opportunity to play and react at all the relevant points in the game.

[9] In Section 3.5 we provide an example of a game with a delayed switch on the equilibrium path.

The maximal fineness of the grid allowed α depends on how far apart the critical points for the two players are. It also depends on the slope of the cost function at the points of delayed switches, if those exist. Within a stage, a player switches at most once. Thus, all he needs is one opportunity to play at the beginning of each stage. Including more points on the grid does not change his strategic opportunities.

Let us stress the importance of the quantifiers used in the theorem. We state the result *for almost every* grid in order to avoid the multiplicity of equilibria, and *generically for every* Π to guarantee the existence of the limit of the equilibria. The limit may not exist for two reasons. First, we want to rule out those games that have multiple equilibria for any grid. This happens if one of the players is initially indifferent between two different actions. A slight perturbation of the payoff matrix would eliminate such multiplicity. Second, the theorem also rules out an additional measure-zero case, which arises when in the limit the two players have a critical point at the same time. Suppose this common critical point is t^*. Then, for any given grid g, no matter how fine it is, the equilibrium may depend on whether $prev_i(t^*) < prev_j(t^*)$ or the reverse. Therefore, the limit of the equilibria may not exist. A slight perturbation of the payoffs of one of the players separates the critical points for both players, making the problem disappear. For example, a fully symmetric Battle of the Sexes exhibits grid dependence, but any perturbation provides a unique outcome.

3.4 Cost Invariance

The shape of the equilibrium of a game (Π, C) depends on the choice of the cost technology. If one wanted to empirically use this model, information about the cost technology, C, is unlikely to be available. As a result, the model may be unidentified: given Π, different choices of C may give rise to different equilibrium outcomes. Driven by this motivation, we suggest restrictions on the relationship among the cost functions across different players and moves. Under these restrictions the equilibrium is invariant to the exact shape of the cost structure.

The model is independent of the "nominal" units of time used. All that matters is the value of the cost function at each decision node. Consequently, rescaling time has no impact on the spe strategies. Formally, let $C(t)$ be a cost technology, g the grid, and $f(\cdot)$ any strictly increasing function. Then, the games (Π, C, g) and $(\Pi, C(f(t)), f(g))$ have the same equilibrium strategies, outcome, and values (and, consequently, so do the grid invariant games (Π, C) and $(\Pi, C(f(t)))$).

Consider now a cost technology of the form $C_i(a_i \to a_i', t) = \theta_i^{a_i \to a_i'} c(t)\ \forall a_i \neq a_i'$, where $c(t)$ is, as usual, continuous and strictly increasing in t, with $c(0) = 0$ and $c(T) = \infty$. By setting $f(\cdot) = c^{-1}(\cdot)$ and rescaling time, any such game can be thought of as a game with proportional linear costs (across players *and* actions), namely $C_i(a_i \to a_i', t) = \theta_i^{a_i \to a_i'} t$. Thus, the equilibrium only depends on (Π, Θ), where Θ stands for the full matrix of $\theta_i^{a_i \to a_i'}$'s, but not on $c(t)$.

There are two special cases of the former result that are worth mentioning. First, if $C_i(a_i \to a_i', t) = c(t)\ \forall a_i \neq a_i'$ then the equilibrium of (Π, C, g) only depends on Π. In many situations it may be natural to assume that the switching cost technology is the same for all players and across all possible moves. In such cases the model provides an essentially unique equilibrium which only depends on Π. Second, if $C_i(a_i \to a_i', t) = \theta_i c(t)\ \forall a_i \neq a_i'$ then the equilibrium only depends on $\Pi^* = \left\{\frac{\Pi_i}{\theta_i}\right\}_{i=1}^{N}$.[10]

3.5 Additional Results

A natural question at this point is whether there are any easy conditions on the primitives (Π, C) that determine the shape of the outcome. The short answer is no. Even though the equilibrium structure is simple, the combination of incentives along the eight possible stages is sufficient to provide a rich variety of possible dynamic interactions.[11] One can establish some simple results, such as the fact that players' equilibrium payoffs are bounded from below by their maxmin payoffs of the one-shot game. But in order to obtain sharper equilibrium predictions one has to restrict attention to specific families of games. In this manner, one can show that pure coordination games always result in the Pareto efficient outcome, or use the notion of defendability introduced by Lipman and Wang (2000) to provide sufficient conditions for a Nash Equilibrium (of the one-shot game) to be the outcome of the dynamic game. The study of the entry deterrence situation in Section 4 is another example of this approach.

Finally, we provide an illustration of the strategic use of delayed moves. Denote action a_i as **super-dominant** for player i if $\underset{a_j}{Min}\ \Pi_i(a_i, a_j) > \underset{a_j}{Max}\ \Pi_i(a_i', a_j)\ \ \forall a_i' \neq a_i$. By the maxmin argument, it is clear that if player i has a super-dominant action, this has to be his final action. This may lead us to think that in equilibrium player j best-responds to player i's super-dominant action. If such a response leads to the best outcome for player i then this is indeed true. However, when there is a conflict, and player j's best-response works against player i's incentives, player i may be able to "discipline" player j and force him to choose the other action. The following game (with symmetric switching cost functions, $c(t)$) illustrates this case:

	L	R
U	$13, 3$	$1, 10$
D	$0, 5$	$0, 0$

Although U is super-dominant for player 1, in equilibrium he starts by playing D. He then switches to his super-dominant strategy U if player j "behaves" and plays L, and only after $t = c^{-1}(7)$, when player j is fully committed to his "disciplined" behavior. This is credible since if player 2 played R, switching to U would not justify the costs. It is also profitable: player 1's payoffs, $13 - c(c^{-1}(7)) = 6$, are higher than 1, which he could obtain by playing U throughout.

[10] To see this, note that multiplying a player's switching costs and payoffs by a common factor has no effect on the game, except for a normalization of this player's final payoffs (see equation (1)).

[11] By focusing only on the direction of the incentives for each player at each stage, one can obtain a full taxonomy of all the possible two-by-two games. In this way, all such games can be classified into 75 types.

4 An Application: Entry Deterrence

We now proceed to analyze the standard entry deterrence problem presented in the introduction in which fixed deadlines arise naturally. The expiration of a patent on a certain drug, the introduction of a new hardware technology, or the scheduling decision for the release of a new product are only some examples. The values in the matrix Π capture the resulting payoffs of the ex-post competition. As for the increasing switching costs, imagine that the incumbent (entrant) fights (enters) by investing in, say, physical capital. At any point before the opening of the market, he can contract the delivery of this investment for the opening day, when it is actually needed. Delaying these contracts naturally increases costs. One reason for this can be that the machinery suppliers may be aware of the deadline and may charge the incumbent (entrant) more for it. Alternatively, the supply of these factors may decrease over time because they get committed to other tasks.

Imagine now a firm that has previously contracted the physical asset and now wants to get rid of it. At first glance one may think that this involves no cost, as the firm can sell the asset for its market value. But implicitly there is a cost. As in the previous example, it may be reasonable to assume that the scrap value of the asset diminishes over time. Thus, by leaving the market late a firm loses the money it could have earned had it left earlier. The increasing switching costs capture the diminishing value of the physical capital in the outside market. As the deadline gets closer the wedge between the value of the capital within the firm, if it fights (or enters), and its value elsewhere gradually increases. Additionally, one can interpret these costs in contracting terms. If a contract is nullified, there is a penalty involved in it, which increases as the deadline approaches. Our assumptions also require that writing new contracts or nullifying existing ones becomes sufficiently costly when the opening of the market is close enough.

The entry game that we consider has the following general payoff matrix:

	Entry	No Entry
Fight	$d, -a$	$m, 0$
No Fight	D, b	$M, 0$

where all the parameters are positive and satisfy $M > m > D > d$ and $M > D + b$. For simplicity only, we also assume that the switching costs, $c(t)$, are equal for both parties and across different actions. As described in the introduction, the one-shot sequential games exogenously give all the commitment power to one of the parties. If the entrant plays first the spe outcome is [*No Fight, Entry*], but if the incumbent is able to commit first then the equilibrium is [*Fight, No Entry*].

The following proposition describes all possible equilibrium outcomes. It shows that four possible cases may arise in equilibrium: three outcomes played immediately, and one more which involves a strategic delay. The four cases create a partition of the parameter space.

Proposition 4 *The spe outcome of the entry game is:*

(i) [No Fight, Entry] with no switches $\Leftrightarrow D - d > a$.

(ii) *[Fight, No Entry] with no switches* $\Leftrightarrow D - d < a$, $b > M - m$, *and* $b > Min\{a, m - D\}$.

(iii) *Start with [Fight, No Entry] and switch (by the incumbent) to [No Fight, No Entry] at* $t^* = c^{-1}(b) \Leftrightarrow D - d < a$ *and either* $a < b < M - m$ *or both* $(M - D)/2 < b < Min\{M - m, a\}$ *and* $Min\{a, m - D\} < M - m$.

(iv) *[No Fight, No Entry] with no switches* $\Leftrightarrow D - d < a$ *and either* $b < Min\{M - m, (M - D)/2, a\}$ *or* $Max\{M - m, b\} < Min\{a, m - D\}$.

The proof is a simple application of the limiting version of the algorithm to all the relevant cases. Below we provide some economic intuition. Given that neither of the players wants to stay at [*Fight, Entry*] till time T, and that both would rather have the opponent switch, there is an off-equilibrium war of attrition taking place at this profile. Each player prefers to wait and let the other player move away from it. The party that wins the war of attrition is the first one that can credibly tie himself to that position. Given that we have assumed the same switching cost technology for both parties, the winner is the player with smaller benefits of making the move ($D - d$ for the incumbent and a for the entrant). The other party foresees this and moves away immediately. Thus, when $a < D - d$ the incumbent is forced to accommodate, resulting in [*No Fight, Entry*], the best outcome for the entrant. This is case (i) of the proposition.

If $D - d < a$, the war of attrition is won by the incumbent. The threat, in equilibrium, is sufficient to keep the potential entrant out of the market. However, while the incumbent is happy deterring entry, he can do so at different costs. He could fight till T, but, if possible, he would prefer to either deter entry by not fighting at all, or by switching to *No Fight* later in the game. These different levels of commitment correspond to cases (ii), (iii), and (iv) of the proposition. The intuition for which case arises can be illustrated by examining the parameter b and its impact on the incumbent's strategies at profile [*Fight, No Entry*].

If b is high we are in (ii). In this case, as long as it is still profitable for the incumbent to quit fighting, it is also profitable for the entrant to react by entering. Thus, the only way entry can be deterred is through fighting. In case (iii) the incumbent achieves [*No Fight, No Entry*], but only after paying the cost of strategically delaying the switch to *No Fight* till $c^{-1}(b)$, the point after which the entrant is committed to staying out. This happens when b has an intermediate value. It has to be low enough so that late in the game it is still profitable for the incumbent to switch and stop fighting; but it has to be high enough so that earlier in the game, if the incumbent decided to switch to *No Fight*, the entrant would enter, knowing that the incumbent cannot restart fighting.

Finally, in case (iv) the commitment power of the incumbent is the highest. He can deter entry without ever fighting. This is achieved by maintaining a credible threat to react by fighting whenever the entrant decides to enter. For this threat to be successful, the entrant needs to lack the credibility to enter the market and stay. In other words, as long as the entrant finds it profitable to enter, he still finds it profitable to switch to no entry if he were subsequently fought. This is guaranteed by $b < a$. On top of that, the incumbent must be able to credibly commit to respond by fighting to any entry attempt by the entrant. This occurs either because m is big enough (which can be thought of as a case in which it is quite cheap for the incumbent to fight),[12] or because M is very big (which implies that after deterring entry by fighting, the incumbent still finds it profitable to pay the extra cost to get rid of the additional capacity).

We think that the final two cases of the proposition are sensible and appealing outcomes, which rationalize how an incumbent can deter entry without actually fighting. Similar results were obtained in Milgrom and Roberts (1982) and Kreps and Wilson (1982). Our solution, however, does not rely on the introduction of asymmetric information, as the previous papers do.

5 Discussion

5.1 Robustness to Other Protocols

The results so far are invariant to the choice of the grid. As we highlighted before, this is because players are given the opportunity to react at all the relevant points in the game. In this section we argue that this rationale is robust to other protocol specifications. In particular, we claim that for any sufficient amount of asynchronicity in the timing of actions between players, the qualitative results of the paper would not change.

Maintaining the same switching cost structure, payoffs, and finite grid as before, consider a random device that determines at each node whether or not a player gets to play. This results in periods of both simultaneous and sequential moves.[13] While the analysis of this game is technically more complicated, we believe that for fine grids there is an essentially unique equilibrium which can be described through a stage structure. This equilibrium is the same as the one we obtain with the sequential structure. This stage structure has two features that are new compared to the pure sequential game: (i) At periods at which both players play simultaneously the spe strategies may be mixed. This is because a player may be tempted to save on switching costs by moving at the same time as the other instead of waiting one period for the other to make the move first.[14] As the grid becomes finer, however, this temptation diminishes. Thus, in the limit the mixing probabilities converge to pure strategies. (ii) Late in the stages, the probability that one player may not play again within the stage becomes higher. Once the probability is high enough, the incentives of the players may change. As the fineness of the grid goes to zero, however, the probability to move during any given time interval goes to one. Therefore, in the limit, the equilibrium converges to the one predicted by our sequential-move structure.

The results are different if one considers the game in which players decide simultaneously at every point. As before, one can describe the equilibria with a stage structure. For the same reason as above, these equilibria involve mixed-strategies. In addition, multiple equilibria arise. One of

[12] Note that the example presented in Section 3.1 is covered by this case.

[13] For example, consider an i.i.d. (over time and across players) process such that, at each period, each player plays with probabilty $p < 1$.

[14] For example, consider the game of Section 3.1 played simultaneously at each node. At $t = 4.99 + \varepsilon$ and profile [*No Fight, Entry*] both players mix their actions.

these equilibria corresponds to the one predicted by the sequential (or asynchronous) structure. Moreover, the other equilibria are, in a sense, not robust; they can be eliminated by adding a small amount of asynchronicity. The intuition for these extra equilibria is similar to the one underlying the existence of a Pareto-dominated equilibrium in the repeated pure coordination game. As long as a player expects the other player to choose the "bad" action, it is in his best interest to also do so. In this respect, the introduction of asynchronicity has the same effect as in Lagunoff and Matsui (1997). It breaks down this "cursed" string of beliefs. As long as the grid is fine enough, the probability of a player playing alone on a short period of time is almost certain. This allows the players to "coordinate" and guarantees uniqueness.

5.2 Relationship to Lipman and Wang (2000)

As we already pointed out, this paper is closely related to Lipman and Wang (2000) (henceforth: LW). They analyze the robustness of the results for finitely repeated games to the introduction of small switching costs. Their modeling strategy is, therefore, driven by the repeated game literature. Our purpose is different. We analyze commitment in situations which are essentially played only once. We use switching costs as the commitment device. This dictates a different set of considerations. Among others, we emphasize the need of a framework that exhibits grid-invariance. We envision the "order of moves" as a modelling assumption which tries to capture a more amorphous set of rules. Thereby, to capture the sources of commitment, one needs the equilibrium predictions to be robust to changes of this sort. This property, for instance, is not a concern (and, as we point out below, does not apply) in the LW framework.

Still, some of our specific results for coordination games or the Prisoners' Dilemma resemble theirs. This is not coincidental: the constant payoffs and increasing switching costs in our setting have a similar flavor to the decreasing future payoffs and constant switching costs in LW. Loosely speaking, in LW one compares the switching costs ϵ_i to the future payoffs, namely $(T-t)\pi_i$. One may be tempted to think that this is equivalent to having constant payoffs π_i and increasing switching costs of $c_i(t) = \frac{\epsilon_i}{T-t}$, which would satisfy the assumptions of our paper. This argument is, in general, wrong. Whenever there are delayed switches (on or off the equilibrium path), a short-run vs. long-run trade-off appears in LW, but is not present here. The reason for this is that in our setting a player only cares about his own actions and his opponent's *final* action. In contrast, LW use flow payoffs and therefore players care about the whole sequence of their opponent's actions. This results in different equilibrium outcomes for the two models.

To illustrate, consider the example in Section 3.5. In equilibrium player 1 starts playing his dominated strategy, and switches to his dominant strategy only later. This is not an spe in the LW setup. With flow payoffs, player 1 would lose too much from "waiting" at his dominated strategy. He would rather play his dominant strategy throughout and obtain payoffs of at least 1 for the whole duration of the game. Indeed, $[U,R]$ played throughout the game is the spe in LW.

The difference in the payoff structure has consequences in terms of the grid invariance result as well. LW's model is sensitive to the choice of the exact points in time at which players get to play (in particular, to whether the decision nodes are set equidistantly from each other). Consider for instance the games studied in Theorem 6 and 7 of LW. These very similar games result in different equilibria because of minor differences in (very) short-term incentives that player 1 faces. Changes in the exact timing of play (keeping the simultaneous-move assumption) will change these short-term incentives, and ultimately (as consecutive distances become less equal) affect the equilibrium prediction. This argument is true even in the limit, when the grid is very fine.

A more technical difference between this paper and LW is the timing of actions. LW use a simultaneous move game, while we use a sequential one. As we discuss in the previous section, a sequential structure eliminates the need to deal with multiplicity of equilibria and mixed strategies.

5.3 General Action Spaces and N players

While the model was constructed for N players and arbitrary action spaces, the analysis has focused only on two-by-two games. This was done for several reasons. First, these games illustrate the richness of the dynamic framework considered. Second, the proofs are more tractable. Finally, the grid invariance result fails for generic games of more than two players.

Consider first games of two players but larger action spaces. If these are finite, we believe that the game has a unique grid-invariant equilibrium, which can be described by a stage structure. The reason for this is the same as before: as long as the critical points for the two players do not coincide (and this happens generically), for fine grids the players will have the opportunity to carry out all their relevant decisions. To prove it, one could use the same techniques as for the two-by-two case, namely to construct an algorithm and then check its convergence properties. Unfortunately, with an arbitrary number of actions we have not been able to rule out the possibility that the algorithm never ends, namely that the number of stages is infinite.[15]

While the stage structure obviously does not extend to continuous action spaces, one can use our framework for specific games and obtain interesting equilibrium outcomes. Consider first a symmetric two-player homogeneous product price (Bertrand) competition. Applying our framework with identical switching costs $c(t)$ across actions and players results in a unique symmetric equilibrium in which both parties are able to coordinate on the monopolistic price. Interestingly, if one considers any asymmetry (across players) in the switching cost technologies, the equilibrium reverts to perfect competition.

The second example is a bargaining game.[16] Each of the two players has to decide how much

[15] We should note that the algorithm succesfully computed the equilibrium for thousands of simulated random games. A related interesting question regards the maximal number of stages in the equilibrium of a game with K actions. It is easy to see that this number has to be at least K^2. Due to the possibility of delayed switches, the actual number is much greater than this lower bound. For $K = 2$ there are eight stages; simulations for $K = 3$ results in about 30 stages, and for $K = 5$ the number is already above 150.

[16] See Caruana, Einav, and Quint (forthcoming) for a full analysis of a multi-player version of this game, as well as for proofs of the results.

of a pie he wants. If the final demands add up to no more than the size of the total pie, each player gets his demand. If the sum is higher, however, no one gets anything. Specifically, let $A_i = [0, 1]$ and let the payoffs be a_i if $a_i + a_j \leq 1$ and 0 otherwise. With identical switching costs across actions the final agreement is achieved immediately, avoiding any switching costs. With symmetric players, each obtains an equal share of the pie. For asymmetric cases, the higher a player's cost function the earlier he is able to commit not to decrease his demand, so the higher the stake he obtains.[17] We can interpret the cost function as a measure of players' bargaining power. Somewhat surprisingly though, the outcome depends only on the values of the cost functions at a particular point, t^*, but not on their overall shape.

Consider now games with more than two players. Given a particular grid the equilibrium still exhibits a stage structure.[18] But while the exact point in time at which a player plays is not important, it turns out that the identity of the player who plays next may be crucial. To gain intuition, consider a three player game in which players play sequentially in a pre-specified order that repeats itself. Imagine a stage at which, confronted with profile a, both player 1 and player 2 want to switch their actions immediately. Now, imagine that we are at profile (a'_3, a_{-3}) with $a'_3 \neq a_3$. It is player 3's turn to move and he is considering switching his action to a_3. He likes the consequences of player 1's switch from a, but not those of 2's. Here the order of play is key. If player 3 has the opportunity to move right before player 1, he will move to a_3. If player 3 gets to play only before player 2, however, he will prefer not to switch to a_3. This change in player 3's incentives can have drastic consequences on the overall shape of the equilibrium. Notice that this grid dependence property persists no matter how fine the grid is.

Despite this negative result, we should stress that there are interesting families of N-player games which are robust to changes in the grid. We believe that price competition, quantity competition, public good games, and bargaining are among those. In Caruana, Einav, and Quint (forthcoming) we extend the analysis of the bargaining game for more than two players. Obtaining a unique equilibrium in this setting is interesting as it is well known that bargaining models with more than two players are very sensitive to the choice of the protocol, and are often not precise in their predictions. Quint and Einav (2005) analyze a simple N-player entry game with asymmetric costs and a similar dynamic structure and obtain the socially efficient outcome as its unique (and grid-invariant) equilibrium. We also conjecture that games with convex and compact action spaces and continuous and concave payoff functions should exhibit grid invariance.

6 Conclusions

Commitment is typically modeled by giving one of the players the opportunity to take an initial binding action. Although this approach has proven to be useful, it cannot address questions

[17] Muthoo (1996) presents a two-player two-period bargaining model in which he obtains a similar result.

[18] For $N > 2$ one has to modify the stage definition to accommodate the following pattern. There are stages in which, at a given profile, the strategies, instead of being constant, follow a cyclical pattern.

such as how this commitment is achieved, or which party is able to enjoy it. We consider a dynamic model in which players announce their intended final actions and incur switching costs if they change their minds. Given that changing their previous announcements has costs, these announcements are not simply cheap talk. Thus, the switching costs serve as a mechanism by which announcements can be made credible and commitment achieved.

Players are allowed to play very often. Despite this, the equilibrium can be described by a small number of stages, with the property that within each stage players' strategies remain constant. This stage structure does not change when the order of the moves is altered or more decision nodes are added; it is also robust to various changes in the protocol as long as some asynchronicity exists. In this sense, the relevant order by which parties get to commit is endogenously determined.

Moreover, our analysis suggests that the notion that commitment is achieved "once and for all" is too simplistic. Early on players are completely flexible. Late in the game they are fully committed. In between, however, commitment depends on the actual positions of the players. This is why we describe our equilibrium as a "commitment ladder," according to which players are able to bind themselves to certain actions only gradually. This allows for a richer range of possible dynamic stories. The entry deterrence case provides a good example. On top of the two outcomes that arise when one applies the simple one-shot sequential analysis, the model provides a rationale for entry deterrence with no actual fight. This is achieved by a credible threat to fight in retaliation to entry. In this manner, our framework provides an umbrella that covers dynamic interactions that were previously captured only with different models.

The model has several additional desirable features. First, if one assumes that switching costs are identical across players, the equilibrium is invariant to the specific choice of the cost structure. Second, if one thinks that players have some control over their switching cost technology, this can be incorporated by simply increasing the players' action spaces. Third, the framework is flexible enough to accommodate many different strategic situations. We have studied entry and bargaining, and suggested elections, political conflicts, and competition in time as other potential applications. Fourth, we believe that the model may be attractive for empirical work. The uniqueness of equilibrium is important in the empirical analysis of discrete games, in which relying on first order conditions is impossible. On top of this, the algorithm we provide significantly reduces the computational burden of estimating the model.

Finally, let us mention two potential directions for future research. First, we think that the protocol invariance property is attractive, so it may be interesting to search for other frameworks which satisfy it. For example, one could explore games in which players can build up stocks of strategy-specific investments. Consider, for instance, a firm announcing its intention to enter a new market by acquiring some industry-specific human capital. If it later decides not to enter, this human capital cannot simply be erased, as it is implied by the current model. Second, we think that the notion of strategic delays deserves more attention. In a world of imperfect information delaying an action has an option value. In our (perfect information) model delays still occur in equilibrium, but for pure strategic reasons.[19] Delays are costly, but allow players to make threats credible. It would be interesting to introduce incomplete information into the framework and analyze the interaction of these two types of incentives to delay.[20]

References

[1] Admati, Anat, and Motty Perry (1991), "Joint Projects without Commitment," *The Review of Economic Studies*, 58, 259-276.

[2] Caruana, Guillermo, Liran Einav, and Daniel Quint (forthcoming), "Multilateral Bargaining with Concession Costs," *Journal of Economic Theory*, forthcoming.

[3] Einav, Liran (2003), "Not All Rivals Look Alike: Estimating an Equilibrium Model of The Release Date Timing Game," mimeo, Stanford University.

[4] Esty, Benjamin C., and Pankaj Ghemawat (2002), "Airbus vs. Boeing in Superjumbos: A Case of Failed Preemption," Harvard Business School Working Paper Series, No. 02-061.

[5] Fudenberg, Drew, and Jean Tirole (1985), "Preemption and Rent Equalization in the Adoption of New Technology," *The Review of Economic Studies*, 52, 383-402.

[6] Gale, Douglas (1995), "Dynamic Coordination Games," *Economic Theory*, 5, 1-18.

[7] Gale, Douglas (2001), "Monotone Games with Positive Spillovers," *Games and Economic Behavior*, 37, 295-320.

[8] Henkel, Joachim (2002), "The 1.5th Mover Advantage," *Rand Journal of Economics*, 33(1), 156-170.

[9] Judd, Kenneth L. (1985), "Credible Spatial Preemption," *Rand Journal of Economics*, 16, 153-166.

[10] Kreps, David M., and Robert Wilson (1982), "Reputation and Imperfect Information," *Journal of Economic Theory*, 27, 253-279.

[11] Lagunoff, Roger, and Akihiko Matsui (1997), "Asynchronous Choice in Repeated Coordination Games," *Econometrica,* 65, 1467-1477.

[12] Lipman, Barton L., and Ruqu Wang (2000), "Switching Costs in Frequently Repeated Games," *Journal of Economic Theory,* 93(2), 149-190.

[13] Maggi, Giovanni (1997), "Endogenous Leadership in a New Market," *Rand Journal of Economics*, 27(4), 641-659.

[14] Milgrom, Paul, and John Roberts (1982), "Predation, Reputation and Entry Deterrence," *Journal of Economic Theory*, 27, 280–312.

[15] Muthoo, Abhinay (1996), "A Bargaining Model Based on the Commitment Tactic," *Journal of Economic Theory*, 69, 134-152.

[19] Henkel (2002) and Gale (1995) obtain a similar result, although in the latter this is driven by a coordination motive.

[20] Maggi (1996) extends Saloner (1987) to obtain an interesting interaction between commitment and uncertainty.

[16] Quint, Daniel, and Liran Einav (2005), "Efficient Entry," *Economics Letters*, 88(2), 278-283.

[17] Rosenthal, Robert W. (1991), "A Note on Robustness of Equilibria with Respect to Commitment Opportunities," *Games and Economic Behavior*, 3(2), 237-243.

[18] Saloner, Garth (1987), "Cournot Duopoly with Two Production Periods," *Journal of Economic Theory*, 42, 183-187.

[19] Schelling, Thomas C. (1960), *The Strategy of Conflict*, Cambridge, MA: Harvard University press.

[20] Schmalensee, Richard (1978), "Entry Deterrence in the Ready-to-Eat Breakfast Cereal Industry," *Bell Journal of Economics*, 9, 305-327.

[21] Stromberg, David (2002), "Optimal Campaigning in Presidential Elections: The Probability of Being Florida," Seminar Paper No. 706, IIES, Stockholm University.

[22] Van Damme, Eric, and Sjaak Hurkens (1996), "Commitment Robust of Equilibria and Endogenous Timing," *Games and Economic Behavior*, 15(2), 290-311.

Appendix A: Proofs

Before we tackle the proofs of the propositions, we first establish a definition and two useful lemmas that are used throughout the rest of the appendix.

Definition 5 *For a given player i, $t \in g_i$, and action profile a, if $s_i(a,t) = a'_i \neq a_i$ then player i has an **active switch** at (a,t).*

Lemma 1 *If $s_i(a,t) = a'_i \neq a_i$ then $s_i((a'_i, a_{-i}), t) = a'_i$.*

Proof. Let $V_i(a,t)$ be the continuation value of player i at decision node (a,t). By $s_i(a,t) = a'_i$ we know that $V_i((a'_i, a_j), next(t)) - C_i(a_i \to a'_i, t) \geq V_i((a_i, a_j), next(t))$. Since costs are non-negative, this trivially implies $V_i((a'_i, a_j), next(t)) \geq V_i((a_i, a_j), next(t)) - C_i(a'_i \to a_i, t)$ proving the the lemma. ■

Lemma 2 *If there exists a player i and a point in time $t \in g_i$ such that $s_i((a_i, a_j), t) = s_i((a'_i, a_j), t)$ $\forall a_i, a'_i \in A_i$, $a_j \in A_j$, then for both players and for any $t' \in g_j$ such that $t' \leq t$, the strategies are independent of a_i, i.e. $s_j((a_j, a_i), t') = s_j((a_j, a'_i), t')$ and $s_i((a_j, a_i), t') = s_i((a_j, a'_i), t')$ $\forall a_i, a'_i \in A_i$, $\forall a_j \in A_j$. Moreover, there are at most three stages in the interval $[0, t]$.*

Proof. We prove this by induction on the level of the game tree, starting at t and going backwards. By assumption, the statement holds at t, i.e. that $s_i((a_i, a_j), t) = s_i((a'_i, a_j), t)$ $\forall a_i, a'_i \in A_i$. Now, suppose it holds for time $t' \leq t$. We have to show that it holds for $t'' = prev(t')$. Let player j be the player who plays at time t'', i.e. $t'' \in g_j$. We check two cases: first, when $j \neq i$, and second, when $j = i$.

If $j \neq i$ all we have to check is that player j's continuation values just after his move at t'', $V_j(a', t')$, are independent of a'_i. By the induction assumption from time t' until time t, no player's strategy depends on player i's action. Thus, the actions of both players evolve independently of it and $a^t_j(a^{t'}) = a^t_j(a^{t'}_j)$. Moreover, player i plays at t and therefore $a^t_i(a^{t'}) = a^t_i(a^{t'}_j)$. This implies that player j's continuation values satisfy $V_j(a_i, a'_j, t') = V_j(a'_i, a'_j, t')$ $\forall a_i, a'_i \in A_i$.

If $j = i$, the induction assumption implies that $a^t(a^{t'}) = a^t(a^{t'}_{-i})$ $\forall a^{t'} \in A$. Thus, player i knows that, independently of his action at t'', he will end up at subgame $(a^t_i(a^{t''}_j), a^t_j(a^{t''}_j), t)$. Therefore, in order to save on switching costs, he would best move to his time t strategy with respect to the play of the other player, i.e. $s_i(a, t'') = s_i(a^t_j(a^{t''}_j), t)$, which is independent of a_i. This concludes the first part of the lemma.

We prove now the second part of the lemma. Denote by t^* the last point in time at which the hypothesis of the lemma is satisfied. Note that t^* is a critical point. Given that player j's continuation values at t^*

depend only on a_j, we denote them by $V_j(a_j, t^*)$. Player j can obtain the outcome that is more favorable for him just by playing $\widetilde{a}_j = \underset{a_j}{ArgMax} V_j(a_j, t^*)$ at his first grid point and not switching ever until t^*. At any profile a with $a_j = \widetilde{a}'_j$, player j switches immediately to $\widetilde{a}_j$ if and only if $V_j(\widetilde{a}_j, t^*) > V_j(\widetilde{a}'_j, t^*) - C_j(\widetilde{a}'_j \to \widetilde{a}_j, t)$. Clearly, early in the game such a switch is profitable, but as we approach t^* it may not be. Denote by $t^{**} \in g_j$ the last point before t^* at which this switch would be made. Player i's strategy at each time before t^* mimics his strategy at t^*, with respect to player j's anticipated action. To summarize, prior to t^{**} both players play $(s_i(\widetilde{a}_j, t^*), \widetilde{a}_j)$ at any profile, and between t^{**} and t^* the strategies of both players at profile a are $(s_i(a_j, t^*), a_j)$. Thus, we have at most three stages and three critical points: $prev_i(t^{**})$, t^{**}, and t^*. The critical point at $prev_i(t^{**})$ does not always exist. It appears only when player i needs to re-adjust to the expected move at t^{**} by player j. This happens when $s_i(\widetilde{a}_j,t^*) \neq s_i(\widetilde{a}'_j,t^*)$. ■

Proposition 1 *$\forall i \in \{1, 2\}$ $\forall t, t' \in g_i$, if t,t' are in the same stage then $s_i(a, t) = s_i(a, t')$ $\forall a \in A$.*

Proof. We prove by contradiction that there are no two consecutive decision nodes for player i, $t, next_i(t) \in g_i$, within a stage satisfying $s_i((a_i, a_j), t) = a_i$ and $s_i((a_i, a_j), next_i(t)) = a'_i$ for a given $a \in A$. If player j does not move between t and $next_i(t)$ the contradiction is immediate. Consider the case in which he does. W.l.o.g. suppose that he plays only once in between and does it at t'. We consider different cases depending on what player j does at profiles $((a_i, a_j), t')$ and $((a'_i, a_j), t')$:

1. If $s_j((a_i, a_j), t') = a_j$ and $s_j((a'_i, a_j), t') = a_j$ player i can deviate from the proposed equilibrium and increase his profits by playing a'_i at (a, t), which leads to a contradiction.

2. If $s_j((a_i, a_j), t') = a_j$ and $s_j((a'_i, a_j), t') = a'_j$, and given that t' is not the end of a stage, we know that $s_j((a'_i, a_j), next_j(t')) = a'_j$. This implies that the equilibrium path starting at (a, t) leads to $((a'_i, a'_j), next_i(next_j(t')))$. But player i can get there at a lower cost by deviating and playing a'_i at (a, t) and not switching until $next_j(t')$. This provides the contradiction.

3. If $s_j((a_i, a_j), t') = a'_j$ and $s_j((a'_i, a_j), t') = a'_j$, and given that t' is not the end of a stage, we know that $s_j((a_i, a_j), next_j(t')) = s_j((a'_i, a_j), next_j(t')) = a'_j$. Using Lemma 2 we know that $s_i((a_i, a_j), next_i(t)) = s_i((a_i, a'_j), next_i(t)) = a'_i$. Now it is easy to see that player i can improve by deviating at (a, t) and playing a'_i. Again, this leads to a contradiction.

4. Finally, if $s_j((a_i, a_j), t') = a'_j$ and $s_j((a'_i, a_j), t') = a_j$ and given that t' is not the end of a stage, we know that $s_j((a_i, a_j), next_j(t')) = a'_j$. Consider what player i does at $((a_i, a'_j), next_i(t))$. If $s_i((a_i, a'_j), next_i(t)) = a_i$, one can check that player i can benefit from playing a'_i at (a, t), providing a contradiction. If $s_i((a_i, a'_j), next_i(t)) = a'_i = s_i((a_i, a_j), next_i(t))$, by Lemma 2 we have that $s_j((a_i, a_j), t') = s_j((a'_i, a_j), t')$, which is a contradiction. ■

Proposition 2 *Given a cost structure C, generically for every (Π, g), the unique spe of (Π, C, g) is completely characterized by $\overline{m} \leq 7$ critical points $\{t^*_m\}_{m=1}^{\overline{m}}$ and the corresponding stage strategies.*

Proof. The proof makes use of the algorithm (Appendix B). The proof applies only generically to avoid those cases in which the algorithm aborts. This happens when a player is indifferent about what to play at a node. Given (Π, C, g), the algorithm provides the following output $(t^*_m, S_g(i, a, m), V_m, AM_m)_{m=0}^{\overline{m}}$. All we have to show is that the algorithm replicates the spe. More precisely, that

$$\widetilde{s}_i(a, t) = S_g(i, a, \widetilde{m}(t)) \text{ where } \widetilde{m}(t) = \{m | t \in (t^*_{m+1}, t^*_m]\}$$

are indeed the spe strategies. We will also see that the following definition of $\widetilde{V}_p(a, t)$ coincides with the continuation values of the game for player i at node (a, t)

$$\widetilde{V}_i(a, t) \equiv \begin{cases} V^{new}(V_{\widetilde{m}(t)-1}, AM_{\widetilde{m}(t)-1}, t, i) \text{ evaluated at } (a, i) \text{ if } t \in g_i \\ V^{new}(V_{\widetilde{m}(t)-1}, AM_{\widetilde{m}(t)-1}, t, j) \text{ evaluated at } (a, i) \text{ if } t \notin g_i \end{cases} \tag{4}$$

where $V^{new}(V, AM, t, i)$ is defined in Appendix B, part 4.

We prove the proposition by induction on the level of the game tree, starting at T and going backwards. The induction base is straight forward: as time approaches T the costs go to infinity. Therefore, provided that the grid is fine enough, the cost of switching at the final decision node is too high. The algorithm initializes with $AM_0(a, i) = 0$ for all a, i. Thus, $\widetilde{s}_i(a, T) = a_i$ and $\widetilde{V}_i(a, T) = \Pi_i(a)$ which coincide with the equilibrium strategies and continuation values.

Suppose now that the statement is true for $next(t)$. We will show that it is true for t as well. Fix a profile a. As before, once we have proven that the proposed strategy $\widetilde{s}_i(a, t)$ is indeed optimal, verifying the update of the continuation values is immediate. Because of the induction hypothesis we know that $\widetilde{V}_i(a, next(t))$ are the continuation values of the game. Therefore the spe strategy is the solution to

$$s_i(a, t) = \underset{a_i' \in A_i}{ArgMax}\{\widetilde{V}_i((a_i', a_j), next(t)) - C_i(a_i \to a_i', t)\}$$

Proving that $\widetilde{s}_i(a, t) = s_i(a, t)$ is equivalent to proving that

$$AM_{\widetilde{m}(t)}(a, i) = 1 \Longleftrightarrow \widetilde{V}_i((a_i', a_j), next(t)) - C_i(a_i \to a_i', t) > \widetilde{V}_i(a, next(t)) \tag{5}$$

The advantage of using equation (5) is that it only involves functions defined in the algorithm. Therefore, the problem is reduced to an algebraic check. This is simple but tedious, as it involves many different cases. First, $\widetilde{V}$ is defined piecewise and recursively, thus it can have eight different expressions depending on the values of AM and FS. Second, the statement deals with $\widetilde{m}(t)$ and $\widetilde{m}(next(t))$, which may take the same or different values. Potentially, thirty two cases have to be checked. Many of the cases can be ruled out as impossible or easily grouped and checked together. Including a full check for all the cases in the Appendix would be too long, and would not provide much intuition. Still, we present one case to show how easy each check is. Consider a point $t \in g_i$ in the middle of a stage. Suppose that only player i has an active move (at profile a) on this stage. These conditions translate into $\widetilde{m}(t) = \widetilde{m}(next(t))$ and all the $AM_{\widetilde{m}(t)}$'s are equal to zero except for $AM_{\widetilde{m}(t)}(a, i) = 1$. In this case, applying equation (4), we have that

$$\widetilde{V}_i(a, next(t)) = V_{\widetilde{m}(next(t))}((a_i', a_j), i) - C_i(a_i \to a_i', next_i(t))$$
$$\widetilde{V}_i((a_i', a_j), next(t)) = V_{\widetilde{m}(next(t))}((a_i', a_j), i)$$

Now one can easily check that equation (5) is satisfied. Moreover, once we know that the algorithm solves for the unique spe of the game, then, as a direct application of Remark 1 (in the end of Appendix B), we get that the equilibrium has no more than eight stages. ∎

Proposition 3 *Given a cost structure C, generically for every (Π, g), the unique spe strategies of (Π, C, g) are such that the equilibrium path of any subgame contains at most one delayed switch.*

Proof. Consider a subgame $(\widetilde{a}, t_0)$ with $t_0 \in g_i$ with a delayed switch by player i at $t > t_0$. First, we show that on the equilibrium path of this subgame player j will never switch after t. Suppose towards contradiction that player j switches at $t_1 > t$. W.l.o.g. assume that the last delayed switch by player i before j's first switch is at (a, t). Thus, player i switches from a_i to a_i' , after which player j at $((a_i', a_j), t_1)$ switches to a_j'. By Lemma 1, player j plays a_j' at $((a_i', a_j'), t_1)$ as well. This means that at $(\widetilde{a}, t_0)$ player i has a profitable deviation: by always playing a_i' he obtains the same outcome with lower switching costs.

Next, we show that player i will not switch after t either. We prove it by contradiction. Without loss of generality, assume that $t = next_i(t_0)$ and that the last two delayed switches by player i are at (a, t) from a_i to a_i' and at $((a_j, a_i'), t_1)$ from a_i' to a_i. Note that we are making use of the first part of the proposition, which guarantees that player j does not switch after t. Denote the possible continuation values for player j at t_1 by $A \equiv V_j((a_i, a_j'), t_1), C \equiv V_j((a_i, a_j), t_1) = V_j((a_i', a_j), t_1)$. Observe that the delayed switch of player i at t_1 implies that player j switches from a_j to a_j' at $((a_i, a_j), prev_j(t_1))$, implying $A - C_j(a_j \to a_j', prev_j(t_1)) > C$.

Now, player j must play a_j' at $((a_i', \widetilde{a}_j), next_j(t_0))$, otherwise there would not have been any reason for player i to delay the switch at $((a_i, a_j), t_0)$. Thus, $V_j((a_i', a_j'), next_j(t_0)) - C_j(\widetilde{a}_j \to a_j', next_j(t_0)) > C - C_j(\widetilde{a}_j \to a_j, next_j(t_0))$. Observe also that for $t_0 < t' \leq t_1$ at $((a_i', a_j'), t')$ player j always sticks to a_j', otherwise player i could play a_i' at $((a_i, a_j), t_0)$ instead of delaying. Denote by $\bar{t}$ the first time, if any, that player i plays a'_i at $((a_i, a'_j), \bar{t})$ for $t_0 < \bar{t} \leq t_1$. If $\bar{t}$ does not exist, the following is a profitable deviation for player j: play a'_j at $((a_i, \tilde{a}_j), next_j(t_0))$ and stick to a'_j at any $t_0 < t' \leq t_1$. This strategy would yield payoffs of $A - C_j(\tilde{a}_j \to a'_j, next_j(t_0))$, which are greater than $C - C_j(\tilde{a}_j \to a_j, next_j(t_0))$ (player j's value from playing a_j at $next_j(t_0)$), and hence provides a contradiction. If $\bar{t}$ exists then the following is a profitable deviation for player j: play a'_j at $((a_i, \tilde{a}_j), next_j(t_0))$ and after that mimic the spe strategy at every node. It is easy to check that this results in payoffs of at least $V_j((a'_i, a'_j), next_j(t_0)) - C_j(\tilde{a}_j \to a'_j, next_j(t_0))$, which are greater than $C - C_j(\tilde{a}_j \to a_j, next_j(t_0))$, as shown before. Thus, leading to a contradiction. The reason for this is that, given the switch by player i at $\bar{t}$, the only case in which $V_j((a'_i, a'_j), next_j(t_0)) \neq V_j((a_i, a'_j), next_j(t_0))$ is if player j switches to a_j at $((a_i, a'_j), t')$ for $next_j(t_0) < t' < \bar{t}$. But if this happens, by revealed preferences we know that $V_j((a'_i, a'_j), next_j(t_0)) < V_j((a_i, a'_j), next_j(t_0))$. ■

Theorem 1 *Given C, generically for every Π there exists $\alpha > 0$ such that for almost every $g \in G$, $\varphi(g) < \alpha$ the spe equilibria of (Π, C, g) are essentially the same.*

Proof. It is sufficient to show that generically for every (Π, C) the limit of the equilibria of the finite games, taking $\varphi(g) \to 0$, exists and is independent of the order of moves. Precisely we will prove that $\lim_{\varphi(g)\to 0} S_g(i, a, m) = S(i, a, m)$ where $S_g(i, a, m)$ and $S(i, a, m)$ are defined in Appendix B.

First, note that the statement of the theorem is generic to avoid the cases for which the limiting version of the algorithm aborts. This rules out the cases in which the critical points are the same for both players.

We prove the statement above recursively on the stages of the algorithm. For a given m we check the convergence of the functions used in the algorithm $(t, a^*, p^*, t^*_m, AM_m, V_m, FS_m)$. This task has to be done in the same order in which the algorithm proceeds. It is sufficient to realize that each function is piecewise defined by continuous transformations of (i) other functions for which the convergence has already been checked (because of the recursive procedure); or (ii) the cost function, which is continuous. Finally, the cutoff points in the piecewise functions also converge. This is so because the mutually exclusive conditions that define the cutoff points are (except for the case of $t(a, i)$) functions with a finite range (and for which the recursive procedure applies). For the case of $t(a, i)$ the cutoff is determined by $\Delta V = 0$, at which there is no discontinuity. This essentially finishes the proof of the theorem. The existence of α is an immediate consequence of the fact that the range of $S_g(i, a, m)$ is finite. ■

Proposition 4 *The spe outcome of the entry game is:*

(i) [No Fight, Entry] with no switches $\Leftrightarrow D - d > a$.

(ii) [Fight, No Entry] with no switches $\Leftrightarrow D - d < a$, $b > M - m$, *and* $b > Min\{a, m - D\}$.

(iii) Start with [Fight, No Entry] and switch (by the incumbent) to [No Fight, No Entry] at $t^* = c^{-1}(b)$ $\Leftrightarrow$ $D - d < a$ *and either* $a < b < M - m$ *or both* $(M - D)/2 < b < Min\{M - m, a\}$ *and* $Min\{a, m - D\} < M - m$.

(iv) [No Fight, No Entry] with no switches $\Leftrightarrow D - d < a$ *and either* $b < Min\{M - m, (M - D)/2, a\}$ *or* $Max\{M - m, b\} < Min\{a, m - D\}$.

Proof. If $D - d > a$, the equilibrium outcome is [*No Fight, Entry*]. To see this, consider the following strategy by the entrant: start by entering and do not exit after that. Clearly, the best response by the incumbent to such a strategy is not to fight, so that final payoffs would be (D, b). We need to show that this strategy is subgame perfect for the entrant. Consider the decisions made at profile [*Fight, Entry*]. If $t > c^{-1}(D - d) > c^{-1}(a)$ none of the players would consider switching, as the the switching costs are already too high. When t is between $c^{-1}(a)$ and $c^{-1}(D - d)$ the entrant is already committed not to leave, but the incumbent finds it still profitable to stop fighting. Finally, before $c^{-1}(a)$ the entrant knows that by sticking to entry he will eventually achieve his maximum payoffs, b, and hence it is credible for him to do so.

Now we consider the case of $D - d < a$, and mechanically go over all possible cases. The reader can verify that these cases match the different restrictions stated in the proposition. Starting at T, and going backwards, note first that the first action that becomes active, i.e. the first action for which one of the players finds it beneficial to pay the switching costs, is related to which of the three following numbers is higher: a, b, or $M - m$. We analyze each case in turn:

1. If b is the maximum, the action "Enter if the incumbent is not fighting" is the first move to become active (at $t = c^{-1}(b)$). Thus, in the last stage of the game (between $c^{-1}(b)$ and T) no player moves, and in the previous stage this is the only action active. The next differences to be considered are a and $m - D$.

   (a) If $a > m - D$, in the next stage the entrant activates his other action, "exit if the incumbent fights". In the termination stage the incumbent ignores the entrant's actions (recall Proposition 2) and best replies to the entrant's strategy at $t = c^{-1}(a)$. Given that $m > D$, the best reply is to fight, thus leading to an equilibrium in which [*Fight, No Entry*] is played throughout the game.

   (b) If $m - D > a$, in the next stage the incumbent starts fighting if the entrant is out, knowing that if he does not do so, the entrant will enter. This does not change the incentives for the entrant, who still wants to enter whenever the incumbent is not fighting. This creates an off-equilibrium war of attrition at the profile [*Fight, Entry*], which, given that $a > D - d$, is won by the incumbent. Therefore the equilibrium is [*Fight, No Entry*].

2. If $M - m$ is the maximum, at $t = c^{-1}(M - m)$ the action "do not fight if the entrant is out" is the first one to become active. The next differences to be considered are a and b.

   (a) If $a > b$, at $c^{-1}(a)$ the action "exit if the incumbent fights" becomes active. The next switch depends on the comparison between b and $(M - D)/2$. Consider what happens at [*No Fight, No Entry*]. If the entrant enters expecting the incumbent not to react he will gain b. If the incumbent decided to fight the entry, that would prompt a chain reaction, in which the entrant would react by exiting, after which the incumbent would stop fighting. The incumbent, therefore, understands that in order to obtain M, he has to make two consecutive and almost immediate (recall that $\varphi(g) \to 0$) switches. Initiating this sequence of moves is beneficial as long as $2c(t) < M - D$, or that $t < c^{-1}((M - D)/2)$. Thus:

      i. If $(M - D)/2 > b$ then the equilibrium is [*No Fight, No Entry*]. The entry deterrence is achieved through a credible threat to fight in case the entrant decided to enter.

      ii. If $b > (M-D)/2$ then right before $t = c^{-1}(b)$ the entrant would enter if the incumbent were not fighting. In anticipation, the incumbent has two options. He can either accommodate entry and obtain D, or use a delayed switch strategy: start by fighting, make the entrant exit, and wait until $t = c^{-1}(b)$ to stop fighting, knowing that after $t = c^{-1}(b)$ the entrant will not enter any longer. In this case, the incumbent payoffs would be $M - c(c^{-1}(b)) = M - b$. Given that we assumed that $M - b > D$, this second option is preferred by the incumbent. Thus, on equilibrium players start with [*Fight, No Entry*] and at $t = c^{-1}(b)$ the incumbent switches to [*No Fight, No Entry*].

   (b) If $b > a$, at $c^{-1}(b)$ the action "enter if the incumbent is not fighting" becomes active. In response, the incumbent's switch to *No Fight* becomes inactive right before this. He can still achieve M by fighting until $t = c^{-1}(b)$, and switching to *No Fight* only after that. This will give him payoffs of $M - c(c^{-1}(b)) = M - b$. As we have assumed that $M - b > D$, this strategy is better than simply accommodating entry. Therefore, the equilibrium involves initial actions of [*Fight, No Entry*] and a later switch to [*No Fight, No Entry*] at $t = c^{-1}(b)$.

3. If a is the maximum, at $t = c^{-1}(a)$ the action "exit if the incumbent is fighting" becomes the first active move. We now need to compare between $M - m$, $m - D$, and b.

   (a) If b is the maximum, this case is analogous to case 1.a above. Given that the entrant is (unconditionally) more flexible than the incumbent, the incumbent prefers to fight. Thus, on equilibrium [*Fight, No Entry*] is played throughout the game.

   (b) If $m - D$ is the maximum, the equilibrium is [*No Fight, No Entry*]. This is a case in which the incumbent can easily commit to fight entry: it is not very costly for him (as m is large), but it hurts the entrant a lot (as a is large).

(c) If $M-m$ is the maximum, this case is analogous to 2.a above. If $(M-D)/2 > b$ the incumbent can credibly threat to fight entry and the equilibrium outcome is [*No Fight, No Entry*]. If, however, $b > (M-D)/2$ the equilibrium involves playing [*Fight, No Entry*] until $t = c^{-1}(b)$, in which the the incumbent switches to [*No Fight, No Entry*]. ■

Appendix B: Algorithm

Here we describe the algorithm, which is essential for the proof of Theorem 1. In the proof we also refer to the limiting version of the algorithm, that is, as the fineness of the grid $\varphi(g)$ goes to zero. Since the switching cost technology is continuous, the limiting version is identical to the finite version of the algorithm, with the only changes affecting parts 2 and 4, in which $next_i(t)$ and $prev_i(t)$ are replaced by t. A Matlab code for the limiting version of the algorithm is available at *http://www.stanford.edu/~leinav.*

In the end of this appendix we prove that the algorithm terminates in a small and finite number of steps, for any grid. Finally, in what follows, if p is one player we use $\sim p$ to denote the other player. Given a particular game (Π, C, g) the algorithm steps are described below.

Initialization: Set $m = 0$ (stage counter, starting from the end); $t_0^* = T$ (the last critical time encountered); $V_0(a,p) = \Pi$ (continuation value of player p at profile a just after t_m^*); $AM_0(a,p) = 0$ (an indicator function; it equals one iff there is an active switch at time t_m^* by player p from profile a); and $IM = \{(a,p) | a \in A,\ p = 1,2\}$ (the set of inactive moves).

Update (m, V_m, AM_m)**:**

1. $m = m + 1$
2. Find the next critical time t_m^*, and the action a^* and player p^* associated with it. This is done by comparing the potential benefits and costs for each move. We use some auxiliary definitions:

   (a) Let $SM_{m-1}(a,p)$ be an ordered set of action profiles $(a^0, a^1, ..., a^{k-1}, a^k)$ such that $a^0 = a$ and

   $$a^{i+1} = \begin{cases} (a_p^i, a_{\sim p}^{i\prime}) & \text{if } AM_{m-1}(a^i, \sim p) = 1 \\ (a_p^{i\prime}, a_{\sim p}^i) & \text{if } AM_{m-1}(a^i, p) = 1 \text{ and } AM_{m-1}(a^i, \sim p) = 0 \end{cases}$$

   for $i \geq 0$. This defines the sequence of consecutive switches within stage $m-1$ that start at a and ends at a profile from which there is no active move. We denote this final node by $\overline{SM}_{m-1}(a,p)$. The sequence is finite, contains up to three switches, and is solely a function of AM_{m-1}.

   (b) Given $SM_{m-1}(a,p) = (a^0, ..., a^k)$, define $FS_{m-1}(a,p) = \sum_{i=1}^{k} I(a_p^{i-1} \neq a_p^i)$ where $I(\cdot)$ is the indicator function (FS_{m-1} computes the number of switches by player p in the $SM_{m-1}(a,p)$ sequence).

   (c) Let $\Delta V_{m-1}(a,p) \equiv V_{m-1}(\overline{SM}_{m-1}((a_p', a_{\sim p}), p)) - V_{m-1}(\overline{SM}_{m-1}(a,p))$. This difference in values stands for the potential benefits of each move at profile a by player p.

   Now, compute the critical time associated with each move. This involves four different cases, as shown below. The first is when the move gives negative value. The second is a case in which if player p does not move, he will be moving at his next turn (because the other player will move to a profile in which player p prefers to move). This means that player p prefers to move right away, rather than delaying his move, so the critical time kicks in immediately before the next critical time. The third case is the "standard" case, in which the critical time is the last time at which the cost of switching is less than its benefit. The last case is similar, but takes into account that the move involves an extra immediate switch at the next period.

$$t_m(a,p)^{21} = \begin{cases} 0 \\ \quad \text{if } \Delta V_{m-1}(a,p) < 0 \\ prev_p(t^*_{m-1}) \\ \quad \text{if } \Delta V_{m-1}(a,p) \geq 0 \text{ and } FS_{m-1}(a,p) > 0 \\ Max\left\{t \in g_p, t < t^*_{m-1} | C_p(a_p \to a'_p, t) \leq \Delta V_{m-1}(a,p)\right\} \\ \quad \text{if } \Delta V_{m-1}(a,p) \geq 0 \text{ and } FS_{m-1}(a,p) = 0 \text{ and } FS_{m-1}((a'_p, a_{\sim p}), p) = 0 \\ Max\left\{t \in g_p, t < t^*_{m-1} | C_p(a_p \to a'_p, t) + C_p(a'_p \to a_p, next_p(t)) \leq \Delta V_{m-1}(a,p)\right\} \\ \quad \text{if } \Delta V_{m-1}(a,p) \geq 0 \text{ and } FS_{m-1}(a,p) = 0 \text{ and } FS_{m-1}((a'_p, a_{\sim p}), p) > 0 \end{cases}$$

The next critical time is the one associated with the move that maximizes the above, out of the moves that are not active yet.

$(a^*, p^*) = \underset{(a,p) \in IM}{ArgMax}\{t_m(a,p)\}$

Abort if $|p^*| > 1$.[22] Equal critical times for different players (*the solution is not grid invariant*).
If not, set $t^*_m = t_m(a^*, p^*)$
Abort if $t^*_m = 0$ (*a player is indifferent between two actions at* $t = 0$)
If not, set $p^*_m = p^*$

3. Update the set of active moves. First, activate the move associated with the new critical time. Second, deactivate moves by the other player that originate from the same action profile, but only if $m = 2$ or if we are in the early part of the game. The third case involves a move whose destination is the origin of the new active move. Such a move is deleted and reevaluated in the next iteration. Finally, the rest of the moves remain as they were before.

$$AM_m(a,p) = \begin{cases} 1 & \text{if } (a,p) \in (a^*, p^*) \\ 0 & \text{if } (a,p) \in (a^*, \sim p^*) \text{ and } (m = 2 \text{ or } AM_{m-1}(a',p) = 1) \\ 0 & \text{if } (a,p) \in ((a^*_{p^*}, \sim a^*_{\sim p^*}), \sim p^*) \\ AM_{m-1}(a,p) & \text{otherwise} \end{cases}$$

4. Compute the continuation values of the players just after t^*_m. This is done by using the value at the terminal node of an active sequence of consecutive moves (as defined in part 2), and subtracting the switching costs incurred by the player along this sequence. These switching costs are incurred just after t^*_m. First, define the following mapping

$$V^{new}(V^{old}, AM, t, \overline{p})(a,p) = V^{new}(V^{old}, SM(AM), t, \overline{p})(a,p) = V^{old}(\overline{SM}(a,\overline{p})) - CC(SM(a,\overline{p}), t, a, p)$$

where CC is recursively defined as follows:

$$CC(SM(a,\overline{p}), t, a, p) = \begin{cases} 0 & \text{if } \overline{SM}(a,\overline{p}) = (a) \\ CC(SM(a^1, \sim p), next_{\sim p}(t), a^1, p) & \text{if } a_p = a^1_p \text{ where } (a^0, ..., a^k) = SM(a,\overline{p}) \\ CC(SM(a^1, p), next_p(t), a^1, p) + \\ \quad + C_p(a_p \to a^1_p, next_p(t)) & \text{if } a_p \neq a^1_p \text{ where } (a^0, ..., a^k) = SM(a,\overline{p}) \end{cases}$$

Now, compute the continuation values by $V_m = V^{new}(V_{m-1}, AM_{m-1}, t^*_m, p^*_m)$.

5. Let $IM = \left\{(a,p) | AM_m(a,p) = 0 \text{ and } AM_m((a'_p, a_{\sim p}), p) = 0\right\}$.

6. **Terminate** if $\#IM = 0$ (all moves are active), and let $\overline{m} = m$, $t^*_{\overline{m}+1} = 0$. Otherwise, go to part 1.

[21] Note that by having weak inequalities we implicitly assume that a player switches whenever he is indifferent between switching or not.

[22] $ArgMax$ is a correspondence. This is why we use '$\in$' rather than equalities in part 3 of the algorithm. Given the way we construct $t_m(a,p)$, the multiple solutions must be associated with a unique p^* for any finite grid. In the limiting case, this is the only generically case. This is why the algorithm may abort in non-generic cases.

Output: The essential information of the algorithm consists of the number of stages of the game, $\overline{m}$, the critical points that define the end of each stage, $(t^*_m)_{m=0}^{\overline{m}}$, and the strategies at every stage

$$S_g(p,a,m) = \begin{cases} a_p & \text{if} \quad AM_m(a,p) = 0 \\ a'_p & \text{if} \quad AM_m(a,p) = 1 \end{cases}$$

Nevertheless, for practical reasons we define the output of the algorithm to be

$$(t^*_m, S_g(p,a,m), V_m, AM_m)_{m=0}^{\overline{m}}$$

In the limiting case, we use the notation $S(p,a,m)$ instead of $S_g(p,a,m)$.

Lemma 3 *For any* (Π, C, g)*, the algorithm ends in a finite number of stages, and in particular* $\overline{m} \leq 8$.

Proof. The algorithm finishes when $\#IM = 0$. Observe that:

1. If $AM_m(a,p) = 1$ then $AM_m((a'_p, a_{\sim p}), p) = 0$ and vice versa, thus $\#IM = 0$ implies that $\#AM = 4$.
2. Whenever $\exists p, m$ s.t.$\sum_a AM_m(a,p) = 2$ we get into a "termination phase" (which corresponds to Lemma 2) and the algorithm is guaranteed to terminate within at most two more stages. It can be verified that $\sum_a AM_{m+1}(a, \sim p) = 2$ and that both active moves by player $\sim p$ are in the same direction. Therefore, player p's two moves immediately become active at stage $m+2$, without any deletion of an active move by player $\sim p$, terminating the algorithm.
3. $\#AM$ is non-decreasing in m: each iteration adds an active move $(AM(a^*, p^*))$ and may potentially remove at most one active move.[23]
4. For $m > 2$, and before reaching the "termination phase," an active move (a,p) is deleted only when $(a,p) \in ((a^*_{p^*}, a'^*_{\sim p^*}), \sim p^*)$. In particular, at stage m, a deleted move must belong to player $\sim p^*_m$.
5. Observations 2 and 4 imply that once $\#AM = 2$ the algorithm terminates within at most 3 stages. If the two active moves are by the same player then we can use observation 2. If they are by different players, observation 4 guarantees that in the next stage one player will have 2 active moves.

Using all the above, all we need to show is that it is not possible to have an infinite sequence of stages with only one active move in each of them. That is, such that any move that becomes active at stage m, becomes inactive at stage $m+1$. Suppose, toward contradiction, that such an infinite sequence exists. Without loss of generality, consider $m = 2$, in which $AM_2(a,p) = 1$ for some (a,p), and $AM_2(\widetilde{a}, p') = 0$ for any $(\widetilde{a}, p') \neq (a,p)$. If (a,p) is deleted at $m = 3$, it must be that the new active move is such that $AM_3((\widetilde{a}_p, a_{\sim p}), \sim p) = 1$. Similarly, we obtain that $AM_4(a', p) = 1$ and that $AM_5((a_p, a'_{\sim p}), \sim p) = 1$. This gives the following contradiction. By $AM_2(a,p) = 1$ we know that $V_3(a, \sim p) = V_3((a'_p, a_{\sim p}), \sim p)$. By $AM_3((a'_p, a_{\sim p}), \sim p) = 1$ we know that $V_3((a'_p, a_{\sim p}), \sim p) < V_3(a', \sim p) - C_{\sim p}(a_{\sim p} \to a'_{\sim p}, t)$ for any $t < t^*_3$. It is easy to see that $t^*_4 < t^*_3$, so the above implies that $V_5(a, \sim p) = V_3((a'_p, a_{\sim p}), \sim p) < V_3(a', \sim p) - C_{\sim p}(a_{\sim p} \to a'_{\sim p}, t^*_4) = V_5((a'_p, a_{\sim p}), \sim p)$, while by $AM_4(a', p) = 1$ we also know that $V_5(a', \sim p) = V_5((a_p, a'_{\sim p}), \sim p)$. The two last equations imply that $\Delta V_5(a', \sim p) > \Delta V_5((a_p, a'_{\sim p}), \sim p)$, which is a contradiction to the fact that $(a^*, p^*) = ((a_p, a'_{\sim p}), \sim p)$ at $m = 5$. This, together with observation 5 above, also shows that $\overline{m} \leq 8$. ■

Remark 1 *In fact, it can be shown that* $\overline{m} \leq 7$ *because a deletion at* $m = 2$ *according to* $(a,p) \in (a^*, \sim p^*)$ *and* $m = 2$ *implies that there can be only one (rather than two) additional deletions later on.*

[23] Whenever the *ArgMax* is not a singleton, then it is easy to see that we add two active moves by the same player, thus we are done by observation 2 above.