A Theory of Endogenous Commitment

Guillermo Caruana
CEMFI

Liran Einav
Stanford University

Working Paper No. 0206
September 2002

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 this February. We will miss him very much. We have also benefited from discussions with Estelle Cantillon, Hsueh-Ling Huynh, Shachar Kariv, Asim Khwaja, David Laibson, John Leahy, Bart Lipman, Markus Möbius, Dilip Mookherjee, Al Roth, Balazs Szentes, Elena Zoido, participants at the 2002 Annual Congress of the EEA in Venice, and seminar participants at Alicante, Autonoma, BU, Carlos III, CEMFI, Duke, Harvard, Illinois, Kellogg, NYU Stern, Pompeu Fabra, Rutgers, and Tel-Aviv. The usual disclaimer applies.
(Email addresses: caruana@cemfi.es, leinav@stanford.edu).

CEMFI, Casado del Alisal 5, 28014 Madrid, Spain
www.cemfi.es

Abstract

Commitment is typically modeled by giving one of the players the opportunity to take an initial binding action. The drawback to this approach is that the fundamental question of who has the opportunity to commit is driven by a modeling decision. This paper presents a framework in which commitment power arises naturally 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 want, but pay a switching cost if they do so. We show that for two-player games there is a unique subgame perfect equilibrium with a simple structure. This equilibrium is independent of the order 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 many different environments. In particular, we study an entry-deterrence situation and a bargaining setting. The predictions for these are intuitive and illustrate how commitment power is endogenously determined.

JEL classification: C7, C73
Keywords: Commitment, entry, switching costs.

1 Introduction

Commitment is 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. This allows 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. In this paper the set of commitment possibilities is not imposed; it arises naturally from the fundamentals of the model. Thus, issues such as bargaining power, credibility, and leadership can be addressed.

Let us first illustrate with a simple example from the industrial organization literature. Consider a standard entry situation. A potential entrant is considering 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 overcapacity. If we let the entrant commit to an action before the incumbent decides, he will enter the market, forcing the incumbent to accommodate him. 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. Clearly, this model cannot answer the question of who has the opportunity to commit earlier, as it is assumed. The purpose of this paper is to construct a model of commitment that does not rely on an exogenously specified choice of the order of moves. Our framework will allow us to answer which of the two previous outcomes is more likely to arise. In fact, in our context, the two previous outcomes will be two special cases of a wider range of possible equilibrium outcomes. Thus, the model has the additional desirable feature of bringing under a unified umbrella situations that were previously captured only by employing different models.

We consider 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. They can change their minds 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 become so high that the players are completely committed to their announced actions.

The model generally has a unique subgame perfect equilibrium (henceforth spe). 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 as often as they like, in equilibrium they seldom do so. In particular, on the equilibrium path of games with two players and two actions at most one player switches, and when he does so he does it only once.

Our results for two-player games are independent of the order of moves. As long as both players can revise their announcements frequently enough, the exact configuration of when they get to move or in which order has no impact on the shape of the equilibrium. This accomplishes the task described above, namely to lay out a theory in which the commitment power is not driven by modeling assumptions. Moreover, we study the robustness of the equilibrium to changes in the protocol. Throughout the paper we use protocols according to which players move sequentially in a pre-specified order. This restriction simplifies the proofs but does not drive the results. We claim that as long as there exists a sufficient amount of asynchronicity in the timing of actions between players, the qualitative results of the paper would not change. More precisely, if we allow for the order of play to be determined by a random process, as in Lagunoff and Matsui (1997), so that periods of both simultaneous and sequential moves arise, the model would still provide the same unique equilibrium outcome.

The order independence result does not generally extend to more than two players. Nevertheless, we provide an interesting family of N-player games for which it does. In a bargaining setting we obtain an intuitively appealing unique equilibrium outcome. It is worth mentioning that N-player bargaining models normally lack predictive power and are notoriously sensitive to the protocol assumed.[1]

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 of the players–cannot fit all possible scenarios. This framework, however, is flexible enough to accommodate a wide range of interesting economic situations. We have already introduced the example of entry. Consider any new market which is to be opened at a pre-specified date (e.g. as a result of a patent expiration,

[1] For example, in Rubinstein's (1982) model with three or more players, any split of the pie is an subgame perfect equilibrium (Herrero (1985)).

the introduction of a new technology, deregulation, etc.). 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. 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 (e.g. the Fourth of July), 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] This type of competition is also present in other situations, such as network TV programming and the timing of sales promotions. Finally, the model may be also applied to elections and other political conflicts (e.g. the UN ultimatum in the Gulf War) in which a deadline is present, or, as we have already mentioned, to bargaining situations.

This paper shares certain similarities with Admati and Perry (1991) and Gale (1995, 2001). These models have the feature that investment decisions are irreversible. This provides these actions with a commitment component. In our model, however, there is reversibility, but we impose an increasing cost to it. So, eventually players get locked into their actions as well. Our interest, though, lies in studying how this commitment is achieved as time goes by. This feature and the general approach also distinguish our paper from Rosenthal (1991) and Van Damme and Hurkens (1996), who, like us, address the issue of commitment and its timing. Henkel (2002) has similar motivation to ours, and some of his results are related (e.g. the potential for strategic delays), but his approach in extending the one-shot sequential game is somewhat more restricted (only one more move added). More importantly, in his work the players' roles (the leader vs. the passive player) are exogenously imposed, while our main goal is to let these roles arise endogenously from the fundamentals of the model.

Despite being different in its spirit and intentions, the paper most similar to ours is probably Lipman and Wang (2000). Both papers analyze finite games with switching costs. Therefore, both use similar techniques and share similarities in the structure of their equilibria. Lipman and Wang's purpose, 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 big and payoffs that only depend on the final decisions. As a consequence, the intermediate decisions taken by the players do not directly

[2] See Einav (2001) for an empirical analysis of the release date timing game.

impact the final payoffs in our framework, as they do in theirs. In general, this results in different predictions for both models.

Let us finally stress the potential that this model may have for empirical research. The properties of the spe, namely its uniqueness and robustness to the protocol used, make it attractive for the analysis of strategic interactions that are focused on discrete decisions. In such situations multiplicity of equilibria is a burden since it is usually difficult to find properties that are common to all equilibria. On top of that, we provide an efficient algorithm that solves for the spe, which could be of use in the estimation strategy.

The paper continues as follows. Section 2 presents the general model for N players and K actions. Section 3 presents the results for two-by-two games. In Section 4 we deal with several extensions. First we discuss how the results generalize to K actions and how they only partially extend to games with more than two players. Then we stress the robustness of the spe structure to other protocols. Finally we analyze the consequences of imposing some sensible restrictions on the switching cost technology. Section 5 applies the model. We devote special attention to two families of games: entry games and bargaining games. Section 6 concludes. All proofs and the analysis of the examples are relegated to the Appendix, which also contains an algorithm that solves for the spe and a practical demonstration of how to use it.

2 The Model

The model is constructed for N players and finite action spaces. The dynamic structure of the model is as follows. 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 payoffs for each player. Between $t = 0$ and $t = T$ players will have to make up their minds about their final decisions, taking into account that any time they change their minds they have to pay a switching cost. A complete 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. How small it needs to be will be specified later. The idea is that players have many opportunities to switch their decisions.[3] Note that this specification of the grid allows for different configurations of the order in which players move. A natural example of a grid (of fineness $N\tau$) is $g_i = \{(kN+i)\tau \leq T \mid k = 0, 1, 2, ...\}$. However, the general specification may take many other forms.

We introduce at this point two operators that we use often in the analysis. Given a game grid g and a point in time t, denote the next point on the grid at which player i gets to play by $next_i(t) = Min\{t' \in g_i | t' > t\}$. Similarly we have $prev_i(t) = Max\{t' \in g_i | t' < t\}$.

Each player has a finite action space A_i. When player i gets to play at t, he has to take an action from A_i, i.e. $A_i(t) = A_i$. At every point in time all players have perfect information about the previous moves made by everyone.

The very first move by each player, taken at t_1^i, is costless. However, if a player changes his action at t from a_i to a_i', he has to pay a switching cost $C_i(a_i \to a_i', t)$.[4] We impose the following assumptions on the cost 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. There is no cost if there is no switch: $C_i(a_i \to a_i, t) = 0 \ \forall a_i \in A_i$, $\forall i \in \{1, ..., N\}$ $\forall t \in [0, T]$.

3. Costs are very low early on: $C_i(a_i \to a_i', 0) = 0 \ \forall a_i, a_i' \in A_i$, $\forall i \in \{1, ..., N\}$.

4. Costs are very high towards the end: $C_i(a_i \to a_i', T) = \infty \ \forall a_i, a_i' \neq a_i \ \forall i \in \{1, ..., N\}$.

5. Triangle inequality:

$$C_i(a_i \to a_i', t) + C_i(a_i' \to a_i'', t) > C_i(a_i \to a_i'', t)$$

$$\forall t \in (0, T) \ \ \forall a_i, a_i', a_i'' \text{ s. t. } a_i \neq a_i', a_i \neq a_i'', a_i' \neq a_i'', \ \forall i \in \{1, ..., N\}$$

Note that the switching costs do not depend on the actions taken by other players. However, the switching costs across the different players' own possible moves, $a_i \to a_i'$, may be different from each other.

[3] 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. We address this issue later in the paper.

[4] Note that the switching costs are defined over $[0,T]$, independently of the specific grid.

Finally, all that remains to be specified are the payoffs for the different players. Given the actions of all players $\overline{a} = (\overline{a}_i)$, where player i's actions are $\overline{a}_i = (a_i(t))_{t \in g_i}$ and $a_i(t) \in A_i$, and the final actions by all players $a^* = (a_i(t^i_{L_i}))$, player i's payoffs are:

$$U_i(\overline{a}) = \Pi_i(a^*) - \sum_{t \in g_i - \{t^i_1\}} C_i(a_i(prev_i(t)) \to a_i(t), t)$$

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 of all the players, minus all the switching costs he incurred in the process, which depend only on player i's own 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 simply 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.[5] For this reason and to simplify the analysis we abstract from these cases. We will discuss, however, 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 too costly compared to any possible extra benefit achieved in the final payoffs. More precisely, if we define

$$\Delta\Pi_i(a_i \to a'_i, a_{-i}) = \Pi_i(a'_i, a_{-i}) - \Pi_i(a_i, a_{-i})$$

all we need is that

$$C_i(a_i \to a'_i, T) > \underset{a_{-i}}{Max}\ \Delta\Pi_i(a_i \to a'_i, a_{-i}) \quad \forall i, a_i, a'_i$$

It is easy to see that in equilibrium after

$$\overline{t} = \underset{i, a_i, a'_i, a_{-i}}{Max}\ C_i^{-1}(a_i \to a'_i, \Delta\Pi_i(a_i \to a'_i, a_{-i})) \tag{1}$$

no player will ever switch.

[5]In the paper we will use the following measures: (i) for the space of g_i's the following measure: $\mu(B) = \sum_{n=1}^{\infty} \mu_n(B \cap G_n)$, 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.

Second, after a player switches the incurred switching cost is sunk. Given that this is a finite game of perfect information, in the absence of indifference points history is 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 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 $s_i(a,t) \in A_i$ $\forall a \in A$, $\forall t \in g_i$.[6]

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 is implicitly handled by 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 he 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.

3 Analysis of Two-By-Two Games

We start by studying the equilibrium for a given grid, i.e. for a given (Π, C, g). Then we prove the grid invariance property; that is, we show that for all sufficiently fine grids the equilibrium has generically the same structure. From that point on, we are able to abstract from the grid and attach a unique equilibrium to any given (Π, C). Making use of this, we conclude the section with a taxonomy of two-by-two games and other characterizations. We use the following notation. If $i \in \{1, 2\}$ is one of the players, the other player is $\tilde{\imath}$; and if a_i is one of player i's actions, the other is $\tilde{a}_i$.

3.1 An Example

It is useful to start with 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	$5, 3$	$12, 0$

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

Assume that the switching costs are equal across different actions and for both players; that is, $C_i(a_i \rightarrow \tilde{a}_i, t) = c(t)\ \forall a_i \in A_i\ \forall i \in \{1,2\}$. For simplicity, let $c(t) = t$. Suppose that T is big (e.g. $T = 100$) and that players alternate and decide every .01 increment in the following way: the entrant plays at ε, the incumbent at $.01 + \varepsilon$, the entrant at $.02 + \varepsilon$, and so on. The constant ε is a small number ($0 < \varepsilon < .01$) to avoid indifferences.

The game is solved using backward induction. Late in the game the switching costs are so high that no player wants to change his action. When deciding whether to switch or not this late in the game, a player knows that doing so would be a unilateral and final change of actions. Using equation (1) we can find the latest date at which a player would make such a switch. In this example this happens at $9.98 + \varepsilon$, the first node before $\bar{t} = c^{-1}(10) = 10$ at which the entrant plays and at action profile [*Fight, Entry*] (that is, if the most recently announced actions were *Fight* by the incumbent and *Entry* by the entrant). In this case, the entrant would decide to switch and exit the market. At all other action profiles he would not switch at this or any later time.

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. If he decides to switch to *Fight*, he foresees that the entrant will react by exiting the market, which guarantees 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 the one-shot sequential game.

The table below presents the complete equilibrium strategies. In the second column we indicate 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.

Time	Switches
Initial actions	NE by the entrant; NF by the incumbent
$[0.02+\varepsilon, 1.99+\varepsilon]$	$(NF,E) \to (F,E)$ and $(NF,NE) \to (F,NE)$ 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

In section 5.1 we will analyze in detail a generalized version of this entry game. We defer to that point the economic interpretation of the equilibrium. Our focus now is on the equilibrium structure. 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 our 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, these properties are not particular to this game, but common to any given (Π, C, g).

3.2 Structure of the Equilibrium Strategies

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 *$t^* \in g_i$ is a* ***critical point*** *if there exists an action profile $a = (a_i, a_{\tilde{}i})$ such that $s_i((a_i, a_{\tilde{}i}), t^*) = \tilde{}a_i$ and $s_i((a_i, a_{\tilde{}i}), next_i(t^*)) = a_i$.*

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 changes his response at profile a 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 the other player is going to do something new, which in turn changes player i's incentives. In the example above $t = 9.98 + \varepsilon$ is a critical point of the first type and $t = 4.98 + \varepsilon$ of the second.

Note that players could potentially build up very complicated strategies, which could result in different decisions at every single point in time. As a result, there could potentially be as many stages in the game as points in the grid. In what follows we show that this is generally not the case, and that the number of stages in any game is quite limited.

As in the example, we can use equation (1) to find the last point and profile at which a player finds it profitable to make a switch. After this point we enter the last stage of the game, in which no player makes any switches. Next, we characterize the early stages of the game. First, we state a simple lemma.

Lemma 1 *If $s_i((a_i, a_j), t) = \tilde{a}_i$ then $s_i((\tilde{a}_i, a_j), t) = \tilde{a}_i$ $\forall a \in A$, $\forall t \in g_i$, $\forall i \in \{1, 2\}$.*

In other words, a player never wants to switch both from one profile to another and the reverse at the same time. If he prefers to switch, he does so in one direction only. Thus, if player j is playing a_j, player i's strategies can be described either as always playing a_i, always playing $\tilde{a}_i$, or sticking to his previous action. In the first two cases, player i's strategy, given a_j, is independent of his own previous action. In such a case, we say that player i has an *active switch* at a_j. The following proposition shows that if a player has an active switch for each of the opponent's actions at some point of the game, then the same is true at all earlier points in the game. Furthermore, prior to this point, his actions are ignored by his opponent.

Proposition 1 *If there exists a player i and a point in time $t \in g_i$ such that $s_i((a_i, a_j), t) = s_i((\tilde{a}_i, a_j), t)$ $\forall a_j \in A_j$, then for both players $p = 1, 2$ and for any $t' \in g_p$ such that $t' < t$ the strategies are independent of a_i, i.e. $s_p((a_i, a_j), t') = s_p((\tilde{a}_i, a_j), t')$. Moreover, there are at most three stages in the interval $[0, t]$.*

Another way to express this result is as follows. As long as the switching costs for one player are low enough, he remains completely flexible. Given that his actions are therefore not binding, they are not taken seriously by the other player–they are cheap talk. Thus, there are no strategic interactions during this early period of the game: for one player, his switches are being ignored, and for the other player, who faces an individual decision problem, it is always optimal to decide on his action as early as possible, saving on switching costs. Given this, it is not surprising that the number of stages in this early part is small. In the example above, $t = 4.98 + \varepsilon$ is the first point at which a player (the entrant in this case) has two active switches. Prior to this point, the reader can check that the strategies of both players are not contingent on the entrant's previous action and that this early part of the game has two stages.

Next we obtain the following lemma.

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

The lemma states that the strategies for both players are held constant throughout a stage. This was not directly implied by the definition of a stage; inside a stage there could have potentially been a player i, a point $t \in g_i$, and a profile $a = (a_i, a_{\sim i})$, for which $s_i((a_i, a_{\sim i}), t) = a_i$ and $s_i((a_i, a_{\sim i}), next_i(t)) \neq a_i$. In the proof we show that such behavior would only make sense if the consequences of either staying or switching had changed between t and $next_i(t)$. This would reflect the fact that player j has a critical point between these two points, contradicting the fact that t and $next_i(t)$ were in the same stage. Note that an important consequence of this lemma is that on the equilibrium path of any subgame switches occur only at the beginning of each stage.

The next proposition shows that the simple structure of the equilibrium at the beginning and at the end of the game is also present in its intermediate part. We show that given a game (Π, C, g), its spe can be fully described by a small number of stages. Moreover, the total number of stages does not exceed eight.

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

We provide here some intuition for why the number of stages is bounded. Be aware that all the following arguments are done using backward reasoning. We have already shown (Proposition 1) that if the play has reached a point at which one player is completely flexible, there are at most three stages before that. Consider now a stage in which both players have a profile at which they switch. To pass to the previous stage there has to be a change in the strategy for, say, player i. This can, potentially, eliminate the active switch by player j, but cannot change player i's incentive to continue switching at the profile he was switching from already. Thus, this new switch would make us move into the early part of the game discussed in Proposition 1.

We also know that there is a final stage in which players are completely committed, and that this stage starts at $\bar{t}$, the final point at which one player decides to switch actions. Right before this point, we are at a stage with one active switch. Once player i wants to switch at a given profile, he would also do so in the previous stage (his costs go down) unless there is a switch by the other player that becomes active and changes player i's incentives. Therefore, the key is that a player's incentive to switch from a given profile is changed only as a result of a change in his opponent's strategy.

Combining the arguments above, all that needs to be shown is that there cannot be a long sequence of stages in which only one player switches, and only at one profile. Going backwards, in order to have such a sequence, every time a new switch (by a given player from a given profile) becomes active, it has to eliminate the incentives of the other player to switch. It can be shown that after a small number of switches that follow each other no additional eliminations are possible.

The proposition's statement is generic because our argument assumes no indifference at all $t \in g$. In fact, it would be easy to show that whenever there are multiple equilibria the proposition still holds for each of the equilibria separately. Multiple equilibria can be generated by indifference at different points in the game: either (i) at a decision node (a, t) in the middle of the game, such that $t \in g_i$ and $t \neq Min\{t' \in g_i\}$, in which player i is indifferent between switching or not; or (ii) at the first decision of player i in the game, in which he is indifferent between playing a_i and playing $\tilde{}a_i$. Notice that in case (i) the multiplicity can be avoided by slightly perturbing the grid. Even if we considered the different equilibria, as long as the grid is fine enough, all of them have the same stage structure. The only difference among them is that the critical point is slightly shifted: either player i switches for the last time at the point where he is indifferent or he does it just before, at $prev_i(t)$. Thus, the strategies played in each of the stages, and therefore the outcome of the game, remain the same. Case (ii), however, may provide equilibria with different outcomes. If a player is indifferent at his initial move, the spe can take very different paths. This happens, for instance, when a player has the same payoff values in all the cells of the matrix Π.

Finally, the proof makes use of an algorithm (see Appendix B.1 for its description) that computes the stages and the corresponding strategies played within each of them. The algorithm is of interest in its own right, because it finds the equilibrium strategies in an efficient way, without the need to apply backward induction at every node. The algorithm solves the game backwards by computing continuation values only after each critical point, and skipping all other decision nodes within a stage.

In sum, this section has shown how commitment in this model is achieved through switching costs. 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–each new critical point introduces a new commitment possibility for a player.

3.3 Strategic Delays

In this subsection we take a different approach to the characterization of the spe strategies, using 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, one would rather do it early than late in order to save on costs. On top of that, given the perfect information nature of the game, waiting has no option value. Nevertheless, we show that although delays are costly, they may occur in equilibrium for strategic reasons. However, they are very limited, and cannot occur very often.

Consider a point $t \in g_i$ in the middle of a stage and a profile a such that player i switches at (a, t). Given that t is in the middle of the stage, it is implied that player i also switches at $(a, prev_i(t))$. Thus, the switch at (a, t) is not delayed. It is just an off-equilibrium adjustment by player i. If player i found himself at (a, t) he would immediately switch. Below we formally define a delayed switch.

Definition 3 *Consider a decision node (a, t) for $t \in g_i$ at which player i switches, i.e. $s_i(a, t) = \tilde{a}_i$. This switch is a **delayed switch** if there exists a decision node (a', t') with $t' < t$ such that $t' \in g_i$ and (a, t) is on the equilibrium path of the subgame (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 idea is that a delayed switch is a switch by player i that could have been made before, but was delayed for strategic reasons. 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 the game (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, such elimination leaves us with a unique outcome, so there are no further switches.

Note that the previous proposition reduces the total number of switches for any subgame to a maximum of three (there are at most two non-delayed switches and one which is delayed). In particular, this excludes the possibility of a cycle that visits the four profiles. More importantly, if we apply this result to the equilibrium path of the full game, we obtain the following:

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.*

While we have shown above that at any subgame the equilibrium path includes at most one delay, it may still be the case that different subgames have different delays. It turns out that for some type of delays this is not the case, so that the delayed switch is always taken at the beginning of a certain stage of the game. To make this claim, it is useful to distinguish between two different types of delays, according to the following definition.

Definition 4 *Consider a decision node (a, t) for $t \in g_i$ at which player i switches, i.e. $s_i(a, t) = \tilde{}a_i$. This switch is a* ***real delayed switch*** *if (a, t) is on the equilibrium path of the subgame $(a, prev_i(t))$. A delayed switch which is not a real delayed switch is said to be an* ***immediate delayed switch****.*

An immediate delayed switch is not "really" delayed. The player waits for only one period, at a different profile, in order to make his opponent switch first. A real delayed switch, however, captures the idea of a player really waiting at a given profile. He waits at this profile until the point where his opponent becomes committed. Unlike immediate delayed switches, this wait lasts longer than a single period. Notice also that while real delayed switches may occur on the equilibrium path, immediate delayed switches only happen off the equilibrium path.

The next proposition shows that there is at most one real delayed switch by each player in the whole game. On the equilibrium path though, as we saw in Corollary 1, only one of these switches can be visited.

Proposition 4 *Given a cost structure C, generically for every (Π, g), the unique spe strategies of the game (Π, C, g) are such that if (a, t) is a real delayed switch for player i, then there does not exist another real delayed switch $(a', t') \neq (a, t)$ for the same player.*

Such a switch by, say, player i, is always done at his first opportunity after a critical point of player j. Player i switches only once he knows that player j will not react. This type of strategic delays correspond to the sequence of one-switch stages that we discussed in the previous subsection. The fact that each player has only one real delayed switch in the game implies that there can only be two eliminations of switches. Therefore, a long sequence of one-switch stages cannot occur in equilibrium. This is another way to convey intuition for why the number of stages in equilibrium is bounded and small.

3.4 Grid Invariance

We want to compare the equilibria of a game with different grids. Clearly, the exact position of the critical points depends on the grid chosen, so in general they cannot be the same for different grids. However, we will show 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 first define a notion of equivalence between two equilibria.

Definition 5 *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 can now 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 a 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. The limit does not exist when a critical point for the two players coincide. We discuss this below, when we list the non-generic cases for which the theorem fails.

The minimal fineness of the grid we allow, α, crucially depends on how far apart from each other the critical points for the two players are. It also depends on the slope of the cost function at the points of real delayed switches, whenever they exist. Within a stage, a player switches at most once. Thus, all he needs is one opportunity to play at every stage. Including more points on the grid cannot change his strategic incentives. It may slightly change the incurred switching costs, but not the overall structure of the equilibrium. Therefore, once both players have an opportunity to move between every pair of consecutive critical points, the grid has no impact.

In the proof we make use of another algorithm (see Appendix B.2), which solves for the limit of the equilibria by taking the fineness of the grid to zero. Independently of its role in the proof, this algorithm is useful because it efficiently solves for the invariant equilibrium of any game (Π, C).[7]

Let us stress the importance of the quantifiers used in the theorem. First, we state the result *for almost every* grid. As before, this is done in order to avoid the multiplicity of equilibria.[8] We state that the grid invariance property is satisfied for *generically* every Π. This is done in order 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, as we already mentioned before, if one of the players is initially indifferent between two different actions. Clearly, a slight perturbation of the payoff matrix would eliminate such multiplicity, making it clear why these points are of measure zero. Second, the theorem also rules out an additional measure-zero case. This case 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 a given grid g, the equilibrium may depend on whether $prev_i(t^*) < prev_j(t^*)$ or the reverse. Clearly, even for fine grids, this inequality can go either way. This has the consequence that the equilibrium depends on the order of the moves in the neighborhood of t^*. Therefore the limit of the equilibria does not exist.[9] A Slight perturbation of the payoffs of one of the players separates the critical points for both players, making the problem disappear. In the symmetric game of Battle of the Sexes

[7] See Appendix B.4 for an example on how to solve the equilibrium of a game using the Limit Algorithm.

[8] We actually could have stated the Theorem "for every grid", because, as we argued in the previous section, as long as the grid is fine enough, all (multiple) equilibria are essentially the same. We prefer to state it using "for almost every grid" for consistency (we previously decided to ignore the multiplicity cases for convenience).

[9] For completeness, we should note that in some cases, even when the critical times are common for both players, the equilibrium outcome is unique, and the limit does exist. This heppans when the common critical time is associated with two switches that are "unrelated", i.e. when the origin of one is neither the origin nor the destination of the other. This is illustrated by the Prisoners' Dilemma example in Appendix C.2.

(example C.1) or in the symmetric Matching Pennies (example C.3) we observe grid dependence, but any perturbation provides a unique outcome.

Given that our results hold when the grid is fine enough, is natural to ask whether it is possible to directly construct an analogous game defined in continuous time, as the limit version of the game suggests to do. The reason we use a finite game is to avoid the typical problem of non-existence of equilibrium that is common to continuous games. In our framework this happens when a player wants to create a threat to switch twice. This threat is initiated just before a critical time. Initiating the chain of switches exactly at the critical time would imply that the second switch would be taken too late, making the whole threat too costly. On the other hand, initiating the chain reaction before the critical time would not be optimal, because one could do better by delaying the threat a little bit longer. This situation does not arise for all (Π, C).[10] If it does not, we conjecture that the equilibrium exists, that it is unique, and that it coincides with the limit of our equilibria.

3.5 Taxonomy and Characterizations

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. Unfortunately, the short answer to this is no. There does not exist a simple "reduced-form" function that maps the fundamentals of the game into its outcome. We have shown that the equilibrium has a lot of structure. Still, the incentives in the eight possible stages may get combined in different ways, providing a rich variety of possible dynamic interactions. In this section, we further discuss these issues and provide some special cases in which "shortcuts" are possible.

It is easy to see that the equilibrium payoffs for each player i are at least his maxmin payoffs of the one-shot game. Player i could always guarantee himself at least his maxmin payoffs by playing his maxmin strategy and never switching.[11] This simple argument implies that some of the potential outcomes cannot result in equilibrium. This may suggest that in some cases one could easily predict the equilibrium outcome. It turns out to be somewhat more subtle. For example, consider the following definition.

[10] It arises when in the grid-invariant equilibrium there is a subgame equilibrium path that involves three consecutive switches. Note that the last switch in such a chain is always an immediate delayed switch.

[11] A similar argument for minmax payoffs is not true. The reason for this is that, in the dynamic setting, a player cannot react costlessly to the actions taken by his opponent.

Definition 6 *An action a_i is called **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 the super-dominant play by player i. 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, forcing him to choose the other action by using a delayed switch. On the equilibrium path, player i starts by playing his dominated strategy, switching to his super-dominant strategy only if player j "behaves", and only after player j is fully committed to his "disciplined" behavior. Example C.4 illustrates this case.

In addition, note that a dominant strategy, which is not super-dominant, is not necessarily played in equilibrium. See example C.2 of the Prisoners' Dilemma game, in which cooperation may be achieved in equilibrium.

Another simple characterization uses the following definition of a defendable profile, and is a generalization of a similar definition used in Lipman and Wang (2000).

Definition 7 *The profile a^* is **defendable** by player i if $a^* = \underset{a}{ArgMax}\ \Pi_i(a)$ and*

$$C_j^{-1}(a_j \to a_j^*, \Pi_j(a^*) - \Pi_j(a_i^*, a_j)) > C_i^{-1}(a_i^* \to a_i', \Pi_i(a^*) - \Pi_i(a_i^*, a_j))\ \ \forall a_j \in A_j, a_i' \in A_i$$

Note that by construction a defendable profile is always a Nash Equilibrium of the one-shot game. Clearly, player i's best-response to a_j^* is a_i^*. In addition, it is easy to check that a_j^* is the best-response to a_i^* (note that $\Pi_j(a^*) - \Pi_j(a_i^*, a_j)$ is always positive). It can also be verified that there is at most one defendable profile in a game. It should be noted, however, that this is a very strong condition, so the scope of the next proposition is quite limited. The proposition states that when such a defendable profile exists, the spe of the dynamic game will have this profile being played throughout the dynamic game, with no delays. Thus, the proposition provides a sufficient condition for a Nash Equilibrium of the one-shot game to be the spe of the dynamic game.

Proposition 5 *Given a game (Π, C), if there exists a defendable profile a^*, then, on the equilibrium path, this profile is played throughout the game.*

The defendable profile is achieved without delay because the opponent cannot credibly threat to play something different from a_j^*. Once player i plays a_i^*, for any strategy of player j there exists a time after which player i is fully committed to a_i^* while player j still prefers to switch to a_j^*. Player j understands this, and avoids incurring any fruitless costs by directly playing a_j^* at the beginning of the game. Note that the proposition provides only a sufficient condition for the best outcome of a player to be achievable. This condition is clearly not necessary. Consider for instance the Battle of the Sexes game discussed in Section 5.2.

Finally, the previously introduced notion of two games having *essentially* the same equilibrium allows us to construct a classification of all two-by-two games. If we abstract from quantitative considerations and just focus on the direction of the incentives for each player at each stage (to switch or not to switch), given the finite nature of the action spaces there is a limited number of configurations that a stage can adopt. By combining this with the fact that the number of stages is always less or equal to eight, this provides a full taxonomy of all the possible two-by-two games. One can show that, subject to relabeling of the players and strategies, all games can be grouped into exactly 75 types of dynamic interactions, which are fundamentally different from each other.[12]

4 Extensions

4.1 K Actions, N Players

The model was constructed for N players and any finite action space, but the results in the previous section were presented only for two-by-two games. This was done mainly for two reasons. First, the two-by-two case captures all the richness of the model. Analyzing more general cases just adds technical complexity. Some of the results become more difficult and cumbersome to prove. The second reason is more fundamental, and involves games with more than two players, for which the grid-invariance property generally fails. In this subsection we argue, in a more suggestive than formal way, how the analysis for the two-by-two case is fully generalizable to accommodate two-player games with bigger action spaces and what are the limitations for N-player games. We follow the same outline that we used to present the results for two-by-two games.

Take a generic (Π, C, g) for N players and K actions. The game has a unique equilibrium, which allows us to define our concept of stages and critical points again. The structure of the equilibrium is similar, i.e. there is a finite number of critical points and corresponding stages,

[12] For one-shot two-by-two games, the analogous classification would provide four types of games: (i) games with both players having a dominant strategy; (ii) games with only one of the players having a dominant strategy; (iii) games with no dominant strategy, such that there are 2 pure-strategy Nash Equilibria; (iv) games with no pure strategy Nash Equilibria.

within which the strategies remain constant for long periods of time.[13] Lemma 1 and Proposition 1 are easily generalizable, and are stated and proved for the general case in the Appendix. If for a given profile a_{-i} played by the other players, player i has $K-1$ active switches, then (using this general version of Lemma 1) there exists a unique action a_i such that $s_i(a_i', a_{-i}) = a_i$ for all a_i'. Early in the game, when this is true for all a_{-i} (which means that player i switches at $K^{N-1}(K-1)$ profiles), we say that player i is fully flexible. Therefore his announcements have no credibility and are ignored by everyone. Once this happens for all players but one, we enter into the early part of the game.

The number of stages is also bounded, so that Proposition 2 is generalizable. Thus, we obtain that the structure of the game can be summarized in a finite and small (compared to the potential size of the grid) number of stages, and that this number does not depend on the grid used, as long as it is fine enough. A proof of the finiteness of the number of stages could be constructed using the same techniques as for the two-by-two case, i.e. with the help of an algorithm. A yet unresolved question regards the maximal number of stages in the equilibrium of a game with N players and K actions. It is easy to show that this number has to be at least the number of possible action profiles, K^N. The actual number is greater than this bound. For $(N, K) = (2, 2)$ we have already eight stages; simulations for $(N, K) = (2, 3)$ results in about 20 to 30 stages; and for $(N, K) = (3, 5)$ the number is already above 500 stages.

The notion of strategic delays is also generalizable. The number of switches on the path of any subgame is bounded, and each player has a limited number of real delayed switches. Again, this result would provide an alternative way to show that the number of stages is bounded.

The characterizations provided in 3.5 can be generalized as well. Obviously, even with K actions and N players, each player can guarantee himself his maxmin payoffs. In addition, the defendability result holds for K actions (as stated and proved in the Appendix), and it is quite easy to construct a modified definition of defendability in order to show that the results hold for N players as well.

Let us now discuss the grid invariance property. For two-player games this property extends nicely. 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. One way to prove it formally is to use the same techniques

[13] While this is true for $N = 2$ and $K > 2$ actions, for $N > 2$ one has to modify the definition of stage to accommodate the following pattern. For $N > 2$, one can construct games for which there is a stage in which at a given profile the strategies, instead of being constant, follow a cyclical pattern that repeats itself.

as for the two-by-two case, namely to construct an algorithm and then check its convergence properties. As we have already mentioned, the grid invariance property fails for $N > 2$. The reason for this is that the order of the moves plays a crucial role in this case. 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 1, he will move to a_3, knowing that player 1 will move next. However, if player 3 gets to play only before player 2, then 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 even if the grid is very fine. Although this is a negative result, we would like to stress that there are interesting families of N-player games with an equilibrium which is robust to changes in the grid. An interesting example is the bargaining game that we present later.

Finally, it is worth pointing out why we have only considered finite action spaces. The reason for this is that with an infinite number of actions, we could have an infinite number of critical points. Given that one cannot separate an infinite number of points with a finite grid, the grid-invariance theorem fails: in general the equilibrium may depend on the grid, even when the grid is very fine.

4.2 Robustness to Other Protocols

As we have already highlighted before, our results for two-player games are invariant to the choice of the grid because players are given the opportunity to react at all the relevant points in the game. Now, we claim that this argument is robust to other protocol specifications. In particular, as long as there exists a sufficient amount of asynchronicity in the timing of actions between players, the qualitative results of the paper would not change.

Maintaining the same switching costs and payoff structure, consider the following general protocol. There is a discrete grid at which players might get an opportunity to play. Once the game reaches each node, a random device decides whether a player gets to play or not.[14] This protocol captures the notion that players may play very often, but that the exact moment at

[14] For instance, consider an i.i.d. process such that, each period, each player plays with probabilty $p < 1$.

which they play next is unknown. We require, however, that the event in which only one player decides occurs with a positive probability.

In this section we suggest that this framework would provide an essentially unique equilibrium, and that this equilibrium is qualitatively the same as the one presented in the rest of the paper. We accomplish this task by first analyzing a game with full synchronicity (simultaneous play in all periods). We suggest that the unique equilibrium of our sequential structure is also an equilibrium in the simultaneous-move game. We then describe how multiple equilibria arise in this case, and how the introduction of asynchronicity solves the multiplicity problem.

The analysis of the simultaneous protocol game involves some extra technical complications, but is similar in spirit to the one performed in this paper. Thus, one could prove that all the multiple equilibria have a stage structure. The number of stages is bounded, and once the (common) grid is fine enough, the stage structure is independent of the grid chosen. More importantly for us, the unique equilibrium provided by the sequential protocol is always one of these equilibria. We argue this by comparing the structure of the equilibrium for the simultaneous case with the structure of the sequential. The comparison is done for each of the possible different stage configurations in the sequential case.

First, consider a stage in which one of the players does not switch at any profile of the sequential-move game. It is easy to verify that these are also his simultaneous-move strategies, and that the opponent's strategies remain unchanged. Similarly, when each player has exactly one profile in which he switches but these two switches are "unrelated" (so that the origin of one is neither the origin nor the destination of the other), or when one of the players has two switches, the simultaneous-move strategies remain the same.

Consider now a stage for which, under the sequential protocol, each player has exactly one profile at which he switches. Suppose also that the two switches follow each other, so one's destination profile is the origin profile of the other. In this case, with a simultaneous structure, the equilibrium play involves mixed strategies at certain nodes. Let a be the profile at which, say, player 1 initiates the sequence of two consecutive moves (the last point at which he moves at that profile). Player 2 switches immediately after, at profile $(\tilde{a}_1, a_2)$. With simultaneous moves, one can show that the strategies at profile a (and only at profile a) will be mixed. If player 1 switches at a (as in the sequential game) then player 2 wants to switch immediately as well, rather than wait and switch in the next period. But if player 2 switches at a, player 1 would rather fool him and not switch himself.[15] Hence, the equilibrium strategies have to involve mixing. It can be shown, however, that as the grid becomes finer, player 2's cost of waiting one period diminishes, so that the mixing probabilities at profile a converge to the pure strategies of the sequential structure. Namely, at profile a player 1 switches with probability that goes to one, and player 2 switches with probability that goes to zero. Note that the mixing may also occur on the equilibrium path, if (and only if) the sequential-move game has a strategic delay. Therefore, for a given finite grid the equilibrium outcome becomes random. Still, as the fineness of the grid approaches zero, the mixed strategies converge to the pure strategies of the sequential-move game.

There is one remaining type of stage to cover. This is the one in which, under the sequential structure, the only active switches are made by both players at the same profile a. This situation is the potential source of multiple equilibria in the simultaneous structure. It is possible that at any point on the grid during this stage, on equilibrium, either one player switches, or the other does, but not both. 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 in which the low coordination profile is chosen at every period. 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 in our setup has the same effects as in Lagunoff and Matsui (1997). It breaks down this "cursed" string of beliefs and restores uniqueness. In our setting, 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 restores uniqueness. On top of this, the probability of both players playing at every stage becomes almost certain as well, and this assures that the decisions made early in a stage are the same as in our previous analysis. Late in a stage, the probability that one player may not play again in 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, so in the limit, the equilibrium structure (strategies, stages, outcomes, etc.) remains the same.

Finally, it is worth mentioning how this reduction in the number of equilibria applies also to other models. In particular, Lipman and Wang (2000) consider a simultaneous-move structure. They obtain multiple equilibria and have to deal with mixed strategies. We conjecture that a sequential structure would restore uniqueness in their model, just as it does in ours.

[15] This is guaranteed by the fact that under the sequential structure player 1 does not switch at $(a_1, \tilde{a}_2)$.

4.3 Invariance to the Switching Cost Structure

The shape of the equilibrium of a game (Π, C) depends crucially on the choice of the cost function. It is clear that an increase in the switching costs of one player has the consequence of increasing his commitment power. We show below that simple restrictions on the cost function, or, more precisely, restrictions on the relationship between the cost functions across different players and different moves, make the results of the model invariant to changes in the underlying cost functions. For instance, in the example presented in Section 3.1, the reader can check that the actual choice of the common cost function did not play any role in the determination of the stage structure of the equilibrium.

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. More precisely, consider $C_i(a_i \to a_i', t) = c(t)\ \forall a_i \neq a_i'$, where $c(t)$ is, as usual, continuous and strictly increasing in t, with $c(0) = 0$ and $c(T) = \infty$. Clearly, the critical points in time are going to vary for different choices of cost functions, but the key point is that the stages are going to remain exactly the same. Let us provide some intuition on why this is so. The order of the different stages of the game is a consequence of the order of the critical points, which are all of the form $c^{-1}(\Delta V)$, where ΔV is a difference between two continuation values. Given that the cost function is monotone in t, its inverse is also monotone in ΔV. Therefore, the ordering of the critical points is invariant to the particular choice of the cost function. This argument is not complete. As we already know, whenever there is a delayed switch (on or off the equilibrium path), the continuation values are updated by subtracting the switching cost at the time of the critical move. The level of the switching cost paid is sensitive to the choice of $c(t)$. This is where it is crucial to have the same $c(t)$ for both players. A delayed move by, say, player i is always done after the passage of a critical point that is associated with the other player, player j. This critical point satisfies $t^* = c^{-1}(\Delta V_j)$, so in the limit and right at t^* the value of executing the delayed switch for player i will involve subtracting a term of the form $c(t^*)$, which is equal to $c(c^{-1}(\Delta V_j)) = \Delta V_j$. Hence, the continuation values for player i will depend only on his previous continuation values and on the difference in continuation values for the other player, but *not* on the cost function. This also implies that the equilibrium payoffs are not affected by the choice of the cost function, even if switching costs are incurred in equilibrium.

Notice the power of the result. Given a normal form game Π, if one is willing to accept these restrictions on the switching cost technology, the model provides an essentially unique equilibrium which is independent of the cost technology and of the protocol used. The equilibrium of (Π, C, g) depends only on Π.

The essence of this strong result is also maintained under weaker restrictions on the switching cost technology. If the costs are identical across moves, but can differ *proportionally* across players, i.e. $C_i(a_i \to a_i', t) = \theta_i c(t)\ \forall a_i \neq a_i'$, then the equilibrium only depends on (Π, Θ), where Θ stands for the full vector of θ's. This can be easily proven by simply realizing that the equilibrium structure of such a game is the same as the one that has payoff matrix $\Pi^* = \left\{\frac{\Pi_i}{\theta_i}\right\}_{i=1}^N$ and the same cost function, $c(t)$, for all players and moves. This fact, combined with the previous analysis gives us the result.

In fact, the previous statement can be further generalized. As long as the cost function is of the form $C_i(a_i \to a_i', t) = \theta_i^{a_i \to a_i'} c(t)\ \forall a_i \neq a_i'$, so that all costs in the game are proportional but may vary across players *and* actions, we can still get an essentially unique equilibrium for a given (Π, Θ). Note, however, that there is no simple normalization of the payoff matrix that can be applied for this case. This is due to the fact that the typical normalization across rows or columns which has no impact on the determination of the Nash Equilibria of one-shot normal form games, cannot be performed in this framework. The reason is simple: given that players foresee further moves by other players, their comparisons might involve payoffs from different rows or columns.

5 Applications

5.1 Entry Game

We now proceed to analyze the standard entry deterrence problem presented in the introduction using the framework constructed. There are several reasons why we have chosen this specific problem. How entry decisions are taken is one of the central questions in the industrial organization literature. More important for the purpose of this paper, this is a situation in which commitment plays a central role. Furthermore, the incentives in the entry game are rich enough to create a wide variety of dynamic stories, resulting in different and sensible outcomes. It also illustrates how different commitment possibilities endogenously arise as a function of the parameters of the model.

Let us first explain in more detail how an entry problem can be described within our framework. Clearly, there has to be some type of fixed deadline present. Such fixed deadlines naturally arise in many settings. One may think, for example, on the expiration of a patent on a certain drug, on the introduction of a new hardware technology, or on the scheduling decision for the release of a new product. The values in the matrix Π capture the resulting payoffs of the ex-post competition given the final decisions taken. As for the increasing switching costs, one story consistent with this assumption is the following. Imagine that the incumbent fights by investing in some type of expansion (say, in machinery). In order to be able to do so, at any point before the opening of the market, the incumbent needs to contract the delivery of the required investment for the opening day, when it is actually needed. The sooner these contracts are written, the less costly they are. One reason for this can be that the machinery suppliers may be aware of the deadline and may charge the incumbent more for it. Alternatively, the supply of the factors that are available at the deadline decreases over time, because they become committed to other tasks. For the same reason, if these contracts have to be nullified, the later they are nullified, the higher the penalty the incumbent has to pay. Similarly, the entrant writes contracts regarding the rental of new office buildings, equipment and employees, and writing or nullifying these contracts becomes more expensive as the deadline approaches. Once the opening of the market is close enough, it is reasonable to assume that writing new contracts or nullifying the existing ones becomes too costly.

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 the parameters satisfy the following assumptions:

- The monopolistic payoffs are greater than the duopolistic ones: $M > D$, $m > d$.
- Conditional on the entrant decision, fighting is costly: $D > d$, $M > m$.
- By entering, the entrant earns positive payoffs only if he is not fought: $a > 0$, $b > 0$.
- The monopolistic profits are greater than the sum of the duopolistic ones: $M > D + b$.
- The incumbent would rather fight and deter entry than accommodate entry: $m > D$.

For simplicity only, we assume that the switching costs are equal for both parties, and across different actions. All the (qualitative) results hold for any general cost structure that satisfies the conditions given in Section 2. Denote the switching cost function by $c(t)$, and its inverse by $c^{-1}(v)$.

As described in the introduction, by considering the one-shot sequential games, one exogenously gives all the commitment power to one party or the other, obtaining two different spe outcomes. 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*].

Now we analyze the problem using our dynamic framework. From Theorem 1, we know that, as long as the grid is fine enough, the solution is going to remain invariant. What are the possible outcomes that can arise? Each player can guarantee himself his maxmin payoffs. Thus the incumbent obtains at least D, and the entrant obtains at least payoffs of zero. Therefore [*Fight, Entry*] cannot be the final spe outcome of the game. This leaves three possible outcomes and the question of whether these outcomes are achieved immediately or only after a strategic delay.

The following proposition completely characterizes all possible equilibrium outcomes. It shows that only four possible cases may arise in equilibrium: the three outcomes played immediately, and one more which involves a strategic delay. The four cases create a partition of the parameter space.[16]

Proposition 6 *The spe outcome of the entry game is:*

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

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

Given that neither of the players wants to stay at [*Fight, Entry*] forever, and that both would rather have the opponent moving than themselves, 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 and commit not to move away from it. Given that we have assumed the same switching cost technology for both parties, the winner is determined by comparing the benefits of making the move ($D-d$ for the incumbent and a for the entrant). The player with smaller benefits is

[16] More precisely, this partition excludes the boundaries of the four sets specified in the proposition. Clearly, this set has measure zero.

then able to commit first. The other party foresees this and moves away immediately. Thus, when $a < D - d$, the entrant is able to commit first. This forces the incumbent to accommodate, resulting in [*No Fight, Entry*], the best outcome for the entrant. This is case (i) of the proposition.

If $D - d < a$ (cases (ii) to (iv) of the proposition), 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, unlike the case in which the entrant wins, this case introduces an additional conflict. While the incumbent is happy deterring entry, he can do so at different costs. He could just fight forever, but this is quite costly, so, if possible, the incumbent would prefer to either deter entry by not fighting at all, or to stop fighting late in the game. These are exactly the three cases described in (ii) to (iv), which correspond to different levels of commitment power by the incumbent.

Much of the intuition regarding the parameters that give rise to the respective cases comes from the relative size of the parameter b and its impact on the incumbent's strategies at action profile [*Fight, No entry*]. In case (ii), the key is that b is very high, so that the incentives for the entrant to enter, once the incumbent does not fight, are high. In such a case, at every point in time when it is still profitable for the incumbent to quit fighting, it would be profitable for the entrant to react by entering. Thus, the only way entry can be deterred is through fighting. The next two cases are more favorable for the incumbent, allowing for entry deterrence without actually fighting.

In case (iii), [*No Fight, No Entry*] is achieved by the incumbent, but only after paying the cost of strategically delaying the switch to *No Fight.* This happens when b has an intermediate value. On one hand, it is low enough so that late in the game the incumbent can stop fighting, knowing that the entrant is committed to staying out. On the other hand, it is still high enough that, earlier in the game, if the incumbent decided to switch to *No Fight*, the entrant would enter. The entrant would do so because the incumbent cannot credibly threat to restart fighting. Thus, the only way in which the incumbent can force the entrant out of the market is by fighting in the beginning and switching to not fighting later at $c^{-1}(b)$, which is the point after which the entrant is committed to staying out.

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, if he is subsequently fought, he still finds it profitable to exit again. 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),[17] 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 have been 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.

5.2 Some Other Interesting Games

In this section we briefly discuss the results of applying our framework to other important types of two-by-two games.

We start by studying coordination games. Consider the following normal form game in which all payoff parameters are strictly positive.

	L	R
U	a_1, b_1	$0, 0$
D	$0, 0$	a_2, b_2

The one-shot game has three Nash Equilibria, two in pure strategies and one in mixed strategies. For such coordination games, our model is able to select one of the pure equilibria, solving the coordination problem. In particular, if there is no conflict of interests (if, for example, $a_1 > a_2$ and $b_1 > b_2$) the Pareto preferred outcome is the one chosen. The dynamic structure fixes any initial coordination failure: if, for some reason, one player starts by playing the "wrong" action, the other player would keep playing the "good" one, and the former can switch back and "correct himself". More interesting is the analysis of the cases in which there is a conflict of interest among the parties (if, for example, $a_1 > a_2$ and $b_1 < b_2$). This correspond to the well known game of the Battle of the Sexes, which we fully analyze in example C.1. The player who wins and obtains his preferred outcome is the player who is able to commit faster to this action. We can think of the equilibrium strategies as an off-equilibrium war of attrition. If the players kept announcing their favorite "event" repeatedly, who would be more stubborn? This is determined by comparing the players' interest in the other player's event. The one who likes it less is able to commit first not

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

to give up any longer (the extra utility obtained is not worth the switching cost). The other party foresees this in advance and gives up immediately.

The previous examples are cases in which our model can be thought of as a way to select from among the multiple Nash Equilibria of a one-shot game. There are many instances though, in which the spe final actions of our model are not an equilibrium of the corresponding one-shot game. We have already seen an example of this in the entry game. Another prominent example is the Prisoners' Dilemma, which is discussed in detail in example C.2. There we show that cooperation can be achieved under certain circumstances. Just as in infinitely repeated games, cooperation is achieved by the fact that defectors can be punished. Here the punishment is switching to defection if the opponent switched to it. This is not always credible, but when it is, mutual cooperation can be sustained. In the context of a political crisis, when the benefits of responding to an attack are high, both sides would rather keep peace and cooperate. When the one-sided attack generates great benefits to the attacker but gives the enemy little incentives to respond, then both sides foresee the lack of credible punishment strategies, and hence defect (attack) from the beginning.

The games analyzed in this section so far share a common feature. The equilibrium strategies on and off the equilibrium path contain no delayed switches. This is the reason why in these games the outcome predictions of our model and those provided by Lipman and Wang (2000) coincide. In such cases, the shrinking future and constant switching costs framework used by them turns out to have consequences similar to our constant future and increasing switching costs. However, when the spe from either model involves delayed switches, on or off the equilibrium path, the models' predictions differ from each other. The reason for this is that in our model the payoffs are only determined by the final announcement, while in their framework players receive flow payoffs for all their decisions. Therefore, for their model but not for ours, delayed switches involve collecting the payoffs from the profile played *before* the switch.[18]

Finally, a generalization of the game of Matching Pennies is analyzed in example C.3. Here one can see how our model is able to capture the fact that sometimes commitment is disadvantageous, paralleling the notion of a second-mover advantage. In this context, the player with lower switching costs is able to stay flexible longer in the game, and therefore to best-respond to his opponent's final (committed) action. This game also provides a good example of the type of zero measure cases for which our theory does not provide a grid invariant equilibrium. When one considers the

[18]Thus, it is easy to construct examples in which, for instance, one model predicts a delayed switch on the equilibrium path and the other does not, and vice versa.

complete symmetric Matching Pennies game, the second mover advantage is determined by the choice of the grid. The symmetric Battle of the Sexes is another non-generic example. In that case, the grid determines who obtains the first mover advantage.

5.3 Bargaining

Here we present an application of our model to a family of N-player games. We are aware that our dynamic structure cannot adequately describe a "pure" bargaining model. Thinking narrowly of a series of rounds at a bargaining table, it is difficult to have a general justification for the presence of increasing switching costs. Nevertheless, there are situations in which this assumption suits better. Imagine, for instance, a political bargaining situation in which the announcements involve moving army forces, or in which changing one's mind has a reputational cost that is increasing with time.

The payoff function Π we consider is the typical "share the pie" game. There are N players and each of them has to decide how much of the pie he wants. If the final demands add up to no more than the size of the total pie, they get their demands, but if the sum is higher, no one gets anything. Specifically, given an integer $M \in \mathbb{N}$, let $A_i = \{m/M \mid m = 0, ..., M\}$ and let the payoffs be given by:

$$\Pi_i(a) = \begin{cases} a_i & \text{if} \quad \sum a_i \leq 1 \\ 0 & \text{if} \quad \sum a_i > 1 \end{cases}$$

This game, under certain conditions on the cost structure, and despite being played by more than two players, has an essentially unique spe. The final agreement is achieved immediately, avoiding any switching costs. This is so because off the equilibrium path players have credible threats to deal with the potential deviations. How much each player gets depends on his bargaining power, which comes from having relatively higher switching costs. If his cost of switching is higher, a player can commit to an action earlier in the game. This is all expressed in the following proposition (notice the normalization on the θ's).[19]

Proposition 7 *If $C_i(a_i \to a_i', t) = c_i(t)$, let t^* be the unique solution to $\sum c_i(t) = 1$ and $\theta_i = c_i(t^*)$, then the unique spe outcome of the bargaining game is to start by playing $(\theta_1, ..., \theta_N)$ and never switch in equilibrium.*

[*we provide a proof only for* $N \leq 3$]

[19] To be precise, the action space A_i has to be extended to include θ_i as well. Otherwise, the result would only be achieved in the limit, as M goes to infinity.

Let us emphasize some of the features of this result. First, the outcome prediction is intuitively appealing. With N symmetric players, each obtains an equal share of the pie and the comparative statics work in the right direction: the earlier a player is able to commit not to change his mind, the higher the stake he gets. We can interpret the θ_i's as a measure of the players' bargaining power. The higher the θ_i, the more bargaining power player i has. This result is probably more intuitive in the following example. Consider the case in which all cost functions are proportional to each other, i.e. $C_i(a_i \to a_i', t) = \beta_i c(t)$, in such a case we have $\theta_i = \frac{\beta_i}{\sum_j \beta_j}$. The proposition shows, however, that a similar result holds for a more general cost structure. This is because, somewhat surprisingly, the θ_i depends only on the value of the cost function at a particular point, t^*, but not on its shape at other points.[20] .

Muthoo (1996) presents a two-player bargaining model in which he obtains a similar result. He considers a one shot game in which each player makes a partial commitment by announcing a desired share of the pie. If he ends up receiving less, he has to pay a cost. The outcome is determined by the Nash bargaining solution. Muthoo shows that a player's bargaining power increases with the cost of revoking the initial partial commitment.

The capability of an organization to be flexible is generally considered to be a positive feature. In contrast, in our setting a non-flexible decision structure results in a higher bargaining power. This argument provides a rationale for rigid structures as bargaining devices. The difficulty of organizing a board of directors meeting, bureaucratic structures, why prices are posted, or the advantage of having a clerk with no discretion at the shop counter are only some examples.

Finally, we want to stress that we have been able to construct an N-player bargaining model that makes an essentially unique prediction. It is well known that many of the bargaining models in the literature fail this test. This could be attributed to two different reasons. The models may either become less precise in their predictions (Rubinstein's (1982) model with three players has any split of the pie as an equilibrium[21]), or the two-bargainer protocol cannot be naturally extended, running into the problem that the model provides different predictions depending on which extension of the protocol is chosen.

[20] Note, however, that the shape would affect the results of any comparative statics exercises.

[21] See Hererro (1985).

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 such as how this commitment is achieved, or which party is able to enjoy it. By relying only on fundamentals we obtain a unique equilibrium that is not driven by ad hoc assumptions regarding the order of moves. The relevant order by which parties get to commit is endogenously determined.

We present 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. The switching costs serve as a mechanism by which announcements can be made credible and commitment can be 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 more decision nodes are added and is robust to various changes in the protocol. This accomplishes the task of providing a framework which endogenously determines which player has the opportunity to commit.

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 the middle, however, commitment depends on the actions of the other player. 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, we are able to obtain entry deterrence without the need to actually fight. This is achieved by an off-equilibrium credible threat to fight in retaliation to entry. In this manner, our framework provides a unified umbrella that covers dynamic interactions that were previously captured only with different models.

We think that one of the contributions of this paper is to attract attention to the switching cost technology as the real source of commitment. In the bargaining setting, for example, our model suggests that it may be more important to focus on the costs that parties incur if they change their positions, rather than to emphasize protocol details about the order of play.

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.

Given that the structure of the equilibrium depends only on the relative order of the critical points, any transformation of the cost function would not alter this order. 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 efficient solution algorithm we provide significantly reduces the computational burden of estimating the model.

Finally, let us make two more general comments. First, in a world of imperfect information delaying an action has an option value. In our model, despite the perfect information nature of the game, delays still occur in equilibrium. Henkel (2002) and Gale (1995) obtain a similar result, although in the latter this is mainly driven by a coordination motive. In our context, the nature of the delays is completely strategic. They are costly, but allow players to make threats credible. An interesting extension would be to introduce incomplete information in our current framework and analyze how this affects the strategic delays.

Second, the model presents a case in which the equilibrium is robust to details about its protocol. Not only is the unique equilibrium invariant to the protocol details within the universe of sequential games, it is also robust to other changes in the protocol. Such changes allow for the introduction of uncertainty about the precise times in which players get to play, and for the introduction of periods of simultaneous play. As we have pointed out, the key is to have some amount of asynchronicity in order to allow players to unravel long strings of bad coordination and to restore uniqueness.

The protocol invariant structure of the subgame perfect equilibrium is an appealing property. We are currently studying other dynamic frameworks that may share this equilibrium feature. In particular, we are interested in modifying the current setup to allow for players to build up stocks of strategy-specific investment. In our current framework past switching cost are sunk, past announcements are "forgotten", so history does not matter. One may want to consider the possibility that past actions have an impact. Think, for example, on a firm announcing its intention to enter a new market by acquiring some industry-specific human capital. If it suddenly decides not to enter, this human capital cannot simply be erased. Future work should also aim at generalizing the model along its different dimensions. In particular, we are studying how the complexity of the stage structure changes by increasing either the number of players or the action spaces. We are also trying to characterize the families of N-player games which are robust to changes in the protocol.

References

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

[2] Einav, L. (2001), "Seasonality and Competition in Time: An Empirical Analysis of Release Date Decisions in the U.S. Motion Picture Industry," mimeo, Harvard University.

[3] Gale, D. (1995), "Dynamic Coordination Games," *Economic Theory*, Vol. 5, pp. 1-18.

[4] Gale, D. (2001), "Monotone Games with Positive Spillovers," *Games and Economic Behavior*, Vol. 37, pp. 295-320.

[5] Henkel, J., "The 1.5th Mover Advantage," *RAND Journal of Economics*, Vol. 33(1), pp. 156-170.

[6] Herrero, M. (1985), "A Strategic Bargaining Approach to Market Institutions," Ph.D. Thesis, London School of Economics.

[7] Kreps, D., and R. Wilson (1982), "Reputation and Imperfect Information," *Journal of Economic Theory*, Vol. 27, pp. 253-279.

[8] Lagunoff, R., and A. Matsui (1997), "Asynchronous Choice in Repeated Coordination Games," *Econometrica,* Vol. 65, pp. 1467-1477.

[9] Lipman, B. L, and R. Wang (2000), "Switching Costs in Frequently Repeated Games," *Journal of Economic Theory,* Vol. 93 (2), pp. 149-190.

[10] Milgrom, P., and J. Roberts (1982), "Predation, Reputation and Entry Deterrence," *Journal of Economic Theory*, Vol. 27, pp. 280–312.

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

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

[13] Rubinstein, A. (1982), "Perfect Equilibrium in a Bargaining Model," *Econometrica*, Vol. 50, pp. 97-110.

[14] Van Damme, E., and S. Hurkens (1996), "Commitment Robust of Equilibria and Endogenous Timing," *Games and Economic Behavior*, Vol. 15(2), pp. 290-311.

Appendix

A Proofs

Proof of Lemma 1

Let us state and prove a generalized version of Lemma **1:**

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

Let $next(t) = \underset{i}{Min}\{next_i(t)\}$. By $s_i(a,t) = a'_i$ we know that

$$V_i((a'_i, a_{-i}), next(t)) - C_i(a_i \to a'_i, t) \geq V_i((a''_i, a_{-i}), next(t)) - C_i(a_i \to a''_i, t) \ \forall a''_i. \tag{2}$$

By the triangle inequality we know that

$$C_i(a_i \to a'_i, t) + C_i(a'_i \to a''_i, t) \geq C_i(a_i \to a''_i, t).$$

Adding both expressions we get

$$V_i((a'_i, a_{-i}), next(t)) \geq V_i((a''_i, a_{-i}), next(t)) - C_i(a'_i \to a''_i, t) \ \forall a''_i$$

so that $s_i((a'_i, a_{-i}), t) = a'_i$ and the first part of the statement is true.

The second statement is already true for a_i and for a'_i. For $a''_i \neq a_i, a'_i$ we just need to add (2) with $a''_i = a_i$, and to the following triangle inequality

$$C_i(a''_i \to a_i, t) + C_i(a_i \to a'_i, t) > C_i(a''_i \to a'_i, t).$$

We then get

$$V_i((a'_i, a_{-i}), next(t)) - C_i(a''_i \to a'_i, t) > V_i((a_i, a_{-i}), next(t)) - C_i(a''_i \to a_i, t)$$

which implies that a_i is dominated by a'_i , i.e. $s_i((a''_i, a_{-i}), t) \neq a_i$.

Proof of Proposition 1

We prove the first part of the proposition for general K and N**:**

- If there exists a player i and a point in time $t \in g_i$ such that $s_i((a_i, a_{-i}), t) = s_i((a_i', a_{-i}), t)$ $\forall a_i, a_i' \in A_i$, $a_{-i} \in A_{-i}$, then for all players $p = 1, 2, ..., N$ and for any $t' \in g_p$ such that $t' \leq t$, the strategies are independent of a_i, i.e. $s_p((a_{-i}, a_i), t') = s_p((a_{-i}, a_i'), t')$ $\forall a_i, a_i' \in A_i$, $a_{-i} \in A_{-i}$.

We prove this by induction on the level of the game tree, starting at t and going backwards. By the assumption of the proposition, we know that at time t the proposition holds, i.e. that $s_i((a_i, a_{-i}), t) = s_i((a_i', a_{-i}), t)$ $\forall a_i, a_i' \in A_i$. Denote these strategies by $s_i(a_{-i}, t)$. Now, suppose the proposition holds for time $t' \leq t$, so that we have to show that it holds also for the previous decision node in the game, i.e. for $t'' = \underset{k}{Max}\{prev_k(t')\}$. Let player j be the player who plays at time t'', i.e. $t'' \in g_j$. We check two cases, the first is when $j \neq i$ and the second is when $j = i$.

If $j \neq i$ consider player j's continuation values just after his move at t'', $V_j(a, next(t''))$, and player j's continuation values just before player i moves at t, $V_j(a, t)$. Using the proposition assumption, we know that $V_j(a, t) = V_j((s_i(a_{-i}, t), a_{-i}), next(t))$, which is independent of a_i, and using the induction assumption we know that all strategies of all players between t'' and t are independent of a_i, thus it is clear that $V_j(a, next(t''))$ is also independent of a_i, which immediately implies that $s_j(a, t'')$ is independent of a_i.

If $j = i$, consider player i's (j's) continuation values before his move at time t, $V_i(a, t)$. Using the proposition assumption, we know that $V_i(a, t) = V_i((s_i(a_{-i}, t), a_{-i}), t) - C_i(a_i \rightarrow s_i(a_{-i}, t), t)$. Now, observe that by the induction assumption we know that from time t'' until time t, no player's strategy depends on a_i. In particular, the actions of the rest of the players evolve independently of a_i, so that at time t, they are $a_{-i}^t(a_{-i}^{t''})$. Thus, player i's strategy at t'' is reduced to an individual decision problem, in which he tries to save on cost, because his actions has no strategic effect. It is easy to verify that following the properties of the switching costs function (in particular, the monotonicity and the triangle inequality) player i's would best mimic his time t strategy with respect to the play of the other players, i.e. $s_i(a, t'') = s_i(a_{-i}^t(a_{-i}), t)$, which is independent of a_i. This concludes the first part of the proposition.

Now we prove for $N = K = 2$ that the game has at most three stages: Denote by t^* the last point in time at which the hypothesis of the proposition is satisfied. Note that t^* is a critical point. Given that player j's continuation values at t^* depend only on a_j, we can denote them $V_j(a_j, t^*)$. Player j can obtain the outcome that is more favorable for him just by playing $a_j^* = \underset{a_j}{ArgMax} V_j(a_j, t^*)$ at his first grid point and not switching ever until t^*. At any profile a with

$a_j = \tilde{a}_j^*$, player j switches immediately to a_j^* if and only if $V_j(a_j^*, t^*) > V_j(\tilde{a}_j^*, t^*) - C_j(\tilde{a}_j^* \rightarrow a_j^*, t)$. 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, up to t^{**}, play begins by both players playing $(s_i(a_j^*, t^*), a_j^*)$ 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 exit. It appears only when player i needs to re-adjust to the expected move at t^{**} by player j. This happens when $s_i(a_j^*,t^*) \neq s_i(\tilde{a}_j^*,t^*)$.

Proof of Lemma 2

We just need to prove that there are no two consecutive decision nodes for player i, $t, next_i(t) \in g_i$, within a stage satisfying $s_i((a_i, a_{\tilde{i}}), t) = a_i$ and $s_i((a_i, a_{\tilde{i}}), next_i(t)) = \tilde{a}_i$. We prove it by contradiction. If player $\tilde{i}$ were not playing between t and $next_i(t)$ the contradiction is immediate. Consider now the case in which he does. Without any loss of generality let us suppose that he plays only once in between and does it at t'. We consider different cases depending on what player $\tilde{i}$ does at profiles $((a_i, a_{\tilde{i}}), t')$ and $((\tilde{a}_i, a_{\tilde{i}}), t')$:

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

2. If $s_{\tilde{i}}((a_i, a_{\tilde{i}}), t') = a_{\tilde{i}}$ and $s_{\tilde{i}}((\tilde{a}_i, a_{\tilde{i}}), t') = \tilde{a}_{\tilde{i}}$, and given that t' is not the end of a stage, we know that $s_{\tilde{i}}((\tilde{a}_i, a_{\tilde{i}}), next_{\tilde{i}}(t')) = \tilde{a}_{\tilde{i}}$. This implies that the equilibrium path starting at (a, t) takes us to $((\tilde{a}_i, \tilde{a}_{\tilde{i}}), next_i(next_{\tilde{i}}(t')))$. Now, one can check that player i can get there at a lower cost by deviating and playing $\tilde{a}_i$ at (a, t) and not changing at $next_i(t)$. This provides the contradiction.

3. If $s_{\tilde{i}}((a_i, a_{\tilde{i}}), t') = \tilde{a}_{\tilde{i}}$ and $s_{\tilde{i}}((\tilde{a}_i, a_{\tilde{i}}), t') = \tilde{a}_{\tilde{i}}$, and given that t' is not the end of a stage, we know that $s_{\tilde{i}}((a_i, a_{\tilde{i}}), next_{\tilde{i}}(t')) = s_{\tilde{i}}((\tilde{a}_i, a_{\tilde{i}}), next_{\tilde{i}}(t')) = \tilde{a}_{\tilde{i}}$. Using Proposition 1 we can conclude that $s_i((a_i, a_{\tilde{i}}), next_i(t)) = s_i((a_i, \tilde{a}_{\tilde{i}}), next_i(t)) = \tilde{a}_i$. Now it is easy to see that player i can improve by deviating at (a, t) and playing $\tilde{a}_i$. Again, this leads to a contradiction.

4. Finally, if $s_{\tilde{i}}((a_i, a_{\tilde{i}}), t') = \tilde{a}_{\tilde{i}}$ and $s_{\tilde{i}}((\tilde{a}_i, a_{\tilde{i}}), t') = a_{\tilde{i}}$ and given that t' is not the end of a stage, we know that $s_{\tilde{i}}((a_i, a_{\tilde{i}}), next_{\tilde{i}}(t')) = \tilde{a}_{\tilde{i}}$. Now, we need to consider what

player i does at $((a_i, \tilde{\ } a_{\tilde{\ } i}), next_i(t))$. If $s_i((a_i, \tilde{\ } a_{\tilde{\ } i}), next_i(t)) = a_i$, one can check that player i can benefit from playing $\tilde{\ } a_i$ at (a, t), providing a contradiction. If $s_i((a_i, \tilde{\ } a_{\tilde{\ } i}), next_i(t)) = \tilde{\ } a_i = s_i((a_i, a_{\tilde{\ } i}), next_i(t))$, then we can use Proposition 1 to conclude that $s_{\tilde{\ } i}((a_i, a_{\tilde{\ } i}), t') = s_{\tilde{\ } i}((\tilde{\ } a_i, a_{\tilde{\ } i}), t')$, which leads a contradiction.

Proof of Proposition 2

The proof makes extensive use of the Solution Algorithm (Appendix B.1). The proof is generic 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 $(t_m^*, S_g(p, a, m), V_m, AM_m)_{m=0}^{\overline{m}}$. All we have to show is that the Solution Algorithm replicates the backward induction, that is, that

$$\widetilde{s}_p(a, t) = S_g(p, a, \widetilde{m}(t)) \text{ where } \widetilde{m}(t) = \{m | t \in (t_{m+1}^*, t_m^*]\}$$

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 p at node (a, t)

$$\widetilde{V}_p(a, t) \equiv \begin{cases} V^{new}(V_{\widetilde{m}(t)-1}, AM_{\widetilde{m}(t)-1}, t, p) \text{ evaluated at } (a, p) \text{ if } t \in g_p \\ V^{new}(V_{\widetilde{m}(t)-1}, AM_{\widetilde{m}(t)-1}, t, \tilde{\ } p) \text{ evaluated at } (a, p) \text{ if } t \notin g_p \end{cases} \tag{3}$$

where $V^{new}(V, AM, t, p)$ is defined in the Solution Algorithm description in Appendix B.1, part 2.d .

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[22] , the costs at the final decision node of the game are so high that switching is too costly, so that no player switches. In other words, no player has an active move, which is exactly what the algorithm does when it initializes and sets $AM_0(a, p) = 0$ for all a, p. This implies that $\widetilde{s}_p(a, T) = a_p$, which is what we wanted to prove. Now it can easily be verified that $\widetilde{V}_p(a, T)$ captures the continuation values. This finishes the proof for the final decision node.

Now, suppose the statement is true for $next(t) \equiv Min\{next_1(t), next_2(t)\}$, and we will show it is true for t. Fix a profile a. As before, once we have proven that the proposed strategy $\widetilde{s}_p(a, t)$

[22] If the grid is not fine enough, we may have an active move at the last point on the grid, $\underset{i}{Max}\{prev_i(T)\}$, which would imply also that $t_1^* = \underset{i}{Max}\{prev_i(T)\}$, and that the algorithm would update this properly. the proof can be easily extended for such cases as well, but given that our main focus is on grids that are sufficiently fine, we ignore such cases for the rest of the proof.

is correct, verifying the update of the continuation values for both players is an immediate check. To finish, all that remains to be checked is that playing according to $\widetilde{s}_p(a,t)$ at profile a is indeed optimal. Because of the induction hypothesis we know that $\widetilde{V}_p(a, next(t))$ are the continuation values of the game. Therefore the spe strategy, denoted by $a_p^*(a,t)$, can be written as a function of $\widetilde{V}$, and is the solution to

$$a_p^*(a,t) = \underset{a'_p \in A_p}{ArgMax}\{\widetilde{V}_p((a'_p, a_{\sim p}), next(t)) - C_p(a_p \to a'_p, t)\}$$

The next step is to realize that proving that $\widetilde{s}_p(a,t) = a^*(a,t)$ is equivalent to proving the following

$$AM_{\widetilde{m}(t)}(a,p) = 1 \Longleftrightarrow \widetilde{V}_p(({}^{\sim}a_p, a_{\sim p}), next(t)) - C_p(a_p \to {}^{\sim}a_p, t) > \widetilde{V}_p(a, next(t)) \qquad (4)$$

The advantage to rewrite it in this way is that equation (4) only involves functions defined in the algorithm. Therefore, the problem is reduced to an algebra check that does not deal with the game *per se*, but only with variables defined in the Solution Algorithm. The check is simple but tedious, because 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. In general, we need to check thirty two different cases. This is simplified by the fact that many of the cases can be ruled out as impossible, and that some other can easily be grouped and checked together.

It is important to notice that the different cases correspond to the different situations that a player might encounter in a game. For instance, whether he is in a middle of a stage or not, whether the other player has an active move or not, etc.

Including a full check for all the cases in this Appendix would simply be too long, and will not provide much intuition. Still, we present one simple case so the reader can visualize how easy each check is. Consider a point $t \in g_p$ in the middle of a stage. Suppose that only player p 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 $AM_{\widetilde{m}(t)}(a,p) = 1$. In this case, applying Definition 3, we have that

$$\widetilde{V}_p(a, next(t)) = V_{\widetilde{m}(next(t))}(({}^{\sim}a_p, a_{\sim p}), p) - C_p(a_p \to {}^{\sim}a_p, next_p(t))$$
$$\widetilde{V}_p(({}^{\sim}a_p, a_{\sim p}), next(t)) = V_{\widetilde{m}(next(t))}(({}^{\sim}a_p, a_{\sim p}), p)$$

so it is easy to see that (4) is true because of the monotonicity of the cost function, which implies that $C_p(a_p \to {}^{\sim}a_p, next_p(t)) > C_p(a_p \to {}^{\sim}a_p, t)$.

Moreover, once we know that the Solution Algorithm solves for the unique spe of the game, then, as a direct application of Remark 2, we get that the equilibrium has no more than eight stages.

Proof of Proposition 3

Given a specific node (a_0, t_0), consider the subgame that starts at it. Assume there is a delayed move at (a_1, t_1) and wlog., assume it is done by player 2, so that $t_0 \in g_2$. The proof proceeds in two parts. First, we show that on the equilibrium path of this subgame player 1 will never switch thereafter. Then, we show that player 2 will not switch after t_1 either.

To prove the first part, suppose towards contradiction that player 1 switches after t_1. Without loss of generality, assume that the last delayed switch by player 2 is at $((D, R), t_1)$ from R to L for $t_1 > t_0$, after which player 1 moves at $((D, L), t_2)$ from D to U for $t_2 > t_1$. By Lemma 1, player 2 sticks to U at $((U, L), t_2)$, i.e. player 1 plays U at t_2, as long as player 2 is at L. This means that at (a, t_0) player 2 has a profitable deviation–play L always. This gives him the same outcome, with lower (or zero) switching costs.

To prove the second part we will also work by contradiction. Without loss of generality, assume that on the equilibrium path, at $((U, R), next_2(t_0))$ player 2 moves from R to L (after which player 1 does not switch on the equilibrium path, as implied by the first part of the proof), and that at $((U, L), \bar{t})$ player 2 moves from L to R, closing the circle.

Denote the possible continuation values for player 1 at $\bar{t}$ by $A \equiv V_1((D, R), \bar{t})$, $B \equiv V_1((D, L), \bar{t})$, $C \equiv V_1((U, R), \bar{t}) = V_1((U, L), \bar{t})$.

Observe that the delayed switch of player 2 at $\bar{t}$ must mean that player 1 moves from U to D at $((U, R), \bar{t})$, implying $A - C_1(U \to D, \bar{t}) > C$.

We need to consider 2 subcases:

1. $a = (U, R)$ (which, using Lemma 1, also covers the case of $a = (U, L)$), so that player 2 plays R at $((U, R), t_0)$ and player 1 plays U at $((U, R), next_1(t_0))$. This means that player 1 plays D at $((U, L), next_1(t_0))$, otherwise there would not have been any reason for player 2 to delay the switch to L at $((U, R), t_0)$. Observe also that for $t0 < t \leq \bar{t}$ at $((D, L), t)$ player 1 always sticks to D, otherwise player 2 could play L at $((U, R), t_0)$ instead of delaying. Denote by t' the first time, if any, that player 2 plays R at $((D, L), t')$ for $t0 < t' \leq \bar{t}$. Construct the following deviation for player 1–play D at $((U, R), next_1(t_0))$ and stick to D at any $t0 < t \leq t'$. In any other decision node, do the same as before. This strategy would

either yield payoffs of $A - C_1(U \rightarrow D, next_1(t_0))$ if t' does not exist, or payoffs which are equal to $V_1((D,L), next_1(t_0)) - C_1(U \rightarrow D, next_1(t_0))$. Both these expression are more than the payoffs of C that player 1 obtains by playing U at $((U,R), next_1(t_0))$. Thus, we obtained a contradiction.

2. $a = (D,R)$ (which, using Lemma 1, also covers the case of $a = (D,L)$), so that player 2 plays R at $((D,R),t_0)$ and player 1 plays U at $((D,R), next_1(t_0))$. This means that player 1 plays D at $((D,L), next_1(t_0))$, otherwise there would not have been any reason for player 2 to delay the move to L at $((D,R),t_0)$. As in the previous case, it must be that for $t0 < t \leq T$ at $((D,L),t)$ player 1 always sticks to D, otherwise player 2 could play L at $((D,R),t_0)$ instead of delaying. Denote by t' the first time, if any, that player 2 plays R at $((D,L),t')$ for $t0 < t' \leq \bar{t}$. Construct the following deviation for player 1–play D at $((D,R), next_1(t_0))$ and stick to D at any $t0 < t \leq t'$. In any other decision node, do the same as before. This strategy would either yield payoffs of A if t' does not exist, or payoffs which are equal to $V_1((D,L), next_1(t_0))$. Both these expression are more than the payoffs of $C - C_1(D \rightarrow U, next_1(t_0))$ that player 1 gets by playing U at $((D,R), next_1(t_0))$, giving us a contradiction.

Proof of Proposition 4

A delayed switch (in the equilibrium of any subgame) is uniquely defined by the decision node at which it is played (a,t). Consider the first delayed switch, i.e. a delayed switch (a,t) such that there does not exist any delayed switch (a',t') with $t' < t$. Without loss of generality, let this delayed switch be carried by player 2 at $((U,L),t_0)$. By Proposition 3 we know that this implies that both players' equilibrium strategies at $((U,R),t)$ are not to switch, for $t > t_0$. Clearly, this immediately implies that there are no delayed switches of the form $((U,R),t)$ by any of the players. In addition, it is clear that a real delayed switch into (U,R) at t implies an off-equilibrium switch out of (U,R) at $prev_j(t)$, where j is the other player (otherwise, there was no reason to delay the switch to begin with). Given that there are no moves out of (U,R), we know that there cannot be any delayed switches into (U,R), i.e. $((U,L),t)$ by player 2 and $((D,R),t')$ by player 1, for $t,t' > t_0$. In addition, we can rule out a real delayed switch $((D,L),t)$ by player 2–such delayed switch implies a switch by player 1 from D to U at $((D,R), prev_1(t))$. If this is the case, however, there is no motivation for the original delayed switch at $((U,L),t_0)$ by player 2. Player 2 could have just played R to begin with, knowing that player 1 will eventually move to U, thus saving on the costs of the delay. Finally, we can rule out a real delayed switch $((D, R), t)$ by player 2: denote the possible terminal values for player 1 at T by $A \equiv V_1((D, R), T), B \equiv V_1((D, L), T), C \equiv V_1((U, R), T)$. The original delayed switch at $((U, L), t_0)$ by player 2 implies that player 1 moves to D at $((U, R), prev_1(t))$. It is easy to see that at $((D, R), t_0)$ player 2 does not move. By Proposition 3 we know that on the equilibrium path that goes down from $((D, R), t_0)$ there exists at most one delayed switch. It cannot be a switch by player 1, which would have lead him back to (U, R), giving him a continuation value of C, that he could have obtained without any cost by not moving at $((U, R), t_0)$. By the second delayed switch, $((D, R), t)$ by player 2, we know that payoffs of A are not achievable. Thus, we know that the payoffs for player 1 from moving to D at $((U, R), t_0)$ are B, so that $B - C_1(U \to D, next_1(t_0)) > C$. However, it is easy to see that by player 1 moving at $((U, L), t_0)$ and never moving thereafter, he could have also obtained payoffs of $B - C_1(U \to D, next_1(t_0))$, which is better than what he gets in equilibrium, C. Thus we obtain a contradiction.

Hence, potentially there could be only two real delayed switches after t_0–$((U, L), t)$, $((D, L), t)$– both by player 1. We can now analogously apply the first part of the proof in order to obtain that after the earliest of these two potential real delayed moves, only player 2 can potentially have a real delayed switch. We already ruled out, however, any additional real delayed switches by player 2, so there cannot be any more delayed switches in the game, rather than the original one (by player 2) and at most one of these two (by player 1).

Proof of Theorem 1

To prove this, we 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

$$\lim_{\varphi(g) \to 0} S_g(p, a, m) = S(p, a, m)$$

This will be achieved by using another algorithm, which we call the Limit Algorithm (see B.2), that computes the limit of the equilibria.

First, note that the generic statement makes sure that we only consider the cases for which the Limit Algorithm is well defined. Thus, we disregard the cases in which the critical points are the same for different players.

We will prove it recursively on the stages of the algorithm. For any given m we will actually prove the convergence of all the functions used in the Solution Algorithm $(t, a^*, p^*, t^*_m, AM_m, V_m, FS_m)$ to their counterparts in the Limit Algorithm. This task has to be done in the same order in which the algorithm proceeds. Given a particular function in it, all that needs to be taken into account is that this function is piecewise defined by continuous transformations of:

(a)other functions for which the convergence has already been checked (because of the recursive procedure).

(b)the cost function, which is continuous.

The last check is to make sure that in the piecewise defined functions the cutoff points (the points at which the definition changes) also converge and create no problem in the boundary. This is the case because the mutually exclusive conditions that define the cutoff points are either functions with a finite range (and for which the recursive procedure applies), or, in the case of $t(a, p)$, defined using ΔV. Note that, in this latter case, there is no discontinuity of $t(a, p)$ at $\Delta V = 0$, so the convergence check is sufficient.

This essentially finishes the proof of the Order Invariance Theorem. The choice of α is an immediate consequence of the fact that the range of $S_g(p, a, m)$ is a finite set.

Proof of Proposition 5

We prove the proposition of a general action space, for $N = 2$. Generalizing the defendability definition and the proof for $N > 2$ is almost straight forward.

Let $t^*_i(a_j) = Min\,\{t \in g_i | C_i(a^*_i \to a'_i, t) > \Pi_i(a^*) - \Pi_i(a^*_i, a_j)\ \forall a'_i \in A_i\}$. Clearly, player i has no incentive to switch away from (a^*_i, a_j) at any decision node after $t^*_i(a_j)$–the best he can ever hope to gain is less than the required switching costs. Thus, the subgame $((a^*_i, a_j), t)$ reduces to an individual decision making by player j. Observe also that the defendability definition implies also that a^*_j is the best response of player j to a^*_i. Provided that the grid is fine enough, and given that a^* is defendable, we know that there exists a point in time $t' = next_j(t^*_i(a_j)) \in g_j$ such that $\Pi_j(a^*) - C_j(a_j \to a^*_j, t) > \Pi_j(a^*_i, a_j)$. Therefore, in spe player j would switch to a^*_j at decision node $((a^*_i, a_j), t')$. Given that this is true for any a_j, it will never pay for player j to switch to some other action with the hope that player i will switch away from a^*_i. Player i, knowing it in advance, will start by playing a^*_i, and will never move thereafter, knowing that eventually player j will switch to a^*_j, giving player i his maximum possible payoffs. On equilibrium, player j realizes this and starts by playing a^*_j from the beginning.

Proof of Proposition 6

If $D - d > a$ the equilibrium is *[No Fight, Entry]*. To see this, consider the strategy of the entrant to start by entering and never to exit, at any decision node at which he is in. Clearly, the best response by the incumbent to such a strategy is not to fight, so that equilibrium payoffs are (D, b). All we need to show is that this strategy for the entrant is subgame perfect. To see this, consider first all decision nodes for which $t > c^{-1}(D-d) > c^{-1}(a)$, in which neither of the players want to pay the switching cost and to move away from either (Fight, Enter) or from (Don't Fight, Enter). Consider now the last decision node for the incumbent before $t = c^{-1}(D-d)$, in which it still profitable for the incumbent to stop fighting if he fought before, so he would switch and stop fighting, leading to an equilibrium payoffs of $(D - c(t), b)$. The same is true for any time between $c^{-1}(a)$ to $c^{-1}(D-d)$. Thus, even before $c^{-1}(a)$ the entrant knows that by sticking to entry he can achieve his maximum payoffs, b, and hence it is credible for him to do that.

The rest of the proof will consider the case $D - d < a$, and mechanically will go over all the 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 the one associated with maximum of a, b or $M - m$. We analyze each case in turn:

1. b is the maximum, so that at $t = c^{-1}(b)$ the action "Enter if the incumbent is not fighting" becomes active. Thus, in the last stage of the game (between $c^{-1}(b)$ to T) nothing happens, and this action is the only active one in the stage before the last one. Now the new differences to be considered are a and $m - D$. The next action that becomes active is the one associated with the maximum between these two.

    (a) If a is the maximum, then in the next stage the entrant activates his other action (to exit if the incumbent fights), and in the termination stage the incumbent ignores the entrant's actions 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 outcome of *[Fight, No Entry]*. This case is analogous to the two-period sequential game, in which the incumbent enjoys first mover advantage–late in the game he is fully committed, while the entrant remains fully flexible.

    (b) If $m - D$ is the maximum, then in the next stage the incumbent activates a switch from not fighting to 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, Enter), which, given our assumption that $a > D - d$, is won by the incumbent. Therefore the equilibrium is *[Fight, No Entry].*

2. $M - m$ is the maximum, so that at $t = c^{-1}(M - m)$ the action "Do not fight if the entrant is out" becomes active. Thus, in the last stage of the game (between $c^{-1}(M - m)$ to T) nothing happens, and this action is the only active one in the stage before. The new differences to be considered are a and b. The next action that becomes active is the one associated with the maximum of the two.

   (a) If a is the maximum, we get into a stage in which each player has one move active. The next switch that becomes active will determine the outcome of the game. This depends on the comparison between b and $(M - D)/2$. b is the benefit for the entrant if he decides to enter when the incumbent is not fighting–if b is greater then the entrant knows nothing would happen after, so he would gain b for lower costs. To understand why $(M - D)/2$ is the relevant number for the incumbent, we need to consider what happen next. If the entrant is in and the incumbent decides to fight, it will prompt a chain reaction, in which the entrant will exit, after which the incumbent will stop fighting. These three actions will not be delayed. The incumbent, therefore, understands that in order to obtain M , he will have to switch first to fighting, and then switch again to not fighting. These will happen almost immediately (recall that $\varphi(g) \to 0$). Thus, initiating this sequence of moves is beneficial as long as $2c(t) = M - D$, or that $t < c^{-1}((M - D)/2)$.

      i. If $(M - D)/2 > b$ then the equilibrium is *[No Fight, No Entry].* This is obtained through a complicated threat by the incumbent. The threat can be read as "I do not fight today, but in case I see that you decide to enter, I will write the necessary contracts that will allow me to fight you once you enter. This will make you reverse your decision, so you should not even try". Thus, the entry deterrence is achieved through a credible threat to fight in case the entrant decides to enter.

      ii. If $b > (M - D)/2$ then the incumbent is free to best reply to the entrant's strategy at $t = c^{-1}(b)$. Note, however, that his best reply can include now a delayed switch. The incumbent can accommodate, let the entrant enter, and obtain C. Alternatively, the incumbent can fight, make the entrant exit, wait until $t = c^{-1}(b)$, and then stop fighting, knowing that after $t = c^{-1}(b)$, it not beneficial for the entrant to enter anymore. Thus, the two possible equilibrium payoffs for the incumbent are D, if he accommodates, and $M - c(c^{-1}(b)) = M - b$ if he fights, and stops fighting only after $t = c^{-1}(b)$. Given that we assumed that $M - b > D$ the equilibrium is to start with *[Fight, No Entry]* and to switch to *[No Fight, No Entry]* at $t = c^{-1}(b)$.

(b) If b is the maximum, the incumbent's switch to not fighting becomes inactive at $t \leq c^{-1}(b)$. This is because the switch to not fighting at such times will make the entrant enter, denying the incumbent from achieving his big prize of M. Still, the incumbent can potentially achieve M by fighting until $t = c^{-1}(b)$, and switching to not fighting only later on. This will give him payoffs of $M - c(c^{-1}(b)) = M - b$. Thus, the new differences that need to be compared are a and $(M - b) - D$.

i. If a is the maximum, then the incumbent best replies to the entrant's strategy at $t = c^{-1}(a)$, and just as in the last case, given that $M - b > D$, the equilibrium is to start with *[Fight, No Entry]* and to switch to *[No Fight, No Entry]* at $t = c^{-1}(b)$.

ii. If $(M - b) - D$ is the maximum, then by $M - b > D$ and $a > D - d$, the incumbent best replies to the entrant's strategy at $t = c^{-1}(a)$, and chooses to fight in the beginning and then to stop fighting at $t = c^{-1}(b)$, so the equilibrium is to start with *[Fight, No Entry]* and to switch to *[No Fight, No Entry]* at $t = c^{-1}(b)$.

3. a is the maximum, so that at $t = c^{-1}(a)$ the action "Exit if the incumbent is fighting" becomes active. Thus, in the last stage of the game (between $c^{-1}(a)$ to T) nothing happens, and this action is the only active one in the stage before the last one. We now need to compare between $M - m$, $m - D$, and b.

(a) If b is the maximum, then the outcome is *[Fight, No Entry]*–the incumbent best replies to the entrant strategy at $t = c^{-1}(b)$, so he chooses to fight.

(b) If $m - D$ is the maximum, then the equilibrium is *[No Fight, No Entry]*. This is a case in which the incumbent can credibly commit to fight an entrant through the fact that his fighting mechanism is very effective–it hearts the entrant a lot (a is large), but is not very costly (m is large).

(c) If $M - m$ is the maximum, then we are back to case 2a above, where the equilibrium is the following. If $(M-D)/2 > b$ then the incumbent can credibly threat to fight if the entrant enters, thus the equilibrium is *[No Fight, No Entry]*. If, however, $b > (M-D)/2$ then the equilibrium is to start with *[Fight, No Entry]* and to switch to *[No Fight, No Entry]* at $t = c^{-1}(b)$.

Proof of Proposition 7

Strategy of proof: Denote the (yet unknown) spe strategy for player i by $s_i((a_1, ..., a_N), t)$. In Lemma 8 we will argue that given all other players spe strategies player i can get a payoff of θ_i by playing θ_i throughout the game. While this is not player i's spe strategy, it implies that player i gains at least θ_i by playing his spe strategy (otherwise there would be a profitable deviation from the suggested equilibrium). Obtaining this for all players implies that in equilibrium there are no delayed switches (no switching costs are paid, otherwise the sum of the payoffs cannot add up to one), and that each player i announces θ_i throughout (that's the only way each player can obtain at least his θ_i).

Lemma 3 *Consider a subgame $((a_1, ..., a_N), t)$ such that $t > t^*$ and $a_j \geq \theta_j$ $\forall j \neq i$. Then in the equilibrium path of such a subgame player i does not switch.*

Proof Denote by a_i^* the final asking of player i in the equilibrium outcome of this subgame. Suppose, towards contradiction, that player i switches. This implies that $a_i^* > \theta_i$ and hence that $\sum_{j \neq i} a_j^* < \sum_{j \neq i} \theta_j$ (otherwise there is no agreement), which implies that there exists a player $k \neq i$ for whom $a_k^* < \theta_k$, which implies that player k switched at least once on the equilibrium path (recall: $a_k \geq \theta_k$), paying switching costs which are greater than θ_k, and hence obtained negative payoffs at this subgame, compared to at least zero he could have obtained by not switching at all. Therefore, we have a contradiction.

Lemma 4 *Consider a subgame $((a_1, ..., a_N), t)$ such that $t > t^*$ and $a_k \geq \theta_k$ $\forall k \neq i, j$. In the equilibrium path of such a subgame player either i or j (or both) does not switch.*

Proof Denote by a_i^* the final share of player i in the equilibrium outcome of this subgame. Suppose, towards contradiction, that both player i and j switch. This implies that $a_i^* > \theta_i$ and $a_j^* > \theta_j$ and hence that $\sum_{k \neq i,j} a_k^* < \sum_{k \neq i,j} \theta_k$, which implies that there exists a player $l \neq i, j$ for whom $a_l^* < \theta_l$, which implies that player l switched at least once on the equilibrium path (recall: $a_l \geq \theta_l$), paid switching costs which are greater than θ_l, and hence obtained negative payoffs at this subgame, compared to at least zero he could have obtained by not switching at all. Therefore, we have a contradiction.

Lemma 5 *Consider a subgame* $((a_1, ..., a_N), t)$ *such that* $\max_i c_i^{-1}(\theta_i - \frac{1}{L}) < t < t^*$ *and* $a_j \geq \theta_j$ $\forall j \neq i$. *If player* i *switches in the equilibrium path of such a subgame, he switches to* θ_i.

Proof Denote by a_i^* the final share of player i in the equilibrium outcome of this subgame. Suppose, towards contradiction, that player i switches to something different from θ_i. This implies that $a_i^* > \theta_i$ (for L large enough) and hence that $\sum_{j \neq i} a_j^* < \sum_{j \neq i} \theta_j$, so we can continue just as in Lemma 3.

Remark 1 *For two players, the proof is now easy: player* j, *at his last turn before* t^* *and against* θ_i, *knows that by playing* $a_j > \theta_j$ *he will get nothing in the end (follows directly from the Lemmas 1 and 3 above, for any order of play), while by playing* θ_j *he would get positive payoffs. Therefore, player* j *will always prefer to play* $a_j \leq \theta_j$. *We know from Lemma 1 that, as long as player* i *plays* θ_i, *player* j *will not switch, thus guaranteeing an agreement, and payoffs of* θ_i *to player* i.

Lemma 6 *Given* (a, t) *then the final actions taken* a^* *are are such that* $a_i^* \geq \min(a_i, c_i(next_i(t)))$. *Moreover, if* a *is compatible or at least one payer switches, then* $\sum_{i=1}^{N} a_i^* \leq 1$.

Proof Starting with the second part, note that if a player switched then it can be rationalized only if there is an agreement in the end (otherwise it would give him negative payoffs compared to zero he could have obtained by staying put). Moreover, if a is compatible then we get an agreement even if no player switches. The first part is given by the fact that if a player switches then he must obtain at least his switching costs, otherwise he would have been better off not switching.

Corollary 2 *The continuation value for player* i *is bounded between* 0 *and*

$$max \left\{ 0, 1 - \sum_{j \neq i} \min(a_j, c_j(next_i(t))) \right\}.$$

Corollary 3 *The final asking for player* i *is bounded between* $\min(a_i, c_i(next_i(t)))$ *and*

$$max \left\{ \min(a_i, c_i(next_i(t))), 1 - \sum_{j \neq i} \min(a_j, c_j(next_i(t))) \right\}.$$

Lemma 7 *For three players, consider a compatible state* (θ_1, a_{-1}) *at* $t > t^*$, *i.e.* $\theta_1 + a_2 + a_3 \leq 1$. *Then "always* θ_1*" guarantees* θ_1 *(agreement).*

Proof Suppose not. Let (θ_1, a_{-1}) at $t > t^*$ stand for the latest such a subgame in which it does not hold. Look at the equilibrium path out of (θ_1, a_{-1}) at t, and we know that at t we move into an incompatible profile. Wlog suppose this is player 2 who moves into an incompatible profile, and denote her new action by by a_2'. Clearly, the fact that the action was taken in equilibrium implies that $a_3 < \theta_3$ (otherwise 2 should not switch, as stated in Lemma 3), and therefore that $a_2' > \theta_2$ (otherwise (θ_1, a_2', a_3) would remain compatible). This implies that on the equilibrium path player 3 is idle (Lemma 3: both i and j are at their θ's or above). Therefore, player 2 must have made the switch because he expects that in equilibrium player 1 will make a switch, generating eventually an agreement. Note, however, that player 2's final asking is bounded from above by $1 - \theta_1 - a_3$ (Corollary 3), which is less than a_2'. Therefore, we must have on equilibrium path player 2 scaling down his asking later on. This means that on the true equilibrium path of (θ_1, a_{-1}) player 3 is idle, player 2 switches at least twice, and player 1 switches at least once. Thus, the sum of the (θ_1, a_{-1}) subgame continuation values for players 1 and 2 is bounded from above by $v = 1 - a_3 - 2c_2(t) - c_1(next_i(t))$.

Denote by t^{**} the solution to $c_2(next_2(t)) + c_1(next_1(t)) = 1 - a_3$, and let $\theta_1' = c_1(next_1(t^{**}))$ and $\theta_2' = c_2(next_2(t^{**}))$. Note that $t^{**} > t^*$, $\theta_1' > \theta_1$, and $\theta_2' > \theta_2$ (because $a_3 < \theta_3$). If $t \geq t^{**}$ then v is negative, which implies that either 1 or 2 obtains a negative payoff, which leads to a contradiction (she could have got at least zero by not moving). If $t \in (t^*, t^{**})$ then: (i) player 1 can guarantee herself at least $\theta_1' - c_1(next_1(t))$ (by playing θ_1' and not mavin anymore); and (ii) player 2 can guarantee herself at least $\theta_2' - c_2(t)$ (this two claims are proved below). Now, note that $\theta_1' - c_1(next_1(t)) + \theta_2' - c_2(t) = 1 - a_3 - c_2(t) - c_1(next_1(t)) > 1 - a_3 - 2c_2(t) - c_1(next_1(t))$ which was the upper bound for the sum of the payoffs obtained above. This means that at least one of the players is doing worse than he could have guaranteed himself, which is a contradiction to the original story being an equilibrium.

Now, let us prove claim (i) above, i.e., that at $t \in (t^*, t^{**})$ player 1 can guarantee himself $\theta_1' - c_1(next_1(t))$. This is done by switching immediately to θ_1' and playing it throughout, leading us to a profile of (θ_1', a_2', a_3) [23] . By the Lemma 3, player 3 is idle (and is expected to remain idle in equilibrium) as long as player 2 does not switch below θ_2. Then, after t^{**} player 2 does not switch against (θ_1', a_3): by Corollary 3 she knows that she cannot expect to have a final asking higher than $1 - \theta_1' - a_3 = \theta_2'$ unless there is no agreement, and hence a switch is too costly at $t > t^{**}$. Just before t^{**} player 2 is better off playing θ_2' compared to any $a_2 > \theta_2'$, guaranteeing an agreement. Finally, we can rule out the possibility that player 3 switches. As we have mentioned, this would have to be triggered by a move of player 2 below θ_2, but this is incompatible with

[23] More precisely, we could be at a different (θ_1', a_2'', a_3) profile with $a_2'' > \theta_2$, if player 2 had switched again before $next_1(t)$, but the proof is valid also in this case.

Lemma 4.

A proof of claim (ii) is exactly analogous to the previous one, and in which the roles of players 1 and 2 are reversed.

Lemma 8 *For three players, consider a state* (θ_i, a_{-i}) *at* $t < t^*$*. Then "always* θ_i*" guarantees* θ_i*.*

Proof It is sufficient to see what happens just before t^*. Let player k be the one who moves last before t^* (i.e. $prev_j(t^*) < prev_k(t^*)$). It is easy to see that if $a_j > \theta_j$ player k will not switch (player i's final asking is at least θ_i and player j final asking would be greater than θ_j, so player k will obtain negative payoffs by switching), and that if $a_j = \theta_j$ player k will switch to θ_k if $a_k > \theta_k$ or $a_k = 0$ and will not switch otherwise. Therefore, it is easy to see that if player j plays θ_j he guarantees himself positive payoffs (Lemma 7, and that this is better than playing $a_j > \theta_j$ against $a_k \geq \theta_k$ (which would result in a no agreement). This is enough to show that the profile becomes compatible for $a_k \geq \theta_k$ and $a_j \geq \theta_j$. All that remains to show is that they get into a compatible profile in the case of $a_k < \theta_k$ or $a_j < \theta_j$. Let a'_k and a'_j be what they do in equilibrium in their turns to play. Obviously, in equilibrium we cannot have $a'_k > \theta_k$ and $a'_j > \theta_j$: this will result in a no-agreement, while at least one of them switched, which is a contradiction. Suppose, wlog., that player j has $a'_j > \theta_j$. In this case it is easy to follow similar arguments to the proof of the previous lemma in order show that it implies that he has to scale down his asking later on, and that this expects another switch by player i, thus resulting in a contradiction to the total values (player k is idle as before). Finally, note that we completely ignored throughout the order in which player i gets to play just before t^*. This is because it is not relevant: once $t > \max_i c_i^{-1}(\theta_i - \frac{1}{L})$ then player i commits himself (in equilibrium) to a final asking of at least θ_i, which is the only thing we need.

B Algorithms

In this part we describe the Solution Algorithm and the Limit Algorithm, and comment on the differences. It may be useful to go through the different steps of the algorithm in parallel to the specific example we provide in the end of this section. Finally, we prove that the algorithm terminates in a finite time.

In what follows, we use the destination operator $d(i)$: Given a function X, we define $X(a,p)^{d(i)} = X((\tilde{a}_i, a_{-i}), p)$. In words, the value of the function X for player p at the destination of a move by player i from the initial profile a.

B.1 Solution Algorithm

Given a particular game (Π, C, g) the algorithm steps are described below.

1. **Initialization**

   (a) Set $m = 0$, where m stands for the stage of the algorithm (starting from the end of the game).

   (b) Set $t_0^* = T$ to be the last critical time encountered.

   (c) Set $V_0(a, p) = \Pi$, where $V_m(a, p)$ is the continuation value of player p at profile a just after t_m^*.

   (d) Set $AM_0(a, p) = 0$, where $AM_m(a, p)$ is an indicator function. It is equal to one if and only if there is an active switch at time t_m^* by player p from profile a to profile $(\tilde{a}_p, a_{\tilde{p}})$.

   (e) Set $IM = \{(a, p) | a \in A_1 \times A_2,\ p = 1, 2\}$ be the set of inactive moves.

   (f) Go to step 2.

2. **Update** (m, V_m, AM_m)

   (a) $m = m + 1$

   (b) Find the next critical time, and the action (a^*) and player (p^*) that are associated with it. This is done by comparing the potential benefits and costs for each move. Given the monotonicity of the cost function, each beneficial switch is associated with a unique cutoff point in time after which the move is not taken anymore because it is too costly. To compute this, we use some auxiliary definitions:

      i. Let $SM_{m-1}(a, p)$ be an ordered set of decision nodes $\{(a_0, p_0), (a_1, p_1), ..., (a_k, p_k)\}$, such that $(a_0, p_0) = (a, p)$ and $(a_i, p_i), (a_{i+1}, p_{i+1})$ are in the set if and only if $AM_{m-1}(a_i, \tilde{p}_i) = 1$ and $(a_{i+1}, p_{i+1}) = ((a_{p_i}, \tilde{a}_{\tilde{p}_i}), \tilde{p}_i)$, or $AM_{m-1}(a_i, \tilde{p}_i) = 0$, $AM_{m-1}(a_i, p_i) = 1$, and $(a_{i+1}, p_{i+1}) = ((\tilde{a}_{p_i}, a_{\tilde{p}_i}), p_i)$.
      This defines a sequence of consecutive switches, starting after a move of player p into profile a. The sequence ends at a profile from which there is no active move. It is finite and contains up to three switches, because a full circle cannot exist.
      Note that SM_{m-1} is solely a function of AM_{m-1}.

ii. Let $\overline{SM}_{m-1}(a,p)$ denote the final element in $SM_{m-1}(a,p)$, which stands for the node we end up at after a sequence of consecutive switches.

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

iv. Let $FS_{m-1}(a,p) = \sum_{i=1}^{|SM_{m-1}(a,p)|} I(\exists a'_{\tilde{}p} s.t.((a''_p, a'_{\tilde{}p}), p) = (a_i, p_i) \in SM_{m-1}(a,p))$.
where I is the indicator function. This computes how many switches will immediately follow a switch by player p to profile a. For two-by-two games, this is at most one. This takes into account the costs player p will incur if he decides to move to profile a.

Now, we 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, thus it is never taken. 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 in which the cost of switching is less than the benefit. The last case is similar, but takes into account that the move involves immediate switch at the next period.

$$t_m(a,p)^{24} = \begin{cases} 0 \\ \quad \text{if } \Delta V_{m-1}(a,p) < 0 \\ prev_p(t^*_{m-1}) \\ \quad \text{if } \Delta V_{m-1}(a,p) \geq 0 \text{ and } FS_{m-1}(a,p) > 0 \\ Max\left\{t \in g_p, t < t^*_{m-1} | C_p(a_p \to \tilde{}a_p, t) \leq \Delta V_{m-1}(a,p)\right\} \\ \quad \text{if } \Delta V_{m-1}(a,p) \geq 0 \text{ and } FS_{m-1}(a,p) = 0 \text{ and } FS_{m-1}(a,p)^{d(p)} = 0 \\ Max\left\{t \in g_p, t < t^*_{m-1} | C_p(a_p \to \tilde{}a_p, t) + C_p(\tilde{}a_p \to a_p, next_p(t)) \leq \Delta V_{m-1}(a,p)\right\} \\ \quad \text{if } \Delta V_{m-1}(a,p) \geq 0 \text{ and } FS_{m-1}(a,p) = 0 \text{ and } FS_{m-1}(a,p)^{d(p)} > 0 \end{cases}$$

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

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

$(a^*, p^*) = \underset{(a,p)\in IM}{ArgMax}\{t_m(a,p)\}$ [25]

$t_m^* = t_m(a^*, p^*)$

Abort if $t_m^* = 0$. This happens if a player is indifferent between two actions in the beginning of the game.

$p_m^* = p^*$

(c) Update the set of active moves:

We update only when it is necessary. As can be seen below, there are four cases. The first is the new active move, the one associated with the new critical time, that becomes active. The second is a move by the other player that originate from the same action profile. It can be shown that unless $m = 2$ or unless we are in a early part of the game, this move remains active. In the other cases, it has to be reevaluated, so it becomes inactive. The third case is a move whose destination is the origin of the new active move. Such a move has to be reevaluated at all times, so that it gets deleted, and is reevaluated the next iteration. Finally, all other moves are unaffected by the new active move, so they remain active or inactive, as they were before.

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

(d) We compute the continuation values of the players just after the last critical point found 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.b), and subtracting the switching costs incurred by the player along this sequence. These switching costs are incurred just after t_m^*.

Let us, first, define this mapping in general to be $V^{new}(V^{old}, AM, t, \overline{p}) =$ $V^{new}(V^{old}, SM(AM), t, \overline{p})$, such that:

$$V^{new}(a,p) = V^{old}(\overline{SM}(a,\overline{p})) - \sum_{i=1}^{|SM(a,\overline{p})|} I(\exists a'_{\tilde{\ } p} \text{ s.t. } ((a''_p, a'_{\tilde{\ } p}), p) = (a_i, p_i) \in SM(a,\overline{p}))C_p(a'_p \to a''_p, (next_p)^i(t))$$

where $I(\cdot)$ is the indicator function.

[25] Potentially, the $ArgMax$ is a correspondence rather than a function. This is why we use '∈' rather than equalities in part 2.c of the algorithm. Given the way we construct $t_m(a,p)$, the multiple solutions must be associated with a unique p^*, so part 2.b is well-defined.

In particular, let $V_m = V^{new}(V_{m-1}, AM_{m-1}, t^*_m, \tilde{\ } p^*_m)$.

(e) Let $IM = \{(a,p) | AM_m(a,p) = 0 \text{ and } AM_m(a,p)^{d(p)} = 0\}$. If $\#IM = 0$, let $\overline{m} = m$, then set $t^*_{\overline{m}+1} = 0$, and **Terminate** (all moves are active). Otherwise, go back to stage 2.

3. **Output**

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

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

For practical reasons we nevertheless define the output of the Solution Algorithm to be

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

B.2 Limit Algorithm

Given (Π, C), independent of g, we define the Limit Algorithm. This algorithm mimics the Solution Algorithm described in the previous section (see Appendix B.1), with slight differences. Along the way, we describe these differences, and comment on them.

1. **Initialization**–same as in the Solution Algorithm.

2. **Update** (m, V_m, AM_m)

(a) same as in the Solution Algorithm.

(b) We start exactly the same as in the solution algorithm, with slight differences in the definition of the $t_m(\cdot)$. The differences are created by taking the limit as $\varphi(g) \to 0$ of the corresponding definitions in the Solution Algorithm.

$$t_m(a,p) = \begin{cases} 0 \\ \quad \text{if } \Delta V_{m-1}(a,p) < 0 \\ t^*_{m-1} \\ \quad \text{if } \Delta V_{m-1}(a,p) \geq 0 \text{ and } FS_{m-1}(a,p) > 0 \\ Min\left\{t^*_{m-1}, \{t|C_p(a_p \to \tilde{\ } a_p, t) = \Delta V_{m-1}(a,p)\}\right\} \\ \quad \text{if } \Delta V_{m-1}(a,p) \geq 0 \text{ and } FS_{m-1}(a,p) = 0 \text{ and } FS_{m-1}(a,p)^{d(p)} = 0 \\ Min\left\{t^*_{m-1}, \{t|C_p(a_p \to \tilde{\ } a_p, t) + C_p(\tilde{\ } a_p \to a_p, t) = \Delta V_{m-1}(a,p)\}\right\} \\ \quad \text{if } \Delta V_{m-1}(a,p) \geq 0 \text{ and } FS_{m-1}(a,p) = 0 \text{ and } FS_{m-1}(a,p)^{d(p)} > 0 \end{cases}$$

$(a^*, p^*) = \underset{(a,p)\in IM}{ArgMax}\{t_m(a,p)\}$[26]

Abort if $|p^*| > 1$. This means equal critical times for different players. This could not happen in the Solution Algorithm, because each player played in different points in time, so the order and the grid solved the ambiguity. In the Limit Algorithm, this problem may occur. Whenever it occurs, this is a case for which Theorem 1 holds only generically.

$t^*_m = t_m(a^*, p^*)$

Abort if $t^*_m = 0$. This means binding indifference by one of the players.

$p^*_m = p^*$

(c) same as in the Solution Algorithm.

(d) Again, this part is the same as before, with the only difference is that we take the limit of the cost function as $\varphi(g) \to 0$ (which is well-defined given the continuity of the cost function). Hence, the definition for V^{new} becomes

the following (the difference is in the last element):

$V^{new}(a,p) = V^{old}(\overline{SM}(a,p)) - \sum_{i=1}^{|SM(a,p)|} I(\exists a'_{\tilde{\ } p} \text{ s.t. } ((a''_p, a'_{\tilde{\ } p}), p)$
$= (a_i, p_i) \in SM(a,p))C_p(a'_p \to a''_p, t)$

(e) same as in the Solution Algorithm.

[26]Potentially, the *ArgMax* is a correspondence rather than a function. This is why we use '$\in$' rather than equalities in part 2.c of the algorithm. If the multiple solutions are associated with a unique p^*, as in the finite case, part 2.b is well-defined. If, however, these multiple solutions involve (two) different players, then the algorithm terminates - we would classify the game to be one for which we cannot solve the limiting case, or, more precisely, the outcome of the limiting case would not be grid invariant.

3. **Output**

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

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

For practical reasons we nevertheless define the output of the Solution Algorithm to be

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

B.3 Finiteness of the Algorithms

We provide now the proof that both algorithms finish in finite time.

Lemma 9 *The algorithm ends in finite time, and in particular* $\overline{m} \leq 8$ *for any* (Π, C, g).

Proof. (The proof is identical for both the Solution Algorithm and for its limit version) The algorithm finishes when $\#IM = 0$. (1) Observe that if $AM_m(a,p) = 1$ then $AM_m(a,p)^{d(p)} = 0$ and vice versa, thus $\#IM = 0$ implies that $\#AM = 4$. (2) Observe that whenever $\exists p, m$ s.t.$\sum_a AM_m(a,p) = 2$ we get into a "termination phase": the algorithm is guaranteed to terminate within at most two more stages: when this is the case, it can be verified that $\sum_a AM_{m+1}(a, \tilde{\ }p) = 2$ and the two are in the same direction, so that player p's two moves immediately become active at stage $m+2$, without any deletion of an active move by player $\tilde{\ }p$, terminating the algorithm. (3) Observe that $\#AM$ is non-decreasing in m: every iteration we add an active move $(AM(a^*, p^*))$ and may potentially remove at most one active move.[27] (4) Note that for $m > 2$, and before we reach the "termination phase", we delete an active move (a,p) only when $(a,p) \in ((a^*_{p^*}, \tilde{\ }a^*_{\tilde{\ }p^*}), \tilde{\ }p^*)$. In particular, at stage m, if we delete a move, this is a move by player $\tilde{\ }p^*_m$. (5) It is implied by (2) and (4) that once $\#AM = 2$ the algorithm is guaranteed to terminate within at most 3 stages: if the two active moves are by the same player then we are done, and if by different players then the next move, whether it deletes an active move or not, guarantees that in the next stage one player will have 2 active moves. Therefore, all we need to show is that we cannot have an infinite sequence of moves, such that any move that becomes active at stage m, becomes inactive at stage $m+1$. Suppose, toward contradiction, that such

[27] 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 (3) below.

an infinite sequence exists. Without loss of generality, consider $m = 2$, in which $AM_2(a, p) = 1$ for some (a, p), and $AM_2(a', p') = 0$ for any $(a', p') \neq (a, p)$. If (a, p) is deleted at $m = 3$, it must be that the new active move is such that $AM_3((\tilde{a}_p, a_{\tilde{p}}), \tilde{p}) = 1$. Similarly, we obtain that $AM_4(\tilde{a}, p) = 1$ and that $AM_5((a_p, \tilde{a}_{\tilde{p}}), \tilde{p}) = 1$. This gives us the following contradiction. By $AM_2(a, p) = 1$ we know that $V_3(a, \tilde{p}) = V_3((\tilde{a}_p, a_{\tilde{p}}), \tilde{p})$. By $AM_3((\tilde{a}_p, a_{\tilde{p}}), \tilde{p}) = 1$ we know that $V_3((\tilde{a}_p, a_{\tilde{p}}), \tilde{p}) < V_3(\tilde{a}, \tilde{p}) - C_{\tilde{p}}(a_{\tilde{p}} \to \tilde{a}_{\tilde{p}}, t)$ for any $t < t_3^*$. It is easy to see that $t_4^* < t_3^*$, so the above implies that $V_5(a, \tilde{p}) = V_3((\tilde{a}_p, a_{\tilde{p}}), \tilde{p}) < V_3(\tilde{a}, \tilde{p}) - C_{\tilde{p}}(a_{\tilde{p}} \to \tilde{a}_{\tilde{p}}, t_4^*) = V_5((\tilde{a}_p, a_{\tilde{p}}), \tilde{p})$, while by $AM_4(\tilde{a}, p) = 1$ we also know that $V_5(\tilde{a}, \tilde{p}) = V_5((a_p, \tilde{a}_{\tilde{p}}), \tilde{p})$. The two last equations imply that $\Delta V_5(\tilde{a}, \tilde{p}) > \Delta V_5((a_p, \tilde{a}_{\tilde{p}}), \tilde{p})$, which is a contradiction to the fact that $(a^*, p^*) = ((a_p, \tilde{a}_{\tilde{p}}), \tilde{p})$ at $m = 5$. This, together with (5) above, also shows us that we always have $\overline{m} \leq 8$. ■

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

B.4 Using the Limit Algorithm–An Example

Here we will show how to solve for the equilibrium of a specific game using the Limit Algorithm. For this purpose we use example C.4, with a common cost function of $c(t) = t$. The full equilibrium strategies are given in the end of this example.

We begin by searching for the first point in time at which one of the players is willing to pay the switching costs at some profile. This is the solution for $\bar{t}$ in equation (1). The fact that $c(t) = t$ simplifies the algebra, implying that all we need is to search for the greatest difference among the four possible positive-value switches. The differences are 13 and 1 for player 1, and 7 and 5 for player 2. Thus, we have that $t_1^* = 13$, which is associated with player 1 switching from (D, L) to (U, L).

Going backwards, we know that this switch remains active at least until the next critical point. Note that in the search for the next critical point, player 2 understands that his continuation value at profile (D, L) is 3 rather than 5, because it is followed by player 1 switching. Having this in mind, the three differences we have to consider (not including the switch which is already active) are 1 for player 1, and 7 and 3 for player 2. Thus, only before $t_2^* = 7$ is player 2 willing to switch from (U, L) to (U, R). This changes the consequences of switching to (U, L). Therefore, player 1 has to reconsider his action at profile (D, L). For this reason, the active switch by player 1 is *deleted* from the set of active moves, and player 1's continuation value at (D, L) is updated to its previous value minus the cost of the delay. This is because, potentially, player 1 would like to delay the switch into (U, L) until just after $t_2^* = 7$, when player 2 does not react by switching any longer. Hence, the continuation value for player 1 at (D, L) is now $13 - c(t_2^*) = 6$. To recap, the continuation values matrix at time t, just before $t_2^* = 7$, is the following:

	L	R
U	$1, 10 - c(t)$	$1, 10$
D	$6, 3$	$0, 0$

The next critical point is found as the maximum over 5 and 1 for player 1, and 3 for player 2. Thus, $t_3^* = 5$, which is associated with player 1 switching from (U, L) to (D, L). Player 2 still does not like profile (U, L), so his previously active switch remains active. This means that in stage 3 there are two active switches, one by each player. Therefore, we just need to check for the other two possible switches. Thus, the next critical point is the maximum over 1 for player 1, and 3 for player 2. Thus, $t_4^* = 3$, which is associated with player 2 switching from (D, R) to (U, R). At this point, just before $t_4^* = 3$, player 2 is fully flexible–his strategies depend only on his opponent's most recent action–so we are in the early part of the game (see Proposition 1 and the discussion that follows). Hence, at this point it is easy to solve for the remaining of the stages, so the rest of stages in the algorithm are purely technical, solving for the strategies of the players during the early part of the game, which runs from $t = 0$ until $t_4^* = 3$.

The final values of the game are $(6, 3)$, and the initial actions are (D, L). To find the equilibrium path, one needs to start at (D, L) and check for active switches that originate at (D, L). Indeed, there exists a strategic delay, which takes place at (just after) $t = 7$, where player 1 switches to (U, L). The full output of the algorithm is given in the following table and picture.

Stage	Interval	Active Moves
$m=6$	$[0,3]$	$(U,L)\to(D,L)$ and $(U,R)\to(D,R)$ by player 1 $(U,R)\to(U,L)$ and $(D,R)\to(D,L)$ by player 2
$m=5$	3^*	$(U,L)\to(D,L)$ and $(U,R)\to(D,R)$ by player 1
$m=4$	3^*	$(U,L)\to(U,R)$ and $(D,R)\to(D,L)$ by player 2
$m=3$	$(3,5]$	$(U,L)\to(D,L)$ by player 1 $(U,L)\to(U,R)$ by player 2
$m=2$	$(5,7]$	$(U,L)\to(U,R)$ by player 2
$m=1$	$(7,13]$	$(D,L)\to(U,L)$ by player 1
$m=0$	$(13,T]$	None

* Stages 4 and 5 are "instant" stages (one-period stages), in which only one player moves. More precisely, the move of stage 4 is done at $prev_1(3)$, and the move of stage 5 is done at $prev_2(prev_1(3))$.

C Examples

In all the examples below, for ease of exposition we use $C_i(a_i \to a_i', t) = c(t)$ for any i, a_i, a_i'. Everything would work (qualitatively) just the same if we had a more general cost structure that satisfies the conditions given in Section 2.

In all examples the row player is denoted by player 1, and the column player by player 2, and all parameters are positive.

Each example is followed by an informal discussion, highlighting the main features of its Subgame Perfect Equilibrium. By no means is this discussion aimed at providing a complete characterization of the equilibrium of each game.

C.1 Battle of the Sexes

	Boxing	*Opera*
Boxing	A,b	$0,0$
Opera	$0,0$	a,B

where $A > a$, $B > b$.

- The equilibrium outcome is: [*Boxing,Boxing*] if $b > a$, [*Opera,Opera*] if $a > b$. If $a = b$ the equilibrium depends on the grid, even for fine grids.

- Key stages in the equilibrium strategies: off equilibrium, the two players play a war of attrition at profile [*Boxing,Opera*], i.e. at each player's favorite activity. If, for example, $b < a$, there exists a stage of the game that lasts from $c^{-1}(b)$ to $c^{-1}(a)$, in which player 2 is fully committed to playing Opera for the rest of the game, while player 1 is still flexible. At this stage, player 1 prefers to switch to Opera, giving him positive payoffs of $a - c(t)$. Both players foresee this, resulting in player 1 giving up immediately.

- Comments: First, note that it does not matter how much one likes his favorite activity, but what matters is how much one likes his least favorite activity. The player who wins is the one who likes it less, making the threat of "staying at home" credible. Second, note that if $a = b$ the game is not grid invariant–the player who wins is the player who can commit first before the common critical point $t^* = c^{-1}(a) = c^{-1}(b)$, i.e. player i wins if and only if $prev_i(t^*) < prev_j(t^*)$.

C.2 Prisoners' Dilemma

	Defect	*Coop*
Defect	d, d	D, c
Coop	c, D	C, C

where: $D > C > d > c$.

- The equilibrium outcome is: [*Defect,Defect*] if $(D - C) \geq (d - c)$, and [*Coop,Coop*] if $(d - c) > (D - C)$.

- Key stages in the equilibrium strategies: clearly, each player can guarantee at least d by defecting all the time. Hence, we need to check if [*Coop,Coop*] can be supported as an equilibrium. Suppose first that $(D - C) \geq (d - c)$. In such a case, there exist a point in time on the grid, in which at least one of the players (both, in case the inequality is strict) would defect, knowing that the switching costs would be too high for the opponent to react by defecting as well. Both players know it, and hence begin by defecting. Once $(d - c) > (D - C)$ such a point in time does not exist, and [*Coop,Coop*] can be supported as (the unique) equilibrium. Each player understands that as long as it is potentially profitable for him to defect, it is profitable for the other player to retaliate and defect as well.

- Comments: Note that in this case the common critical time does not matter for the equilibrium outcome. For example, if $(d-c) > (D-C)$ one can easily verify that the switches that become active at the same time are "unrelated"–one's origin is neither the origin nor the destination of the other.

C.3 Generalized Matching Pennies

	A	B
A	K, ϵ_1	$0, m$
B	$0, m$	k, ϵ_2

where: $\epsilon_1, \epsilon_2 \approx 0$, $K > m$, $K \geq k$.

- The equilibrium outcome is: [A,A] if $\epsilon_1 > \epsilon_2$, [B,B] if $\epsilon_2 > \epsilon_1$. However, if $k < m - \epsilon_1$ and $\epsilon_1 > \epsilon_2$ then play begins with [B,A] and is switched (by player 1) to [A,A] only after $t = c^{-1}(m - \epsilon_1)$. If $\epsilon_1 = \epsilon_2$ then player 2 is indifferent, and both outcomes are equilibria.

- Key stages in the equilibrium strategies: it is clear that player 1 wins the game because he is the most flexible player, allowing him to enjoy the second-mover advantage of the matching pennies game. After $t = c^{-1}(m - \epsilon_1)$, player 1 can still react while for player 2 it is too costly. However, player 2, knowing he is losing the game either way, can at least guarantee his maxmin payoffs by playing A if $\epsilon_1 > \epsilon_2$, and B otherwise. Thus, even though $K > k$, player 1 cannot obtain K for sure. Moreover, when $k < m - \epsilon_1$ and $\epsilon_1 > \epsilon_2$, even though player 2 plays A, player 1 must delay in equilibrium. This is because there is a stage of the game, between $c^{-1}(k)$ and $c^{-1}(m - \epsilon_1)$, in which player 2 would switch to B at profile [A,A], knowing that it will not be contested by player 1. In order to avoid this player 1 starts by playing B, and switches to A only at $c^{-1}(m - \epsilon_1)$, once such a switch by player 2 is not profitable any longer.

- Comments: In the fully symmetric case, in which $\epsilon_1 = \epsilon_2 = 0$ and $K = k = m$, the grid matters. The player who wins is the player who moves last before the common critical point $t^* = c^{-1}(m) = c^{-1}(K)$, i.e. player i wins if and only if $prev_i(t^*) > prev_j(t^*)$.

C.4 Delay with Super-Dominant Strategy

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

- Note that action U is super-dominant for player 1 in this example (see Definition 6).

- The equilibrium path is to play [D,L] until $t = c^{-1}(7)$, and then switch (by player 1) to [U,L].

- Key stages in the equilibrium strategies: the key stage is between $c^{-1}(3)$ and $c^{-1}(7)$. During this stage, at profile [U,L] player 2 finds it profitable to switch to R. After $c^{-1}(7)$, the switching costs are greater than the benefits of the switch $(10 - 3 = 7)$, thus the switch is not profitable anymore. Before $c^{-1}(3)$, a switch by player 2 would allow player 1 to credibly switch to D. This switch will follow by player 2 switching back to L, taking the game back to its equilibrium path, which ends at [U,L]. Therefore, player 1 avoids this off-equilibrium switch by player 2 by playing first D, and switching to U only after $c^{-1}(7)$. This gives him payoffs of $13 - c(c^{-1}(7)) = 6$, which is more than 1, that he would have obtained by playing U to begin with. Note also that this is credible only because player 1 can commit himself to not switching to U had player 2 played R, because 3 is greater than 1 (there is an off-equilibrium war of attrition at the early stage of the game at the profile [D,R]).

- Comments: the key point is that late enough in the game, but still early for player 2 to react, player 1 can commit himself not to play his super-dominant strategy as long as player 2 does not "cooperate". If this commitment was not attainable (for example, if the payoffs of player 1 at [U,R] were greater than 3), player 2 could simply play R, knowing that player 1 would eventually switch to his super-dominant strategy.