

Switching Costs in Frequently Repeated Games¹

Barton L. Lipman
University of Wisconsin

Ruqu Wang
Queen's University

Current Draft
June 1999

¹The authors thank Ray Deneckere, Eddie Dekel, Debraj Ray, John Riley, Arthur Robson, and seminar participants at Boston University, Guelph, Hebrew University, Hong Kong University of Science and Technology, Humboldt University of Berlin, Northwestern, the Northwestern Microeconomic Theory Conference, Wisconsin, Tel Aviv University, Tulane, University of New Orleans, and Waterloo for helpful suggestions, Northwestern University for enjoyable visits during which this research was begun, and SSHRCC for financial support.

Abstract

We demonstrate the surprising effect of two simple variations on the usual finitely repeated game model. In particular, we add a small cost of changing actions and consider the effect of increasing the frequency of repetitions within a fixed period of time. We illustrate the wide range of commitments this can make credible and the effects these have. Not surprisingly, switching costs can make it credible to not change one's action. However, this effect can occur even for very small switching costs, but only for certain actions. For example, because of this kind of commitment, we get a unique subgame perfect equilibrium in coordination games when the Pareto dominant and risk dominant outcomes coincide, but not otherwise. In other games, switching costs can make certain commitments less credible even as some are made more credible. For example, in the repeated Prisoners' Dilemma, these costs reduce the incentive to deviate from mutual cooperation. However, they also reduce the incentive to switch from cooperation to punish a deviator. As a result, we show that switching costs might or might not enable cooperation in the finitely repeated Prisoners' Dilemma depending on which of these effects dominates. More surprisingly, switching costs can make some very complex threats credible, leading to some unexpected possibilities. For example, we show that in some games, this commitment ability can enable a player to obtain a payoff strictly above his Stackelberg payoff. Also, we show that in some games, actions must be changed along the equilibrium path in the unique subgame perfect equilibrium.

1 Introduction

In this paper, we demonstrate the surprisingly strong effects of two seemingly small changes from the usual finitely repeated game framework. Specifically, we add a small cost to changing actions from one period to another and consider the effect of repeating the game more often in a fixed amount of time. These changes make credible certain commitments, the nature of which depend in complex ways on the game being repeated. As a result, all the standard results for finitely repeated games are overturned: stage games with a unique Nash equilibrium can have multiple subgame perfect equilibria in the repeated game with switching costs, while stage games with many Nash equilibria can have a unique subgame perfect equilibrium in the repeated game. Before explaining these results, we must first describe the model in more detail.

To keep the analysis as close as possible to the standard repeated game model, we treat the cost of switching actions as constant over time and across players and focus on the case where it is “small.” There are several reasons for studying such a cost. First, it is a simple way of capturing one type of bounded rationality. If playing a given action is complex, then changing from one action to another may be “hard.” Second, in many economic contexts, changing actions involves real costs. For example, in related work (Lipman and Wang [1999b]), we consider a game due to Gale [1995] in which firms choose between investing and not investing in each period. It seems quite reasonable to believe to switching between such actions incurs fixed “set up” or “shut down” costs. In price setting games, it is natural to assume that there are menu costs associated with price changes. Because the existence of such costs seems plausible for many economic settings, their inclusion seems is a natural “robustness check” for the standard finitely repeated game model.

The switching cost creates a role for the second factor we consider, namely frequent repetition. To understand the idea, suppose that the game is played in continuous time but that actions can only be changed at fixed intervals. We fix the length of time the overall game is played and vary the number of periods (or dates at which actions can be changed) and hence the length of each period (or the length of time for which actions are fixed). As the frequency of play increases, the length of a period and hence the payoffs in a period shrink relative to the switching cost. To see why this is important, note that if the length of the period is sufficiently small, even a tiny switching cost is too large to make a change of action worthwhile if it only leads to a one-period gain. Hence standard backward induction arguments break down. It is important to emphasize that because the total length of time the repeated game is played is held constant throughout, it is only payoffs *per period* which shrink relative to the switching cost, not payoffs over the entire horizon. In fact, as we explain in the conclusion, our results continue to hold if we increase the length of time the game is played while shrinking the length of each period. Hence the switching cost can go to zero relative to total payoffs without affecting our

analysis.

To be more specific, let ε be the cost to changing actions and Δ the length of a period. The length of time the overall repeated game is played is M , independent of Δ , so the number of periods is M/Δ . Our results all take the form of showing that for any sufficiently small $\varepsilon > 0$, there is a $K > 0$ such that our equilibrium results hold for $\Delta \in (0, K\varepsilon)$. In other words, as long as the switching costs are small enough, we get a particular result as long as the game is repeated frequently enough.

As in the usual finitely repeated game model, subgame perfection identifies the “threats” or “promises” which are credible. The switching costs change which kind of commitments are credible in ways that are sometimes surprising. Some of the effects are obvious: for example, as noted above, the commitment not to change actions in the last period of the game is certainly credible if the length of a period is short enough relative to the switching costs. The more interesting phenomena is that credibility of such commitments late in the game can make credible these or other commitments earlier in the game in a way which dramatically affects the equilibrium. For example, in some games, switching costs can make a commitment to not change from a particular action credible even at the very beginning of the game. In these games, a backward induction argument translates the credibility of the commitment late in the game to earlier stages. The reason this happens is that in these games, a player who will be credibly committed if he plays a certain action tomorrow will, as a consequence, refuse to change his action today. Hence the commitment is credible today and, by induction, at all previous dates. More specifically, we show that in a surprisingly broad class of games, one player will be able to achieve his Stackelberg payoff this way. In particular, in a coordination game where the Pareto dominant and risk dominant outcomes coincide, the unique subgame perfect equilibrium is to repeat this outcome always. In a generic Battle of the Sexes game, the only subgame perfect equilibrium will be to repeat one of the stage game Nash equilibria every period, where which equilibrium gets repeated depends on the payoffs. Thus games which have many subgame perfect equilibria without switching costs may have a unique equilibrium with switching costs.

In some games, switching costs affect outcomes in contradictory ways. Consider, for example, the repeated Prisoners’ Dilemma. Cooperation in the Prisoners’ Dilemma can only be sustained if deviation from cooperation can be punished. Switching costs have two effects in this context. First, it may eliminate the incentive to switch away from cooperation in order to cheat the opponent, an effect which clearly favors cooperation. Second, however, it may eliminate the incentive to switch away from cooperation in order to punish the opponent if he has deviated. Which effect dominates depends on the payoffs of the game, so cooperation can be sustained in some finitely repeated Prisoners’ Dilemma games with switching costs and not in others. When cooperation is sustainable, there is a large set of equilibrium payoffs — in particular, always defecting is still an equilibrium.

Hence games which have a unique subgame perfect equilibrium without switching costs can have multiple equilibria with switching costs.

A still more surprising conclusion is that switching costs can make some very complex threats credible. We give an example in which one player is able to credibly threaten to play a specific action until the opponent changes his action and then to play a different action thereafter, switching back later if the opponent goes back to the original action. Clearly, this is a more powerful threat than simply refusing to change one's action. As a consequence, the threatener can receive a payoff above his Stackelberg payoff. The other player can end up with a payoff below his minmax (though not below his pure strategy maxmin). Even more surprisingly, in the example, this is the *unique* subgame perfect equilibrium and there is a different, also unique, subgame perfect equilibrium without switching costs. We also show that these more complex threats can lead to a situation where the unique subgame perfect equilibrium with switching costs requires one player to switch actions along the equilibrium path.

The rest of this section briefly comments on the related literature. The next section contains the model. Section 3 shows how switching costs can enable one player to credibly commit himself to not changing from a specific action. Section 4 contains the results discussed above for the Prisoners' Dilemma. Section 5 shows in the context of an example how switching costs can make some very complex threats credible. Concluding remarks are offered in Section 6. Proofs not in the text are contained in the Appendix.

Of course, the literature on finitely repeated games (see, *e.g.*, Benoit and Krishna [1985]) is relevant to our work, but we do not attempt a survey here. In addition, there are several other related strands of the literature. First, a number of economic models have studied the effect of switching costs for consumers on competition between firms. See, for example, Beggs and Klemperer [1992], Padilla [1995], or Wang and Wen [1998]. Second, the literature on delay in bargaining and the Coase conjecture (such as Gul and Sonnenschein [1988]) has studied the effect of shrinking the length of the period. Third, there are many results on the robustness of the risk dominant outcome in coordination games, often with emphasis on the case where risk dominance and Pareto dominance coincide. See Carlsson and van Damme [1993], Kandori, Mailath, and Rob [1993], Young [1993], and Robson [1994], for example. Fourth, our paper can be seen as studying a particular stochastic game which is "close" to a repeated game and considering the effect of the dynamic aspect on the set of equilibrium outcomes. As discussed by Dutta [1995a, 1995b], some standard repeated game results do not carry over to the broader class of stochastic games, even to some games arbitrarily close to repeated games. What is new here is the consideration of finite horizons (instead of infinite repetition), the role of frequent repetition, and the particularly simple nature of the dynamic aspect (the switching cost). Even in the infinitely repeated case, our results are not direct corollaries of Dutta's — see Lipman and Wang [1999a] for details. Fifth, it is well known that

the addition of small amounts of incomplete information into a repeated game can have dramatic effects, potentially enabling one party to obtain Stackelberg payoffs, as shown by Fudenberg and Levine [1989]. While our results are reminiscent of theirs, the theorems are very different.

Finally, in more directly related work, Lagunoff and Matsui [1997] consider the effect of changing the usual timing assumptions of repeated games. However, their prime focus is on the case where agents cannot change actions simultaneously, the opposite of what we focus on, and they have no switching costs. Despite this, their results on coordination games are similar to ours and use similar reasoning in some steps. In Lagunoff and Matsui [1997], in Gale [1995], and in our coordination game results, the key step is to show that if one player moves to the risk dominant and Pareto dominant outcome, he can force all subsequent play to that outcome. The models differ in what drives this conclusion, but the analysis given this fact is similar.¹

2 Model

Fix a finite normal form game with two² players, $G = (A, u)$ where A_i is i 's set of pure strategies for i and $u_i : A \rightarrow \mathbf{R}$ i 's payoff function, $i = 1, 2$. This game is finitely repeated during a finite time interval of length $M > 0$. The length of time between periods is denoted Δ , so the number of periods is M/Δ (hence all references to Δ should be understood to involve the assumption that M/Δ is an integer). Formally, for any Δ such that M/Δ is an integer, G_Δ is the game G repeated M/Δ times where the payoffs are taken to be total payoffs divided by Δ . That is, the payoffs in a given period are Δ times the payoff from the matrix. Let G_Δ^ε be the same game as G_Δ but where every change of action “costs” ε . The assumption that both players have the same switching cost is a normalization and hence is without loss of generality. Throughout, we number periods from the end, so period 1 is the last period, 2 is the next to last, etc. We use t to denote a period number.

Most of our results have the same basic form, asserting that “for almost all” ε in a certain range, if $\Delta < K\varepsilon$, then some particular outcome must result. The restriction to “almost all” ε is used only in uniqueness arguments to avoid certain ties in payoffs. In particular, we use this to ensure that we do not have a player indifferent between paying the switching cost and not doing so at a certain key juncture, potentially creating

¹Burdzy, Frankel, and Pauzner [1996], like Lagunoff and Matsui, consider a model without simultaneous changes of actions which generates risk dominance in coordination games. While their framework is much more complex than ours, we suspect that the driving force behind the results is related.

²Of course, one could consider games with switching costs and more than two players. However, we stick to the two player case in all of our results and so state this specifically in our definitions.

additional equilibria. To understand this restriction, recall that the game is played over the time interval $[0, M]$ but changes of action can only occur every Δ units of time. Given ε , there is a key length of time from the end, say ℓ^* , such that the agent would *strictly* prefer not changing his action when the time remaining is *strictly* less than ℓ^* . Intuitively, it would be surprising if the dates at which actions can be changed happened to be such that a decision is made when the time remaining is *exactly* ℓ^* . The restriction to almost all ε is used only to ensure that this does not happen. More formally, we remove a countable set of possible values of ε to guarantee that such ties never occur.

Also, even though the results state $\Delta < K\varepsilon$, we generally only prove that Δ must be sufficiently small given some fixed ε , not the more specific claim that Δ/ε must be sufficiently small. To see why “sufficiently small” must actually be below some constant times ε as stated in the theorem, suppose we have shown the result for pair of parameter values, say ε_1 and Δ_1 . Because our arguments are all based on backward induction, M is irrelevant and so the result holds for any M such that M/Δ_1 is an integer. Suppose we multiply ε_1 , Δ_1 , and M by a constant k . This cannot affect the result, of course, since this simply rescales all the payoffs in the game. Again, the irrelevance of M then means that the result must hold for the original M , switching costs of $k\varepsilon_1$, and a period length of $k\Delta_1$. Consequently, we see that the only relationship between ε and Δ that can be relevant is their ratio.

3 Commitment to a Fixed Action

As discussed in the introduction, when the length of a period is sufficiently small relative to the switching cost, neither player will change actions in the last period, regardless of the actions played up to that point. Hence in this sense, any equilibrium involves a certain commitment to fixed actions late enough in the game. The more interesting situation is where this commitment late in the game implies a commitment throughout the game. In this section, we give several results on games where the switching costs create such a credible commitment for one player and thus dramatically reduces the set of subgame perfect equilibria relative to the game without switching costs.

First, we give a surprisingly simple condition for this kind of commitment to be possible in a two player game. Afterward, we turn to some two-by-two games where some interesting tighter characterizations are possible.

3.1 A General Result

Our more general result gives a condition under which one player can ensure himself his largest feasible payoff in the game. We say that an action profile (a_1^*, a_2^*) is *weakly stable* if it uniquely maximizes the payoff of player i and satisfies

$$u_j(a_1^*, a_2^*) - u_j(a_i^*, a_j) \geq u_i(a_1^*, a_2^*) - u_i(a_i^*, a_j), \quad \forall a_j \in A_j, \quad j \neq i.$$

Intuitively, player i has no incentive to deviate from such a profile since he must do strictly worse at any other profile. Also, if j changes his action away from (a_1^*, a_2^*) , then this change affects him more than it affects i . Loosely, then, j seems to have little incentive to deviate. We use the modifier “weakly” to emphasize the idea that it is not obvious that such a profile is necessarily stable in any sense.

Before stating the implication of this property, we note four important facts about it. First, it may seem odd that we are comparing payoff differences *across* agents. The reason such interpersonal comparisons are meaningful here is because we have, effectively, normalized payoffs by assuming the two players have the same switching cost. If we removed this assumption, the appropriate restatement of the condition above is that the change in player j 's payoff *relative to* his switching cost exceeds the change in i 's payoff relative to i 's switching cost.

Second, if (a_1^*, a_2^*) is weakly stable, then it is a Nash equilibrium of the stage game. By definition, any movement away from this profile makes player i worse off. Hence it is obviously true that his best reply to a_j^* is a_i^* . This assumption also implies that $u_i(a_1^*, a_2^*) - u_i(a_i^*, a_j) > 0$ for any $a_j \neq a_j^*$, so it must be true that $u_j(a_1^*, a_2^*) > u_j(a_i^*, a_j)$ for all $a_j \neq a_j^*$. Hence a_j^* is the unique best reply for j to a_i^* , so (a_1^*, a_2^*) is a (strict) Nash equilibrium. In particular, this also implies that $u_i(a_1^*, a_2^*)$ is i 's Stackelberg payoff.

Third, the assumption that (a_1^*, a_2^*) is weakly stable does *not* imply that (a_1^*, a_2^*) is the only Nash equilibrium of the one-shot game. This condition says nothing whatsoever about payoffs when player i plays an action different from a_i^* , so there could be many Nash equilibria and j 's payoff could be higher than $u_j(a_1^*, a_2^*)$ in some or all of them.

Finally, it is not hard to show that any game has at most one weakly stable profile. To see this, suppose that (a_1^*, a_2^*) and (a_1', a_2') are both weakly stable. Since a weakly stable profile must be the unique best profile for one player, we can assume that (a_1^*, a_2^*) maximizes 1's payoff and (a_1', a_2') maximizes 2's payoff. Weak stability of (a_1^*, a_2^*) then requires

$$u_2(a_1^*, a_2^*) - u_2(a_1^*, a_2') \geq u_1(a_1^*, a_2^*) - u_1(a_1^*, a_2').$$

Because $u_2(a_1', a_2') > u_2(a_1^*, a_2^*)$ and $u_1(a_1^*, a_2^*) > u_1(a_1', a_2')$, this implies

$$u_2(a_1', a_2') - u_2(a_1^*, a_2') > u_1(a_1', a_2') - u_1(a_1^*, a_2').$$

However, weak stability of (a'_1, a'_2) requires the opposite inequality.

We have the following surprisingly strong result.

Theorem 1 *Suppose (a_1^*, a_2^*) is a weakly stable profile in G . Then there is an $\bar{\varepsilon} > 0$ and $K > 0$ such that for almost every $\varepsilon \in (0, \bar{\varepsilon})$, for all $\Delta \in (0, K\varepsilon)$, the unique subgame perfect equilibrium outcome of G_Δ^ε is (a_1^*, a_2^*) every period.*

To see the intuition for this result, consider the following example.

	L	M	R
U	5, 3	4, 1	3, 0
D	2, 5	1, 6	4, 10

In this game, the profile (U, L) uniquely maximizes 1's payoff. Also, if 2 changes his action from L to either M or R against U , the reduction in his payoff is larger than the reduction in 1's payoff (2 loses 2 moving from L to M while 1 loses only 1 and 2 loses 3 moving from L to R while 1 loses only 2). Hence (U, L) is weakly stable. Note that (U, L) is a Nash equilibrium, but (D, R) is as well and the latter is much preferred by 2. However, (D, R) is not weakly stable because if 1 changes actions to U , his payoff falls by only 1, while 2's falls by 10.

So suppose that the length of a period is short relative to the switching cost, so that neither player will change actions in the last period, regardless of the actions played before this. Then we know that there is at least one period t (namely the last one, $t = 1$) in which player 1 is committed to U at t if he played U at $t + 1$. That is, in any subgame perfect equilibrium, if 1 played U at the next to last period, he must play U in the last period. We now show that this sets up an induction which ensures that this is true throughout the game.

To see this, suppose we are at period $t \geq 2$. Suppose from period $t - 1$ onward, if 1 played U at the preceding period, he will never change actions. Suppose 1 played U at period $t + 1$. We wish to show that 1 must play U at period t in any subgame perfect equilibrium, giving us the desired induction. There are two cases to consider. First, suppose that if 1 sticks with U at period t , 2 will definitely play L from period $t - 1$ onward. Given the induction hypothesis, 2 will certainly do this if he plays L at t or if the cost of switching from whatever he did play is not too large relative to the gains to be had. In this case, it is clear that 1 should not change actions at t . If he stays with U , then his payoff must be at least $3\Delta + 5(t - 1)\Delta$. If he changes actions, his payoff certainly cannot be better than $4\Delta + 5(t - 1)\Delta - \varepsilon$. This calculation treats 1 as being able to get 5 from $t - 1$ onward even without paying the cost of switching back to U ,

so it is certainly an overestimate of his payoff! Clearly, if $\Delta < \varepsilon$, 1 is better off sticking with U in period t .

So suppose that 2 will not play L from period $t - 1$ onward even if 1 plays U at t . Because of the induction hypothesis, we know that 1 is certainly not going to change actions from $t - 1$ onward. Hence it must be true that 2 plays either M or R at t and does not find it worthwhile to change actions thereafter. For concreteness, suppose 2 plays M at t (an analogous argument covers the case where he plays R at t). So the fact that 2 won't change actions at period $t - 1$ says that $(t - 1)\Delta > 3(t - 1)\Delta - \varepsilon$ or $\varepsilon > 2(t - 1)\Delta$. So consider player 1's situation at period t . If he continues with U , his payoff is $4t\Delta$ because neither he nor player 2 will change actions from $t - 1$ onward. If he switches to D , his payoff is certainly less than $5t\Delta - \varepsilon$. Hence he certainly will not switch if $4t\Delta > 5t\Delta - \varepsilon$ or $\varepsilon > t\Delta$. We saw above that we must have $\varepsilon > 2(t - 1)\Delta$. Hence this inequality holds if $2(t - 1)\Delta \geq t\Delta$ or $t \geq 2$, which is true by hypothesis.

More intuitively, suppose 1 will be committed to U from next period onward if he plays it today. If this commitment is sufficient to force player 2 to play L from tomorrow onward, then player 1 won't switch away from U today. Hence he is committed to it from today onward. If this commitment is not sufficient to force player 2 to play L from tomorrow onward, it must be true that 2's gain to switching to L is not large enough to make a change of action worthwhile. But 1's gain from changing actions must be less than the hypothetical gain calculated as if a change of action would give him 5 from now on. Hence weak stability tells us that if 2 won't find a switch profitable, 1 won't either. Hence, again, 1 is committed to U from today onward. Either way, we get an induction which tells us that if 1 ever plays U , he will always play U thereafter. Hence he will start with U , making (U, L) every period the unique subgame perfect equilibrium outcome.

3.2 Two by Two Games

In certain two-by-two games, we obtain an interesting tighter characterization of situations where commitment is possible. For this case, we let $A_i = \{L, R\}$, $i = 1, 2$, and denote the payoffs in the stage game G by

$$\begin{array}{cc} & L & R \\ L & a_1, a_2 & d_1, c_2 \\ R & c_1, d_2 & b_1, b_2 \end{array}$$

We consider two more specific formulations. In both, there are two Nash equilibria in the stage game, (L, L) and (R, R) . Hence we assume $a_i > c_i$ and $b_i > d_i$, $i = 1, 2$. Second, we assume (L, L) is the profile which uniquely maximizes 1's payoff so $a_1 > b_1$.

We say that the game is a *symmetric³ coordination game* if $a_1 = a_2 = a$, $b_1 = b_2 = b$, etc. In such games, the Nash equilibria are Pareto ranked, with both players preferring the (L, L) equilibrium. In such a game, (L, L) is said to be *risk dominant* if it is the best reply to a 50–50 mixture by the opponent — that is, if $a - c > b - d$. (R, R) is risk dominant if the opposite strict inequality holds.

We say that the game is a *generic Battle of the Sexes game* if $a_2 < b_2$ and $a_2 - c_2 \neq b_1 - d_1$. That is, Battle of the Sexes games are where the equilibria are not Pareto ranked, with player 1 preferring (L, L) and 2 preferring (R, R) . Note that in this case, b_2 must be 2's largest payoff in the matrix as $b_2 > a_2 > c_2$ and $b_2 > d_2$. We call the payoff difference $b_1 - d_1$ 1's *incentive to blink*. Similarly, we refer to $a_2 - c_2$ as 2's *incentive to blink*. To understand the terminology, imagine each player is playing the action that he uses in his preferred equilibrium. Thus we are at the action profile (L, R) with payoffs (d_1, c_2) . If player 1 “gives in” and moves to player 2's preferred equilibrium, he gains $b_1 - d_1$. If, instead, 2 gives in and moves to 1's preferred equilibrium, he gains $a_2 - c_2$. Hence these payoff differences can be thought of as measuring a player's incentive to surrender in such a confrontation. By our definition, in a generic Battle of the Sexes game, the players cannot have equal incentives to blink.

It is easy to see that if G is either a symmetric coordination game or a generic Battle of the Sexes game, G_Δ has many equilibria. In particular, any (rational) convex combination of the payoffs to (L, L) and (R, R) can be achieved by a subgame perfect equilibrium of G_Δ for Δ sufficiently small. In addition, there are equilibria in which (L, R) or (R, L) are played for many periods.

On the other hand, we have:

Theorem 2 *Assume G is a symmetric coordination game. Then if (L, L) is risk dominant, there is an $\bar{\varepsilon} > 0$ and $K > 0$ such that for almost every $\varepsilon \in (0, \bar{\varepsilon})$, for all $\Delta \in (0, K\varepsilon)$, the unique subgame perfect equilibrium outcome of G_Δ^ε is (L, L) every period. However, if (R, R) is risk dominant, then there is a subgame perfect equilibrium in which (L, L) is played every period and one in which (R, R) is played every period.*

This result is *not* a corollary to Theorem 1. One obvious reason it is not a corollary is that Theorem 2 is a statement about when the equilibrium is *not* unique, not just when it is unique. Even the uniqueness part of the theorem is not a corollary of Theorem 1,

³Similar results hold for asymmetric coordination games, but the conditions needed are less straightforward.

however. For example, in the symmetric coordination game

	L	R
L	4, 4	0, 1
R	1, 0	2, 2

(L, L) is risk dominant, but is not weakly stable. In fact, this is a tighter result (for this class of games) than Theorem 1 gives. Notice that (L, L) is weakly stable if $a - c \geq a - d$. But $a > b$, so this implies $a - c > b - d$, so that (L, L) is risk dominant. Hence weak stability implies risk dominance, but, as the example above shows, the reverse is not true.

We also have:

Theorem 3 *Assume G is a generic Battle of the Sexes game. Then there is an $\bar{\varepsilon} > 0$ and $K > 0$ such that for almost every $\varepsilon \in (0, \bar{\varepsilon})$, for all $\Delta \in (0, K\varepsilon)$, G_{Δ}^{ε} has a unique subgame perfect equilibrium outcome. If 2's incentive to blink is larger than 1's, the unique outcome is (L, L) every period. If 1's incentive to blink is larger than 2's, then the unique outcome is (R, R) every period.*

Again, the result is not a corollary to Theorem 1. It is not hard to show that if 1's preferred equilibrium is weakly stable, then his incentive to blink is higher than 1's and similarly for 2's preferred equilibrium. However, in

	L	R
L	5, 2	0, 0
R	0, 0	1, 5

2's incentive to blink is higher than 1's, but neither player's preferred equilibrium is weakly stable. Hence, again, the result is a tighter characterization for this class of games than Theorem 1 yields.

Full proofs of these results are contained in the Appendix, but it is not hard to see the basic intuition. First, consider the symmetric coordination game given above where (L, L) is risk dominant. Analogously to the proof of Theorem 1, the key is that either player can credibly commit himself to playing L for the entire game. To see this, note that if Δ is small relative to ε and we are in one of the last two periods (that is, at $t \leq 2$), then neither player will switch actions because the gains cannot possibly exceed the switching costs. That is, at $t \leq 2$, if a player played L in the previous period, then he must play L at t in any subgame perfect equilibrium.

This gives the starting point of an induction. So fix some $t \geq 3$ and suppose we know that from $t - 1$ onward, if player i played L at the preceding period, then he must play

L for the rest of the game in any subgame perfect equilibrium. We now show that the same must be true at t . So suppose we are at period t and i played L at the previous period. If he plays L again at t , then by the induction hypothesis, he will always play L thereafter. As with the proof of Theorem 1, there are two cases to consider: where the other player will also play L from $t - 1$ onward and where he will not. In the former case, it is easy to see that i should play L at t . If he does so, his payoff from t onward cannot be worse than $4(t - 1)\Delta$, while if he switches to R at t instead, his payoff certainly must be strictly worse than $2\Delta + 4(t - 1)\Delta - \varepsilon$. If Δ is small relative to ε , $2\Delta - \varepsilon < 0$, so i will play L at t , as claimed.

So let us turn to the other case: where j will not play L from $t - 1$ onward even if i plays L at t . Because of the induction hypothesis, we know that i will not change actions. Hence if j will not play L , it must be because the cost of switching exceeds the gains he can incur over the last $t - 1$ periods. That is, we must have $4(t - 1)\Delta - \varepsilon < (t - 1)\Delta$ or $\varepsilon > 3(t - 1)\Delta$. In this situation, if i plays L at t , his payoff from t onward is 0 at worst (since he must get at least 0 at t and will definitely get 0 from $t - 1$ onward). If, instead, i switches to R at t , he will never again switch back. To see this, note that his largest possible gain from switching from R to L at t is $2(t - 1)\Delta$. We know from the above that $\varepsilon > 3(t - 1)\Delta$, so the gain could not be worth the cost. Hence if i switches to R at t , his payoff from then onward is $2t\Delta - \varepsilon$ at best. But $\varepsilon > 3(t - 1)\Delta$ implies that this is negative whenever $3(t - 1)\Delta \geq 2t\Delta$ or $t \geq 3$, which is true by hypothesis.

Put more intuitively, if playing L at t commits player i to this action from then on, he has no incentive to switch away from L just before this time. Hence he is effectively committed to L earlier than t . By induction, this means that playing L in the very first period commits him to always playing L , a commitment he will wish to take advantage of since this leads to his highest possible payoff.

The result for Battle of the Sexes is based on a similar intuition. Here whichever player has the smaller incentive to blink is the one who can commit himself. To see why *only* this player can commit, consider the Battle of the Sexes payoffs in the example above. Recall that 2's incentive to blink is 2, while 1's incentive to blink is 1. Hence, we claim, player 1 can commit himself to L but player 2 cannot commit himself to R . Very late in the game, neither player will be willing to change actions because there is not enough time left for the change to yield enough gains to make it worth the costs. Hence late in the game, *both* players are committed to not changing action.

The more interesting phase of the game is where the length of time left, $t\Delta$, is between ε and $\varepsilon/2$. At such a point, if a change of action yields a gain of 1 each period, it is not worth doing because $t\Delta < \varepsilon$, while a change of action yielding a gain of 2 each period is worthwhile since $2t\Delta > \varepsilon$. So fix such a t and suppose that (L, R) was played in the previous period. Note that player 1 cannot possibly gain more than 1 per period from t

onward by switching to R since this is the largest difference between any of his payoffs from playing R and any of his payoffs from playing L . Hence he will certainly not switch actions at t or any subsequent date. Given this, if player 2 does not switch to L , his payoff will be 0, while switching to L earns $2t\Delta - \varepsilon$. Hence he must switch to L in any subgame perfect equilibrium. Therefore, player 2 certainly cannot commit himself to playing R and not changing in this interval. However, player 1 is committed to playing L : if he played L at the preceding period for any period in this interval, he is committed to never switching from this action. Just as above, this commitment works backward to the beginning of the game, ensuring that 1 is able to force the (L, L) outcome in every period.

It is natural to wonder why we get multiple equilibria in some symmetric coordination games. The key to understanding the risk dominance condition is that the risk dominant profile is the one which either player can commit himself to in the “middle” phase of the game. That is, we know no player would change actions very late in the game so any action, once played, is an action the player is committed to. Earlier, though, this effect will continue for only one of the two actions — specifically, the action in the risk dominant profile. Hence there will necessarily be a point late in the game at which any “mismatch” of actions leads the players to the risk dominant outcome. That is, analogously to the reasoning above for the Battle of the Sexes, in this critical phase of the game, if one played L and the other played R in the previous period, the one who used the risk dominant action will never change actions, so the other player must switch. When risk dominance and Pareto dominance coincide, each player has an incentive to play the risk dominant action in order to achieve precisely this effect. When risk dominance and Pareto dominance differ, each has an incentive to *avoid* the risk dominant action to avoid this effect. On the other hand, if the opponent is expected to play the risk dominant action, there is no gain to avoiding it oneself, so there are multiple equilibria in this case. Put differently, the players can commit themselves in a certain phase of the game to (R, R) but have no incentive to do so.

Hence we find the surprising conclusion that it is easier to get uniqueness when the players have different preferences over equilibria than when they agree on which equilibrium is best. One way to understand this is to note that the introduction of switching costs tends to favor certain outcomes because of the way switching costs create a commitment to some profiles earlier than for others. When players agree on which profile is best, it may be that the switching costs favor a different outcome. In this case, we cannot obtain uniqueness. On the other hand, when the players disagree, one of the players must have an incentive to exploit the effect created by the switching costs, so we do get uniqueness.

It is worth noting that these results do not require a “large” deviation from the usual finitely repeated game model. Returning to the coordination game payoffs used above for

illustrative purposes, it is not hard to show that the unique subgame perfect equilibrium outcome is (L, L) in every period whenever $\varepsilon > 2\Delta$. In other words, we only require that periods are short enough that a change of action which increases one's payoff from the worst Nash equilibrium payoff (2) to the best (4) but does so only for a single period is not worthwhile.

4 Cooperation in the Prisoners' Dilemma

In a way, it is unsurprising that switching costs might allow cooperation in the finitely repeated Prisoners' Dilemma. After all, as we have emphasized, when Δ is small relative to ε , no player will change actions in the last period. Hence the first step of the usual backward induction argument does not follow, so the standard proof that no one cooperates breaks down.

On the other hand, we know that cooperation can only occur if the threat of punishment for defection is credible. While switching costs aid cooperation in the way they reduce the incentive to move away from mutual cooperation, they make cooperation more difficult in the way they reduce the incentive for moving away from cooperation to punish defection. As we show in this section, whether or not cooperation is possible in the Prisoners' Dilemma with small switching costs is completely determined by which of these two effects is the larger.

In this section, we consider the Prisoners' Dilemma with payoffs

$$\begin{array}{cc} & C & D \\ C & a, a & d, c \\ D & c, d & b, b \end{array}$$

where $c > a > b > d$.⁴ We call $c - a$ the *incentive to cheat* and $b - d$ the *incentive to punish*.

In line with the intuition suggested above, we will show that mutual cooperation can be sustained in a subgame perfect equilibrium for small ε and Δ if and only if the incentive to punish is larger than the incentive to cheat. On the other hand, regardless of which incentive is larger, mutual defection is always an equilibrium.

Theorem 4 *Suppose G is a Prisoners' Dilemma game. If the incentive to punish is strictly greater than the incentive to cheat, then there is a $K > 0$ such that for all*

⁴It is not difficult to generalize this result to asymmetric versions of the Prisoners' Dilemma but it adds little.

$\varepsilon \in (0, M(c - a))$, for all $\Delta \in (0, K\varepsilon)$, there is a subgame perfect equilibrium where both players cooperate in every period.

Proof. Fix ε in the range specified and any Δ small enough that there are dates t such that

$$(c - a)t\Delta < \varepsilon \leq (b - d)t\Delta.$$

Construct a strategy for player i as follows. He begins by cooperating and cooperates in any period in which both players cooperated the previous period. If either player defected in the previous period, then for any t such that $(b - d)t\Delta \geq \varepsilon$, i defects, switching to this action if need be. Finally, for any later period, i does whatever he did the previous period.

To see that it is an equilibrium for each player to follow this strategy, let us verify that i 's strategy is optimal given any history. First, consider a history such that $(b - d)t\Delta < \varepsilon$. At this point, the opponent is expected to never change strategies again. Hence the optimal strategy for i must be to play one fixed action for the rest of the game. If i defected in the previous period, the dominant strategy property obviously implies i should not change actions. If i cooperated in the previous period, it is optimal to stick with cooperation as long as either $at\Delta \geq ct\Delta - \varepsilon$ or $dt\Delta \geq bt\Delta - \varepsilon$, depending on whether the opponent is cooperating. But by assumption, $(a - c)t\Delta < (b - d)t\Delta < \varepsilon$, so both inequalities hold. Hence it is never optimal to change actions at such a period.

Now suppose $(b - d)t\Delta \geq \varepsilon$. There are three relevant cases here. First, suppose either player defected in the previous period. Then the opponent is expected to defect from this point onward. If i defected in the previous period, it is clearly optimal to continue defecting. If i player cooperated in the previous period, it is optimal to switch to defecting as $(b - d)t\Delta \geq \varepsilon$. Hence the specified strategy is optimal.

Second, suppose both players cooperated in the previous period and that $(b - d)t\Delta \geq \varepsilon > (b - d)(t - 1)\varepsilon$. Player i expects the opponent to cooperate from now on regardless of what he does at t . Hence i should either cooperate from this period onward or defect from this period onward. The former is better iff $t\Delta c - \varepsilon \leq t\Delta a$ or $(c - a)t\Delta \leq \varepsilon$. But since t is the last period such that $\varepsilon \leq (b - d)t\Delta$, our assumption on Δ implies $(c - a)t\Delta < \varepsilon$, so this holds.

Third, suppose we are at a period t such that both players cooperated in the previous period but $(b - d)(t - 1)\Delta \geq \varepsilon$. Then if i cooperates at t , his payoff will be $at\Delta$, while if he defects, his payoff is $c\Delta + b(t - 1)\Delta - \varepsilon$. Hence cooperation is optimal if

$$\varepsilon \geq [c - a - (t - 1)(a - b)]\Delta.$$

If the term in brackets on the right is negative, this must hold. If it is positive, then this holds for Δ sufficiently small.

To conclude, consider the first period of the game. If i cooperates, his equilibrium payoff will be aM , while if he defects, his payoff is $c\Delta + (M - \Delta)b$. Since the latter converges to $bM < aM$ as Δ goes to zero, we see that it is optimal for i to begin by cooperating if Δ is sufficiently small. Hence for small Δ , these strategies form a subgame perfect equilibrium. ■

So mutual cooperation can be supported as a subgame perfect equilibrium outcome if the incentive to punish exceeds the incentive to cheat. On the other hand, it is easy to see that mutual defection can *always* be supported. To see this, construct an equilibrium as follows. Player i defects in every period (including the first) *unless* both agents cooperated in the previous period and $(c - a)t\Delta < \varepsilon$ or i alone cooperated in the previous period and $(b - d)t\Delta < \varepsilon$. To see that this is an equilibrium, first note that it is clearly optimal to cooperate under the circumstances specified for cooperation. So consider any other history. If both players defected in the previous period, it is clearly optimal to continue with defection since the opponent is expected to always defect thereafter. If i cooperated in the previous period and $t\Delta$ is large enough, then he expects his opponent to defect at t and thereafter. Hence it is optimal for him to switch to defection. Thus these strategies form a subgame perfect equilibrium.

The slightly more difficult result, proven in the Appendix, is that if the incentive to cheat exceeds the incentive to punish, then mutual defection is the unique subgame perfect equilibrium outcome. More specifically,

Theorem 5 *If the incentive to cheat is larger than the incentive to punish, then for all Δ and almost all ε , the unique subgame perfect equilibrium outcome of G_Δ^ε is mutual defection in every period.*

To understand these results, first note that, regardless of whether which incentive is larger, mutual cooperation is stable sufficiently late in the game. That is, if there is sufficiently little time remaining and both players cooperated in the previous period, then both will cooperate for the rest of the game. Thus as noted at the outset, the usual backward induction arguments do not apply. When the incentive to punish is larger than the incentive to cheat, then whenever it is worthwhile to cheat if one can get away with it — that is, whenever $t\Delta(c - a) > \varepsilon$ — then it must also be worthwhile to punish a cheater as $t\Delta(b - d) > t\Delta(c - a) > \varepsilon$. Hence punishment is credible as long as there is any possible gain to cheating. On the other hand, when the incentive to cheat is larger than the incentive to punish, there will necessarily be a point in the game where deviation from cooperation *cannot* be punished in equilibrium. As a result, each player has an incentive to switch to defection just before this point. Once we know that there is a date at which both players will defect, the usual backward induction reasoning applies and shows that both will always defect.

5 More Complex Commitments

In this section, we show how switching costs can make some much more complex commitments credible, leading to some very surprising results. In particular, we give two examples of games with a unique equilibrium outcome with no switching costs and a different unique equilibrium outcome with such costs. In one example, the unique equilibrium outcome requires actions to be changed along the equilibrium path, despite the switching costs. In the other example, the unique outcome with switching costs is independent of the exact level of the switching cost, as long as it is small. As we explain in Remark 1, this result demonstrates a failure of both upper and lower semicontinuity of the equilibrium correspondence, though this failure is more subtle than it may appear.

For both examples, the stage game G is

	L	R
U	a_1, a_2	d_1, c_2
D	c_1, d_2	b_1, b_2

In both, player 2 has a dominant strategy of L , so $a_2 > c_2$ and $d_2 > b_2$. In both, player 1's best reply to L is U or $a_1 > c_1$, making (U, L) the unique Nash equilibrium of the stage game. Hence in the finitely repeated game with no switching cost, the unique subgame perfect equilibrium outcome is to repeat (U, L) every period. The key to both examples is that player 2 has an incentive to try to force player 1 to play D instead of U . More specifically, d_2 is 2's largest possible payoff.

While there are important differences between the two examples, the reason 2 is able to move play away from (U, L) is similar in both. First, if 2 moves to R , he has little incentive to move back if 1 plays U . That is, $a_2 - c_2$ will be small. In this sense, 2's threat to not move back to L is credible late in the game. Second, if 1 responds to 2's move by changing to D , 2's incentive to return to L will be large. That is, $d_2 - b_2$ will be large. Third, because it will induce 2 to move to L , 1 will gain by moving to D in this situation. In other words, $c_1 - d_1$ is sufficiently large. Because this threat is credible, it induces 1 to play D . Finally, 2 has a large incentive to force play from (U, L) to (U, R) even though he will have to change actions twice to get to (D, L) . That is, $(d_2 - a_2)/2$ is sufficiently large. Hence even when (U, L) is being played, at a certain point, 2 can credibly threaten to switch to R , stay there as long as 1 stays at U , and then switch to L once 1 switches to D .

More specifically, both examples assume

$$d_2 - b_2 > a_2 - c_2, \quad a_1 - c_1, \quad |b_1 - d_1|, \quad c_1 - d_1, \quad (1)$$

$$c_1 - d_1 > a_2 - c_2, \quad a_1 - c_1, \quad (2)$$

and

$$\frac{(d_2 - a_2)}{2} > c_1 - d_1. \quad (3)$$

The key difference between the two examples is the effect this has on 1. Suppose (U, L) were played in the previous period and suppose that 1 knows that if (U, L) is played one more time, 2's threat to switch to R and force play to (D, L) will become credible in the next period. Note that once 2 moves to R , 1 is best off moving to D in the same period to speed the movement to (D, L) . In light of this, should 1 go ahead and play D now or wait? By waiting, 1 gets a_1 in the current period, b_1 in the next, and c_1 thereafter, minus ε for the change of actions. By moving now, 1 gets c_1 from the current period onward minus ε . Hence waiting is better if $a_1 + b_1 > 2c_1$, while moving now is better if the reverse strict inequality holds. The difference between the two examples is that we assume that waiting is better in the first, while moving is better in the second. As we will see, this difference means that 2's threat is credible for the entire game in one case but only for the last part of the game in the other case. In the former situation, then, (D, L) will be played the entire game, while in the latter, it is only played toward the end.

5.1 First Example

As stated above, here we assume that L is strictly dominant for player 2 and that U is 1's unique best reply to L . Also, d_2 is 2's largest payoff in the matrix. Finally, we assume (1), (2), (3), and that $a_1 + b_1 > 2c_1$.

As mentioned earlier, the unique subgame perfect equilibrium without switching costs is to repeat (U, L) every period. In stark contrast, the result with switching costs is:

Theorem 6 *Given the payoff assumptions above, there is an $\bar{\varepsilon} > 0$ and $K > 0$ such that for almost every $\varepsilon \in (0, \bar{\varepsilon})$, for all $\Delta \in (0, K\varepsilon)$, G_Δ^ε has a unique subgame perfect equilibrium. In this equilibrium, (D, L) is played every period.*

To see the intuition, consider the following payoff matrix.

	L	R
U	5, 2	0, 0
D	3, 10	4, 4

It is easy to see that these payoffs satisfy all of our assumptions for this subsection. Note that the largest unilateral payoff gain is when 2 moves from (D, R) to (D, L) . Hence

there is a phase of the game where this is the only change of actions any player would undertake. As we work backward from this phase, the next change of actions a player would consider is that at t such that

$$4\Delta + 3(t-1)\Delta > \varepsilon,$$

when 1 would be willing to move from (U, R) to (D, R) , knowing that this will lead 2 to move to (D, L) .

Unfortunately, at this point, the analysis becomes more complex. To give a somewhat misleading intuition first, let us simply suppose that in this phase of the game, 1 would move to (D, R) from (U, R) . Note that at this point, for Δ sufficiently small, we must have $3t\Delta > \varepsilon$ or $6t\Delta > 2\varepsilon$. This implies that 2 would move from (U, L) to (U, R) to start this process. This would end up leading 2 to change actions twice, but the gain of (approximately) $8t\Delta$ exceeds the 2ε switching cost. Hence we see that if (U, L) were played, we will end up moving to (D, L) . This means that if we are at (D, L) , even very early in the game, 1 will not change actions to move to his apparently better action of U . While this would give him a one period payoff gain of 1, he would end up having to switch actions again soon to complete the move back to (D, L) and so this would not be worthwhile. Anticipating this, 1 finds it optimal to begin the game with D , rather than U , to avoid incurring switching costs. In short, 2 is able to credibly threaten that if 1 switches from D to U , he will then retaliate with R and remain at this action until 1 switches back to D .

This intuition is misleading in one respect: The movement from (U, L) to (D, L) *must* involve mixed strategies. However, one can show that as Δ gets small, the sequence of actions described above must occur “soon” with “large” probability.

Notice that Theorem 6 implies that 2’s payoff is higher than his Stackelberg payoff. If $d_1 > b_1$, then 1’s best reply to R is U , so that 2’s Stackelberg payoff would be a_2 . If $b_1 > d_1$, 2’s Stackelberg payoff would be the larger of a_2 and b_2 . However, by assumption, $d_2 > \max\{a_2, b_2\}$. Hence 2’s payoff in this equilibrium exceeds his Stackelberg payoff, precisely because the switching costs enable 2 to make credible a much more serious threat than the threat to not change actions.

To see this in the matrix above, note that 2’s Stackelberg payoff is 4 in that matrix, while he gets a payoff of 10 in the equilibrium. Intuitively, 2 is able to do better than the Stackelberg payoff because he is able to force 1 away from his myopic best-reply. He is able to do this by credibly threatening to move from (U, L) to (U, R) , forcing 1 to earn a lower payoff until he incurs the switching cost necessary to move play to (D, L) .

In the particular numerical example above, it is also true that 1’s payoff is below his minmax. It is not difficult to show that 1’s minmax payoff is $10/3$, while his payoff in the unique equilibrium is 3. On the other hand, this minmax is achieved only in mixed

strategies and this randomization every period by 1 would lead to a high probability of many changes of action and hence is costly. The more natural comparison is to 1's pure maxmin (that is, where he must choose his strategy "first" and must choose a pure strategy). It is easy to see no player can do worse than his pure maxmin payoff. In this example, player 1's pure maxmin payoff is 3, precisely his equilibrium payoff.

Remark 1 The result of Theorem 6 demonstrates a violation of both upper and lower semicontinuity of the equilibrium correspondence in ε and Δ . The result implies that we can take a sequence of (ε, Δ) converging to $(0, 0)$ satisfying the conditions of the theorem such that the set of equilibrium payoffs converges to the singleton $\{(3, 10)\}$. On the other hand, consider the game when $\varepsilon = \Delta = 0$. While it is clear what it means to define the game for $\varepsilon = 0$, it is less immediate how our definitions apply to the case where $\Delta = 0$. We believe the most natural definition of the game at this point is the infinitely repeated game where payoffs are evaluated by the limit of means criterion. To see why, fix $\varepsilon = 0$ and consider the effect of taking Δ to zero. Payoffs for any $\Delta > 0$ are defined to be the average payoffs when the game is repeated $T = 1/\Delta$ periods. It seems natural then to take the limiting payoff function to be the limit of these average payoffs. Given this definition of the game at the point $\varepsilon = \Delta = 0$, it is clear that the set of equilibrium payoffs at this point is the full feasible, individually rational set. Since player 1's minmax payoff is $10/3$, the point $(3, 10)$ is not contained in this set. Hence the set of payoffs at the limit is disjoint from the limiting set of payoffs, indicating that the equilibrium correspondence is neither upper nor lower semicontinuous.

5.2 Example 2

In this subsection, we make the same assumptions as for Example 1 except that we now assume that $2c_1 > a_1 + b_1$ in place of the reverse strict inequality assumed above. Here the unique equilibrium requires changing actions along the equilibrium path.

Theorem 7 *Given the payoff assumptions above, there is an $\bar{\varepsilon} > 0$ and $K > 0$ such that for almost every $\varepsilon \in (0, \bar{\varepsilon})$, for all $\Delta \in (0, K\varepsilon)$, there is a unique subgame perfect equilibrium. In this equilibrium, (U, L) is played in every period up until period $t^*(\varepsilon, \Delta) + 2$ where $t^*(\varepsilon, \Delta)$ is the smallest t such that*

$$c_1(t-1)\Delta + b_1\Delta - d_1t\Delta > \varepsilon.$$

At period $t^(\varepsilon, \Delta) + 2$ onward, 1 changes actions to D and (D, L) is played in all succeeding periods.*

Note that as $\Delta \downarrow 0$, $(t^* + 2)\Delta$ converges to $\varepsilon/(c_1 - d_1)$. Hence the limiting outcome as Δ goes to zero has (U, L) played from the beginning till a length of time of $\varepsilon/(c_1 - d_1)$ left to go, at which point play moves to (D, L) for the rest of the game. As $\varepsilon \downarrow 0$ as well, the length of this second phase goes to zero. Hence the outcome here is “close” to the outcome in the game without switching costs when ε and Δ are small.

To understand the intuition of this result, consider the following payoffs:

	L	R
U	$5, 2$	$0, 0$
D	$3, 10$	$-1, 4$

It is not hard to see that this matrix satisfies the assumptions of this subsection. Just as in the previous subsection, it is not hard to see that there is a phase late in the game where the only change of actions is from (D, R) to (D, L) . Again, as we back up from this phase, we reach a phase where 1 would move from (U, R) to (D, R) . The t^* defined in Theorem 7 is precisely the last date at which 1 would consider this move.

As with Theorem 6, a precise description of the equilibrium requires considering mixed strategies. For simplicity, then, what follows is an imprecise intuition regarding what happens in this phase of the game. At period t^* , if (U, R) were played in the previous period, we move to (D, R) (with high probability) and then onto (D, L) in the following period. So consider the previous period, $t^* + 1$. Suppose (U, L) were played in the period before, $t^* + 2$. Then 2 will move across to (U, R) to trigger the movement to his favorite outcome. Anticipating this, 1 will simultaneously move to D , so the outcome shifts to (D, R) the next period and then onto (D, L) where it will remain the rest of the game.⁵ In other words, 2 will trigger his threat to move to R until 1 moves to D .

Next consider period $t^* + 2$ and suppose (U, L) were played at the preceding period. We know that 2’s threat to move to R the next period if 1 does not change to D is credible from the above analysis. Hence if 1 does not change actions, we see that the outcome will move to (D, R) the next period and (D, L) thereafter, giving 1 a payoff of

$$5\Delta - \Delta + 3t^*\Delta - \varepsilon = 4\Delta + 3t^*\Delta - \varepsilon.$$

On the other hand, if 1 moves to D at this period, his payoff is

$$3(t^* + 2)\Delta - \varepsilon = 6\Delta + 3t^*\Delta - \varepsilon.$$

Hence 1 should give in to 2’s threat and move to D at period $t^* + 2$. In light of this, 2 will remain at L .

Hence we see that play must move from (U, L) to (D, L) at $t^* + 2$, as asserted in the theorem. It is not hard to use this to show that play must be at (D, L) at period $t^* + 2$

⁵To be more precise, the outcome described here is what the mixed strategies converge to as $\Delta \downarrow 0$.

no matter what actions were played in the preceding period. This conclusion is stronger than what we obtained in the previous example where we could only conclude that play must move to (D, L) “soon” with “high” probability. However, 2’s apparently greater ability to force the outcome to (D, L) unravels as we work backward from $t^* + 2$.

To see the point, consider period $t^* + 3$ and suppose that (U, L) were played in the preceding period. Now player 2 knows that he will play L in the next period, so he certainly will not pay to switch to R , knowing he would have to pay to switch back again tomorrow. Hence he will certainly play L at $t^* + 3$. But then 1 has no incentive to switch to D early. By staying at U , he gets a payoff of 5 one more time before having to switch to getting 3. Hence we remain at (U, L) . In other words, because play *will* move to (D, L) at the next period, 2 cannot credibly threaten to play R at this period. Hence he cannot force the outcome away from (U, L) at $t^* + 3$. Put differently, because his threat is both credible at a certain point in time and very powerful, it is not credible a moment earlier.

Similar reasoning shows that if (U, R) were played in the previous period, then we move to (U, L) at $t^* + 3$. To see this, note that 2 has no incentive not to move to L now. He knows he will move to L in the next period regardless and so may as well switch now to get the best possible payoff in this period. In light of this, again, 1 will not switch to D until the next period.

This reasoning continues backward. At any earlier period, if 1 played U in the previous period, he has no incentive to change his action before $t^* + 2$. Hence 2 will play L , switching to this action if need be. In short, in any earlier period, if 1 played U in the previous period, they play (U, L) . This reasoning works all the way back to the beginning of the game, so that play must start at (U, L) .

Intuitively, in both examples, 2 can force 1 to move to D by credibly threatening to play R until he does so. In the second example, though, this threat can only be credible at a particular point in the game, making it optimal for 1 to play U in the meantime. Hence (U, L) is played up until this point. In the first example, the threat is credible for the entire game, forcing 1 to begin with D and making (D, L) the outcome every period.

6 Conclusion

In summary, two seemingly minor variations on finitely repeated games have some surprising effects. While it seems clear that switching costs might enforce a commitment to not changing actions, it is more surprising that very small switching costs can make such a commitment credible even at the very beginning of the game. As a result, repeated games which have multiple equilibria without switching costs can have a unique

equilibrium with small switching costs. Switching costs can have other effects, though, since we may need the threat to switch actions to be credible for certain equilibria. As a result, we can but do not necessarily get cooperation in the finitely repeated Prisoners' Dilemma with small switching costs. Finally, switching costs can make credible some surprisingly complex threats, leading to unexpected outcomes.

A natural question to ask is whether it is important that we have fixed the total length of time the game is played. It is not hard to see that this assumption is simply not relevant to the analysis. More concretely, recall that the total length of time the game is played is M and the number of periods is M/Δ . M is not relevant to any of the proofs of our results as these are all based on induction from the end of the game. One surprising implication of this fact is that the switching cost can be arbitrarily small relative to total game payoffs and still have the strong effects shown above. To be more concrete, consider, for example, the two-by-two coordination game results in Theorem 2. We fix ε and characterize an interval $(0, \bar{\Delta}_\varepsilon)$ such that for any Δ in this interval, the unique subgame perfect equilibrium outcome is the risk dominant and Pareto dominant outcome repeated every period. It is not hard to see that the interval is independent of M , so we can fix any Δ in the interval and let $M \rightarrow \infty$. All along this sequence of games, we obtain the same equilibrium result even though the switching cost relative to total payoffs is going to zero.

A natural related question is whether the results carry over in some fashion to infinitely repeated games. We discuss this question in more detail in Lipman and Wang [1999a], but give a brief answer here. The key to whether small switching costs affect play in infinitely repeated games is whether we think of each period as being “very small” in the infinitely repeated game. To understand this, suppose we consider a simple variation on the usual discounting formulation, evaluating paths of play by the discounted sum over periods of the payoff in a period minus a switching cost if incurred in that period. This formulation makes it impossible to shrink the payoff in a period relative to the switching cost without making the switching cost large. Since this effect is critical to our analysis, this means that our results will not generalize to this form of the infinitely repeated game.⁶

By contrast, consider the opposite extreme where we use the limit of means to evaluate the payoffs from the game and subtract from this limit the total switching costs incurred. In this case, the payoff in a single period is certainly small relative to the switching cost since a gain in only a single period has no effect! In other words, in this setting, even with a “small” ε , the switching cost exceeds the gain that can be earned in any finite number of periods. In this case, the switching cost has a very large effect, dramatically changing the equilibrium payoff set.

⁶A result like this is contained in Chakrabarti [1990].

Appendix

A A Useful Lemma

Lemma 1 *Let g^* denote the largest gain in payoff possible from a unilateral change of action. That is,*

$$g^* = \max \left\{ \max_{a_1, a'_1, a_2} u_1(a'_1, a_2) - u_1(a_1, a_2), \max_{a_1, a_2, a'_2} u_2(a_1, a'_2) - u_2(a_1, a_2) \right\}.$$

Then for any t such that $t\Delta g^ < \varepsilon$, both players must use the action played at the previous date in any subgame perfect equilibrium.*

The proof is by induction. So consider $t = 1$ and assume that $t\Delta g^* < \varepsilon$ for this t . Then in this subgame, it is a strictly dominant strategy for each player to not change actions since for any action by the opponent, the gain in payoff must be strictly less than the switching cost. Hence the claim is certainly true at $t = 1$. So consider any t such that $t\Delta g^* < \varepsilon$ and suppose we have proved the claim for all smaller t . Then we know that whatever actions are played at t will be played at all subsequent dates. But then just as above, whatever action the opponent plays at t , it must be strictly optimal to not change actions. ■

B Proof of Theorem 1

Suppose (a_1^*, a_2^*) is weakly stable. Without loss of generality, assume that this profile uniquely maximizes 1's payoff and that

$$u_2(a_1^*, a_2^*) - u_2(a_1^*, a_2) \geq u_1(a_1^*, a_2^*) - u_1(a_1^*, a_2), \quad \forall a_2 \in A_2.$$

Let \bar{u}_1 denote the maximum of $u_1(a_1, a_2)$ over $(a_1, a_2) \in A_1 \times A_2$ subject to $a_1 \neq a_1^*$.

The key to the theorem is to prove that that if 1 plays a_1^* at some period, he will never change again, so that 2 will either choose his best reply or not switch from his current action, depending on whether there is enough time remaining to make the gain from switching exceed the cost. That is, for all t , if (a_1^*, a_2) were played at $t + 1$, then 1 plays a_1^* at t and 2 plays a_2^* if

$$t\Delta[u_2(a_1^*, a_2^*) - u_2(a_1^*, a_2)] > \varepsilon$$

and 2 plays a_2 otherwise. We assume that ε is such that we never have equality in this equation.

To show this, let g^* be the largest unilateral gain in payoffs as defined in Lemma 1. Define t^* to be the smallest t such that

$$t\Delta g^* \geq \varepsilon.$$

We assume that ε is such that $t^*\Delta g^* > \varepsilon$. Note that as $\Delta \downarrow 0$, $t^* \rightarrow \infty$. By Lemma 1, we know that in any subgame perfect equilibrium, neither player changes actions at any date $t < t^*$.

Hence the claim we wish to show is trivially true for $t < t^*$. So consider any $t \geq t^*$ and suppose we've demonstrated the result for all smaller t . Suppose (a_1^*, a_2) were played at $t + 1$. We now show that no matter what player 2's strategy at t is, player 1's optimal action at t must be a_1^* . So let 2's strategy at t be a_2' .

First, suppose $a_2' = a_2^*$. In this case, the claim is obvious. If player 1 plays a_1^* , then the induction hypothesis tells us that his payoff will be $t\Delta u_1(a_1^*, a_2^*)$, the highest payoff he can possibly get. Changing actions must lower his payoff.

Next, suppose that $a_2' \neq a_2^*$ but

$$(t - 1)\Delta[u_2(a_1^*, a_2^*) - u_2(a_1^*, a_2')] > \varepsilon.$$

Again, the claim is obvious. The induction hypothesis tells us that if 1 doesn't change actions, his payoff is

$$\Delta u_1(a_1^*, a_2') + (t - 1)\Delta u_1(a_1^*, a_2^*)$$

while if he does change, his payoff certainly cannot be more than

$$t\Delta u_1(a_1^*, a_2^*) - \varepsilon,$$

which is obviously smaller if Δ is sufficiently small.

Finally, then, suppose that $a_2' \neq a_2^*$ and

$$(t - 1)\Delta[u_2(a_1^*, a_2^*) - u_2(a_1^*, a_2')] < \varepsilon. \tag{4}$$

In this case, 1's payoff to a_1^* is $t\Delta u_1(a_1^*, a_2')$. His payoff to changing actions cannot exceed

$$\max\{t\Delta \bar{u}_1 - \varepsilon, t\Delta u_1(a_1^*, a_2^*) - 2\varepsilon\}.$$

Hence he certainly does not switch if

$$t\Delta \max\left\{\bar{u}_1 - u_1(a_1^*, a_2'), \frac{u_1(a_1^*, a_2^*) - u_2(a_1^*, a_2')}{2}\right\} < \varepsilon.$$

Let $R(a'_2)$ be the term in brackets on the left-hand side. Note that

$$R(a'_2) < u_1(a_1^*, a_2^*) - u_1(a'_1, a'_2) \leq u_2(a_1^*, a_2^*) - u_2(a_1^*, a'_2).$$

Let $D(a'_2) > 0$ denote $u_2(a_1^*, a_2^*) - u_2(a_1^*, a'_2) - R(a'_2)$. Then we see that 1 will not change actions if

$$\Delta R(a'_2) - (t-1)\Delta D(a'_2) + (t-1)\Delta[u_2(a_1^*, a_2^*) - u_2(a_1^*, a'_2)] < \varepsilon.$$

By (4), a sufficient condition for this is $R(a'_2) < (t-1)D(a'_2)$. Because $t \geq t^*$, a sufficient condition is $R(a'_2) < (t^* - 1)D(a'_2)$. The only place Δ appears in this equation is that t^* depends on Δ . As $\Delta \downarrow 0$, $t^* \rightarrow \infty$. Hence this must hold for Δ sufficiently small.

Summarizing, we have shown that if the induction hypothesis holds for all later dates, then the description of 1's strategy holds at t . Given this, it is obvious that the description of 2's strategy also holds at t .

In light of this, consider 1's choice of action in the first period. If 1 plays a_1^* in the first period, his payoff must be at least $(M-\Delta)u_1(a_1^*, a_2^*) + \Delta \min_{a_2 \in A_2} u_1(a_1^*, a_2)$, while starting with any other action cannot give him a payoff larger than $\max\{M\bar{u}_1, Mu_1(a_1^*, a_2^*) - \varepsilon\}$ which must be smaller if Δ is sufficiently small. Hence 1 must begin with action a_1^* and so never changes actions. Clearly, 2's best reply is a_2^* in every period. ■

C Proof of Theorem 2

First, assume (L, L) is risk dominant, implying $a - c > b - d$, so $a - c$ is the largest unilateral payoff gain. So by Lemma 1, neither player changes actions at any t such that $(a - c)t\Delta < \varepsilon$. Fix any $\varepsilon \in (0, M(a - c))$.

Assume for the remainder of the proof that ε is such that there is no integer t satisfying $(a - c)t\Delta = \varepsilon$. Also, assume that Δ is sufficiently small that there are values of t satisfying

$$(b - d)t\Delta < \varepsilon < (a - c)t\Delta.$$

We now show that for any t such that $(a - c)t\Delta > \varepsilon$, the following two facts are true. First, if either player used L in period $t + 1$, he never changes actions again. Second, if (R, L) or (L, R) were played at $t + 1$, then the player who used R previously switches to L at t and the two play (L, L) from then on. The proof of this claim is by induction.

To begin, let t^* be the smallest t such that $(a - c)t\Delta > \varepsilon$. By our choice of ε , then, $(a - c)(t^* - 1)\Delta < \varepsilon$, so we know that whatever actions are used at t are used in every

later period. Suppose player i used L at period $t + 1$. Then i 's continuation payoff as a function of his period t action and the period t action of his opponent is

$$\begin{array}{cc} & \begin{array}{c} L \\ R \end{array} \\ \begin{array}{c} L \\ R \end{array} & \begin{array}{cc} at^*\Delta & dt^*\Delta \\ ct^*\Delta - \varepsilon & bt^*\Delta - \varepsilon \end{array} \end{array}$$

Because $a > c$, it is obvious that $at^*\Delta > ct^*\Delta - \varepsilon$. Also, by assumption, $(b - d)t^*\Delta < \varepsilon$, so $dt^*\Delta > bt^*\Delta - \varepsilon$. Hence i must play L at period t^* , establishing the first claim. To show the second, suppose that player i used L and player j used R at period $t + 1$. From the previous argument, i will not change actions ever again. Hence j will switch to L if $at^*\Delta - \varepsilon > ct^*\Delta$ or $(a - c)t^*\Delta > \varepsilon$ which is true by the definition of t^* .

To complete the induction, fix any $t > t^*$ and suppose we have shown our claim for all later periods. Suppose i played L at $t + 1$. Then the worst payoff i could get from playing L at t is $d\Delta + a(t - 1)\Delta$ because j will necessarily play L (switching to it if need be) in the next period and from then onward by the induction hypothesis. Suppose instead that i uses R at t . His payoff certainly cannot exceed $\max\{c, b\}\Delta + a(t - 1)\Delta - \varepsilon$ since this calculation gives him the highest possible payoff for this period, the highest payoff in the matrix thereafter, and only charges him the switching cost once (even though he'd have to switch actions twice to earn a). For Δ small enough, this latter payoff is strictly smaller. Hence i will play L at t . To show the second fact, then, is simple. If i played L and j played R at $t + 1$, then we know from the above i plays L from this point on. Because $(a - c)t\Delta > \varepsilon$, j will switch to L and (L, L) is played from then on. By induction, then, we see that the two facts above are true for the entire game.

Finally, consider the first period of the game. If i plays L , his payoff must be at least $d\Delta + a(M - \Delta)$. For Δ close to zero, this is close to Ma . If he plays R instead, his payoff cannot be larger than

$$\max\{bM, b\Delta + a(M - \Delta) - \varepsilon, c\Delta + a(M - \Delta) - \varepsilon\},$$

which converges to $\max\{Mb, Ma - \varepsilon\}$ as $\Delta \downarrow 0$. Clearly, Ma is strictly larger than this, so for Δ sufficiently small, i must play L in the first period. Hence both must play L in the first period and every period thereafter.

This completes the proof of the implications of (L, L) being risk dominant. So suppose (R, R) is risk dominant. We construct subgame perfect equilibria as follows. For any history such that (L, L) or (R, R) was played in the previous period, no player changes actions. If (L, R) or (R, L) were played, the player who used L changes to R as long as $(b - d)t\Delta \geq \varepsilon$. If (L, R) or (R, L) were played and this inequality is not satisfied, then neither player changes actions. This specifies actions in all periods except the first. To construct a subgame perfect equilibrium with outcome (L, L) in every period, complete the strategies by specifying that both begin with L . To construct an equilibrium

with outcome (R, R) every period, have the strategies begin with R . To see that either specification gives us a subgame perfect equilibrium, first consider any period other than the first. If (L, L) were played in the previous period, clearly neither player has an incentive to change since this gives the highest possible payoff. If (R, R) were played in the previous period, a deviation away from this outcome leads to (L, R) or (R, L) . According to the strategies above, this leads either to no further changes (in which case the deviator is worse off) or to the deviator moving back to (R, R) , in which case the two switching costs again imply that the deviator is strictly worse off. Finally, suppose either (L, R) or (R, L) were played in the previous period. If it is early enough in the game, the player who played L is supposed to change actions. Given this, his opponent certainly does not change since he will only have to change back. Given that the opponent will not change, the player who played L previously finds it optimal to change actions. If it is too late in the game for the player who played L to change actions, then $b - d > a - c$ implies that it is too late for the other player to change actions also. Hence in this situation, neither player will change, just as specified by the equilibrium strategies.

Finally, consider the first period. If the opponent is expected to play L , clearly L is the optimal strategy for i since this leads to a payoff of Ma while playing R leads the opponent to change the next period giving i a payoff of approximately Mb . Hence given the play in the subgames, it is an equilibrium for both to start with L . Similarly, though, it is an equilibrium for both to start with R . If i expects his opponent to start with R , then starting with R gives him a payoff of Mb , while starting with L gives a payoff approximately $Mb - \varepsilon$ which is clearly worse. ■

D Proof of Theorem 3

For concreteness, assume $a_2 - c_2 > b_1 - d_1$. The case where the reverse strict inequality holds is entirely symmetric.

Fix any $\varepsilon \in (0, M \min\{a_1 - c_1, b_1 - d_1, b_2 - d_2\})$. We show by induction that if Δ is sufficiently small, then for all t such that $(a_2 - c_2)t\Delta < \varepsilon$, no one changes actions at t if the action profile at $t + 1$ was (L, L) , (R, R) , or (L, R) . Obviously, if Δ is sufficiently small, this is true at $t = 1$. So consider any period t satisfying this inequality and suppose the result has been shown for all smaller t .

First, suppose player 1 used L at $t + 1$. Given the induction hypothesis, we see that his continuation payoff as a function of the period t profile of actions is

$$\begin{array}{cc} & \begin{array}{cc} L & R \end{array} \\ \begin{array}{c} L \\ R \end{array} & \begin{array}{cc} a_1 t \Delta & d_1 t \Delta \\ y_{1t} - \varepsilon & b_1 t \Delta - \varepsilon \end{array} \end{array}$$

where we do not know what y_{1t} is. However, note that a_1 is 1's highest payoff in the matrix, so $a_1 t \Delta \geq y_{1t}$, so $a_1 t \Delta > y_{1t} - \varepsilon$. Also, $d_1 t \Delta > b_1 t \Delta - \varepsilon$ iff $(b_1 - d_1) t \Delta < \varepsilon$. By assumption $\varepsilon > (a_2 - c_2) t \Delta > (b_1 - d_1) t \Delta$, so this holds. Hence 1 must use L at t and, by the induction hypothesis, in every later period. Given this, we see that if (L, L) were played at $t + 1$, 2 will never change actions again either.

So suppose 2 played R at $t + 1$. Then 2's continuation payoff as a function of the period t actions is

$$\begin{array}{cc} & L & R \\ L & a_2 t \Delta - \varepsilon & c_2 t \Delta \\ R & y_{2t} - \varepsilon & b_2 t \Delta \end{array}$$

Again, the induction hypothesis does not tell us what y_{2t} is, but the fact that b_2 is 2's highest payoff in the matrix tells us that $b_2 t \Delta > y_{2t} - \varepsilon$. Also, $c_2 t \Delta > a_2 t \Delta - \varepsilon$ because $(a_2 - c_2) t \Delta < \varepsilon$ by assumption. Hence 2 must play R at t and, by the induction hypothesis, in every subsequent period. Given this, we see that if (R, R) were played at $t + 1$, 1 will never change actions again either.

Assume ε is such that there is no integer t such that $(a_2 - c_2) t \Delta = \varepsilon$. We claim that if Δ is sufficiently small, then for all t such that $(a_2 - c_2) t \Delta > \varepsilon$, if player 1 used L at $t + 1$, then the outcome is (L, L) at t . The proof of this is again by induction. So first consider the smallest t in this range. Because $(a_2 - c_2)(t - 1) \Delta < \varepsilon$, no one will ever change actions from $t - 1$ onward if the profile at t is (L, L) , (L, R) , or (R, R) . Suppose 1 played L at $t + 1$. The same calculations as above show that his best reply to either action by 2 at t is to play L . Hence 1 plays L at t . Clearly, if both played L at $t + 1$, 2's best reply is to play L as well. So suppose (L, R) was played at $t + 1$. Then 2's best reply is L iff $a_2 t \Delta - \varepsilon > c_2 t \Delta$, which is true by assumption.

To complete the induction, then, consider any t such that $(a_2 - c_2) t \Delta > \varepsilon$ and suppose we have demonstrated the result for all smaller t . Suppose 1 played L at $t + 1$. If he plays L at t , his payoff is, at worst, $d_1 \Delta + a_1(t - 1) \Delta$. If he plays R instead, his payoff certainly cannot be larger than $\max\{b_1, c_1\} \Delta + a_1(t - 1) \Delta - \varepsilon$. Hence L is certainly optimal if Δ is sufficiently small. If 2 played R at $t + 1$, he switches to L at t . If he did not, he would switch at $t - 1$ anyway, so the only effect not switching would have on his payoff is to give him a lower payoff in period t . Hence the outcome at t must be (L, L) , completing the induction argument.

To complete the proof then, we see that if 1 plays L at the first period, his payoff must be approximately Ma_1 , while the approximate payoff to playing R at the first period cannot be larger than $\max\{Mb_1, Mc_1, Ma_1 - \varepsilon\}$. Clearly, $a_1 > b_1, c_1$ implies that 1's equilibrium strategy must be to play L in the first period. Given this, 2 must play L in the first period as well and the outcome is (L, L) in every period. ■

E Proof of Theorem 5

Clearly, $c - a \geq b - d$ implies that the largest unilateral gain in payoffs is from switching to D against C . Hence by Lemma 1, neither player changes actions for any t such that $t\Delta(c - a) < \varepsilon$. Assume that ε is such that there is no integer t satisfying $t\Delta(b - d) = \varepsilon$.

Let t^* denote the smallest integer t such that $t\Delta(b - d) > \varepsilon$. We first show that unless $t^* = 1$, both players defect at t^* and thereafter. There are two cases to consider. First, suppose Δ is such that there are no values of t satisfying $(b - d)t\Delta < \varepsilon \leq (c - a)t\Delta$. In this case, we must have $(c - a)(t^* - 1)\Delta < \varepsilon$. Hence whatever actions are played at t are played in all subsequent periods. So suppose i defected at $t^* + 1$. Clearly, the dominant strategy property implies that it cannot be optimal for him to pay to switch to cooperating, so i will defect at t^* . Suppose then that i cooperated at $t^* + 1$. If his opponent cooperates at t^* , i should defect at t^* as long as

$$ct^*\Delta - \varepsilon > at^*\Delta,$$

or $(c - a)t^*\Delta > \varepsilon$ which holds as $(c - a)t^*\Delta \geq (b - d)t^*\Delta > \varepsilon$. Similarly, if his opponent defects at t^* , i should defect at t^* as long as $(b - d)t^*\Delta > \varepsilon$ which also holds. Hence i should defect at t^* . Hence both players defect at t^* and every subsequent period.

The case where Δ is small enough that there are values of t such that $(b - d)t\Delta < \varepsilon \leq (c - a)t\Delta$ is slightly more complex. To show the statement claimed, we must first characterize behavior in this interval. We claim that for any period t such that $(b - d)t\Delta < \varepsilon \leq (c - a)t\Delta$, any player who defected at period $t + 1$ must defect at t and from then on. Furthermore, if i cooperated at $t + 1$ and his opponent defected, then i continues to cooperate. We show this by induction, so first consider the last t in this interval. By definition, $(c - a)(t - 1)\Delta < \varepsilon$, so whatever actions are played at t will be played in all succeeding periods. So suppose i defected at $t + 1$. Clearly, the dominant strategy property implies that it cannot be optimal for him to pay to switch to cooperating, so i will defect at t . Suppose instead that i cooperated at $t + 1$ and his opponent defected. We have just shown that his opponent will continue to defect. So it is optimal for i to cooperate if $dt\Delta > bt\Delta - \varepsilon$ which holds by assumption.

To complete the induction, suppose we are at a period t such that $(b - d)t\Delta < \varepsilon < (c - a)t\Delta$ and that we know that at all future periods, any player who defected in the past will continue to do so and that any player who cooperates against a defector will continue to do so (so the next period need not be inside this interval). Suppose i defected at $t + 1$. If the opponent defects at t , i is clearly best off defecting as well because whatever actions are played at t will be repeated thereafter. If the opponent cooperates at t and i defects, this is repeated thereafter, yielding a payoff of $ct\Delta$ for i . Clearly, this is the highest possible payoff i could get, so it must be optimal for i to defect in this case as well. Hence i defects at t . Given this, suppose i cooperated and his opponent defected

at the previous period $t + 1$. Then the opponent will defect from t onward. Because $(c - a)t\Delta < \varepsilon$, i will continue cooperating.

We now use this to show our claim that if t^* is the last period t satisfying $(b - d)t\Delta > \varepsilon$ and is not the first period, then both players defect at t^* and thereafter in any subgame perfect equilibrium. The restriction to values of ε such that there is no t with $(b - d)t\Delta = \varepsilon$ implies that $(b - d)(t^* - 1)\Delta < \varepsilon \leq (c - a)(t^* - 1)\Delta$. Hence unless both players cooperate, whatever actions are played at t^* will be played in every subsequent period. First, suppose i expects his opponent to defect at t^* . Then he expects the actions at t^* to be repeated in all periods. Clearly, then, if i defected at $t^* + 1$, he should not pay to switch to cooperating. Similarly, it is easy to use $(b - d)t^*\Delta > \varepsilon$ to show that if he cooperated in the previous period, he should defect at t^* . Suppose then that i expects his opponent to cooperate at t^* . To consider the worst case for proving i should defect, suppose he cooperated at $t^* + 1$. If he defects at t^* , these actions are repeated thereafter so his payoff is $ct^*\Delta - \varepsilon$. If he cooperates, his payoff is certainly no larger than

$$\max\{at^*\Delta, a\Delta + c(t^* - 1)\Delta - \varepsilon\}.$$

Hence defection is optimal if $ct^*\Delta - \varepsilon > at^*\Delta$ (which must hold) and $ct^*\Delta - \varepsilon > a\Delta + c(t^* - 1)\Delta - \varepsilon$ (which is implied by $c > a$). Hence i is better off defecting at t^* regardless of what the opponent is expected to do. Hence both players defect at t^* and in every subsequent period.

The induction from here is straightforward. First, if there is at least one period other than the first such that $(b - d)t\Delta > \varepsilon$, then at the period before t^* , we can use backward induction to see that both players start by defecting and defect in every subsequent period. To see how the induction goes, note that by the induction hypothesis, both players are expected to begin defecting from the next period onward. Hence the dominant strategy property implies that if i defected at the previous period, he should certainly defect at the current period. If i cooperated at the previous period, he may as well pay the switching cost now to switch to defection and get the higher payoff today as he will certainly do so tomorrow otherwise. If this is the first period so there was no previous period, switching costs are irrelevant and the domination implies that i should defect. Second, suppose the second period t has $(c - a)t\Delta < \varepsilon$. In this case, whatever actions are played in the first period are played the rest of the game, so it is like there is only one period. Obviously, in this case, both players defect always. ■

F Proof of Theorems 6 and 7

F.1 Preliminaries

We first establish some results which hold for both examples. Throughout this subsection, we do *not* assume either $a_1 + b_1 > 2c_1$ or the reverse. Assume ε is sufficiently small that each of the t^* 's defined below is less than M/Δ .

By Lemma 1, assumption (1) implies that no player changes actions for any t such that $(d_2 - b_2)t\Delta < \varepsilon$. Assume that there is no integer t such that $(d_2 - b_2)t\Delta = \varepsilon$. Define t_1^* to be the smallest integer t such that $(d_2 - b_2)t\Delta > \varepsilon$. Note that $t_1^*\Delta > \varepsilon/(d_2 - b_2)$ for all Δ . Hence it is bounded away from zero as $\Delta \downarrow 0$. It is easy to see that, at period t_1^* , the only change of action any player would make is that if (D, R) were played at the previous period, 2 will change actions to L . It is not hard to extend this to show by induction that the same is true for all $t \geq t_1^* + 1$ such that

$$(c_1 - d_1)t\Delta + (b_1 - c_1)\Delta < \varepsilon, \quad (5)$$

if Δ is sufficiently small.

Hence the only change of action which occurs at any $t \geq t_1^* + 1$ satisfying (5) is that 2 changes from (D, R) to (D, L) . Assume that there is no integer t such that

$$(c_1 - d_1)t\Delta + (b_1 - c_1)\Delta = \varepsilon.$$

Let t_2^* be the smallest integer t such that

$$(c_1 - d_1)t\Delta + (b_1 - c_1)\Delta > \varepsilon.$$

(Note that t_2^* is the $t^*(\varepsilon, \Delta)$ defined in Theorem 7.) It is easy to see that $t_2^* > t_1^*$ if Δ is sufficiently small.

The only change at $t = t_2^*$ is that if 1 played U at the previous period, his best reply if 2 plays R at t is to play D . To see this, note that the payoffs to the two players from $t = t_2^*$ onward as a function of the period t_2^* actions, not counting switching costs incurred at t_2^* are

	L	R
U	$a_1t\Delta, a_2t\Delta$	$d_1t\Delta, c_2t\Delta$
D	$c_1t\Delta, d_2t\Delta$	$b_1\Delta + c_1(t-1)\Delta, b_2\Delta + d_2(t-1)\Delta - \varepsilon$

By the definition of t_2^* , 1's best reply to R if he played U (or by implication D) in the previous period is D . It is still true at $t = t_2^*$ that if 1 played D at the previous period, we must have (D, L) played at t_2^* , while if (U, L) were played at the previous period, it is played again at t_2^* .

Putting this together, suppose (U, R) were played at the preceding period. We know from the above that U is 1's best reply to L and that D is his best reply to R . Also, it is easy to see that 2's best reply to U is R , while his best reply to D is L . Hence we must have mixing at period t_2^* if (U, R) were played in the previous period. Let $\alpha_{t_2^*}$ denote 1's probability of playing U at this subgame and $\beta_{t_2^*}$ denote 2's probability of playing L . Let $V_i^{t_2^*}$ denote i 's expected payoff in this subgame. So we must have

$$\begin{aligned} V_1^{t_2^*} &= c_1 t_2^* \Delta + (1 - \beta_{t_2^*})(b_1 - c_1)\Delta - \varepsilon \\ &= a_1 t_2^* \Delta \beta_{t_2^*} + d_1 t_2^* \Delta (1 - \beta_{t_2^*}) \end{aligned}$$

and

$$\begin{aligned} V_2^{t_2^*} &= \alpha_{t_2^*} a_2 t_2^* \Delta + (1 - \alpha_{t_2^*}) d_2 t_2^* \Delta - \varepsilon \\ &= \alpha_{t_2^*} c_2 t_2^* \Delta + (1 - \alpha_{t_2^*}) [b_2 \Delta + d_2 (t_2^* - 1)\Delta - \varepsilon]. \end{aligned}$$

It is straightforward to solve for $\beta_{t_2^*}$ and $\alpha_{t_2^*}$ and to show that both converge to zero as $\Delta \downarrow 0$. Hence for Δ small, $V_1^{t_2^*}$ is approximately $b_1 \Delta + c_1 (t_2^* - 1)\Delta - \varepsilon$, while $V_2^{t_2^*}$ is approximately $d_2 t_2^* \Delta - \varepsilon$.

So consider period $t = t_2^* + 1$. Now the payoffs from t onward as a function of the period t actions (not counting switching costs incurred at t) are

	L	R
U	$a_1 t \Delta, a_2 t \Delta$	$d_1 \Delta + V_1^{t_2^*}, c_2 \Delta + V_2^{t_2^*}$
D	$c_1 t \Delta, d_2 t \Delta$	$b_1 \Delta + c_1 (t - 1)\Delta, b_2 \Delta + d_2 (t - 1)\Delta - \varepsilon$

Again, if Δ is sufficiently small, then if 1 played D at the previous period, the outcome at $t_2^* + 1$ is (D, L) . Also, if (U, R) were played, we must again have mixing. The change from the previous case is that we must also have mixing if (U, L) were played. The reason for this is that now 2's best reply to U is R , not L , if Δ is sufficiently small. To see this, note that from the above, if Δ is sufficiently small, 2's payoff to switching to (U, R) is approximately

$$c_2 \Delta + (b_2 - d_2)\Delta + d_2 t_2^* \Delta - 2\varepsilon,$$

while his payoff to continuing with L is $a_2 (t_2^* + 1)\Delta$. The former is larger if

$$(d_2 - a_2) t_2^* \Delta + (b_2 - d_2)\Delta + (c_2 - a_2)\Delta > 2\varepsilon.$$

For Δ small, this holds if

$$\frac{(d_2 - a_2)}{2} t_2^* \Delta > \varepsilon.$$

Recall, though, that for Δ small, $t_2^* \Delta$ is approximately $\varepsilon / (c_1 - d_1)$, so this holds for Δ small if

$$\frac{d_2 - a_2}{2} > c_1 - d_1,$$

which holds by (3).

It is not hard to see that 2's randomization must be the same whether (U, L) or (U, R) was played at the previous period. Since 1's payoffs are unaffected by the difference between these histories and since 2's randomization is chosen to make 1 indifferent, the fact that the randomization that makes 1 indifferent is unique implies that the same randomization must be used on either history. Let $\beta_{t_2^*+1}$ denote 2's probability of L on either history. Let $\alpha_{t_2^*+1}(a)$ be 1's probability of U if (U, a) were played at the preceding period, $a = L, R$. Also, let $V_1^{t_2^*+1}$ denote 1's expected payoff in either of these subgames and let $V_2^{t_2^*+1}(a)$ denote 2's expected payoff in the subgame following (U, a) , $a = L, R$. Then we must have

$$\begin{aligned} V_1^{t_2^*+1} &= c_1(t_2^* + 1)\Delta + (1 - \beta_{t_2^*+1})(b_1 - c_1)\Delta - \varepsilon \\ &= a_1(t_2^* + 1)\Delta\beta_{t_2^*+1} + (1 - \beta_{t_2^*+1})[d_1\Delta + V_1^{t_2^*}], \end{aligned}$$

$$\begin{aligned} V_2^{t_2^*+1}(L) &= a_2(t_2^* + 1)\Delta\alpha_{t_2^*+1}(L) + (1 - \alpha_{t_2^*+1}(L))d_2(t_2^* + 1)\Delta \\ &= \alpha_{t_2^*+1}(L)[c_2\Delta + V_2^{t_2^*}] + (1 - \alpha_{t_2^*+1}(L))[b_2\Delta + d_2t_2^*\Delta - \varepsilon] - \varepsilon, \end{aligned}$$

and

$$\begin{aligned} V_2^{t_2^*+1}(R) &= a_2(t_2^* + 1)\Delta\alpha_{t_2^*+1}(R) + (1 - \alpha_{t_2^*+1}(R))d_2(t_2^* + 1)\Delta - \varepsilon \\ &= \alpha_{t_2^*+1}(R)[c_2\Delta + V_2^{t_2^*}] + (1 - \alpha_{t_2^*+1}(R))[b_2\Delta + d_2t_2^*\Delta - \varepsilon]. \end{aligned}$$

It is again not hard to show that $\beta_{t_2^*+1}$ converges to zero as $\Delta \downarrow 0$.

F.2 Proof of Theorem 7

Because the proof of Theorem 7 relies on less information about these value functions, we complete it first. So we now assume that $a_1 + b_1 < 2c_1$. Given the results above, 1's payoffs from $t_2^* + 2$ onward as a function of the actions at this period are

	L	R
U	$a_1\Delta + V_1^{t_2^*+1}$	$d_1\Delta + V_1^{t_2^*+1}$
D	$c_1(t_2^* + 2)\Delta$	$b_1\Delta + c_1(t_2^* + 1)\Delta$

From above, we know that for Δ small, $V_1^{t_2^*+1}$ is approximately

$$b_1\Delta + c_1t_2^*\Delta - \varepsilon.$$

So suppose 1 played U in the previous period. Then D is the best reply to L for Δ sufficiently small if

$$a_1\Delta + b_1\Delta + c_1t_2^*\Delta - \varepsilon > c_1t_2^*\Delta + 2c_1\Delta - \varepsilon$$

which is implied by $a_1 + b_1 < 2c_1$. Also, D is the best reply to R for Δ sufficiently small if $d_1\Delta < c_1\Delta$ which also holds. Hence if 1 played U at the previous period, he must play

D at $t_2^* + 2$ if Δ is sufficiently small. Clearly, if it is optimal to pay a switching cost to switch to D , it must be optimal to play D if this requires no switching cost. That is, if 1 played D at the previous period, he continues to play it at $t_2^* + 2$. In short, 1 *must* play D at period $t_2^* + 2$ if Δ is sufficiently small, regardless of the history.

We know that L is 2's best reply to D at this point, so this implies that if Δ is sufficiently small, then regardless of the history, (D, L) will be played at $t_2^* + 2$ and every period thereafter in any subgame perfect equilibrium.

Given this, it is easy to show by induction that for all $t \geq t_2^* + 3$ such that

$$a_1(t - t_2^* - 2)\Delta + c_1(t_2^* + 2)\Delta - 2\varepsilon < c_1t\Delta, \quad (6)$$

player 2 plays L at t and player 1 plays whatever action he used the previous period.

Let t_3^* be the smallest integer t violating (6). For simplicity, assume there is no integer t satisfying this equation with equality. It is easy to show that the induction above implies that 2 will play L at t_3^* and that 1 will play U at t_3^* if he played it in the previous period. However, if 1 played D at the previous period, he will change to U at t_3^* since it is better to earn a_1 until $t_2^* + 2$ even at the cost of having to change actions twice. Hence no matter what happened in the previous period, (U, L) is played at t_3^* . It is not hard to show by induction that this implies that (U, L) is played in every previous period, including the first period of the game. This completes the proof of Theorem 7.

F.3 Proof of Theorem 6

We now assume $2c_1 < a_1 + b_1$. One implication of this which we use is that it implies $b_1 > d_1$. To see this, note that (2) implies $2c_1 > a_1 + d_1$, so we must now have $a_1 + b_1 > a_1 + d_1$ or $b_1 > d_1$.

The rest of the proof is by induction. We show that for all $t \geq t_2^* + 2$, if 1 played D at $t + 1$, then the outcome at t is (D, L) , while if 1 played U at $t + 1$, then both players must randomize at t where this randomization is uniquely defined. More specifically, we show by induction that if Δ is sufficiently small, then for all $t \geq t_2^* + 2$,

$$c_1t\Delta - \varepsilon < a_1\Delta + V_1^{t-1} < c_1t\Delta + \varepsilon, \quad (7)$$

$$d_1\Delta + V_1^{t-1} < b_1\Delta + c_1(t-1)\Delta - \varepsilon, \quad (8)$$

and

$$a_2\Delta + V_2^{t-1}(L) < c_2\Delta + V_2^{t-1}(R) - \varepsilon. \quad (9)$$

To understand these claims, note that given our statement of what happens in subsequent periods, payoffs from period t onward as a function of the period t actions, not including costs of switching incurred at t , are

$$\begin{array}{cc} & \begin{array}{c} L \\ R \end{array} \\ \begin{array}{c} U \\ D \end{array} & \begin{array}{cc} a_1\Delta + V_1^{t-1}, a_2\Delta + V_2^{t-1}(L) & d_1\Delta + V_1^{t-1}, c_2\Delta + V_2^{t-1}(R) \\ c_1t\Delta, d_2t\Delta & b_1\Delta + c_1(t-1)\Delta, b_2\Delta + d_2(t-1)\Delta - \varepsilon \end{array} \end{array}$$

where V_1^{t-1} is 1's expected payoff from period $t-1$ onward and $V_2^{t-1}(a)$ is 2's expected payoff from $t-1$ onward given that (U, a) is played at t . As explained above, 1's expected payoff from $t-1$ onward must be the same if (U, R) or (U, L) is played at t . Given this, (7) implies that 1's best reply to L is whatever action he played in the previous period. That is, the difference in payoffs is below ε . Equation (8) says that 1's best reply to R is D even if he has to pay a switching cost to play D . Hence this is his best reply regardless of the action he played at $t+1$. Similarly, (9) states that 2's best reply to U is R regardless of what action 2 used at $t+1$. It is immediate from the payoffs above that 2's best reply to D is L regardless of the action he used at $t+1$.

Hence these inequalities imply that if 1 played D in the previous period, he must play D at t and hence 2 must play L at t . If 1 played U in the preceding period, however, both players must randomize.

Defining $\alpha_t(a)$ analogously to $\alpha_{t^*+2}(a)$ and β_t analogously to β_{t^*+2} , we see that for $t \geq t_2^* + 2$,

$$\begin{aligned} V_1^t &= c_1t\Delta + (1 - \beta_t)(b_1 - c_1)\Delta - \varepsilon \\ &= \beta_t[a_1\Delta + V_1^{t-1}] + (1 - \beta_t)[d_1\Delta + V_1^{t-1}], \\ V_2^t(L) &= \alpha_t(L)[a_2\Delta + V_2^{t-1}(L)] + (1 - \alpha_t(L))d_2t\Delta \\ &= \alpha_t(L)[c_2\Delta + V_2^{t-1}(R)] + (1 - \alpha_t(L))[b_2\Delta + d_2(t-1)\Delta - \varepsilon] - \varepsilon, \end{aligned}$$

and

$$\begin{aligned} V_2^t(R) &= \alpha_t(R)[a_2\Delta + V_2^{t-1}(L)] + (1 - \alpha_t(R))d_2t\Delta - \varepsilon \\ &= \alpha_t(R)[c_2\Delta + V_2^{t-1}(R)] + (1 - \alpha_t(R))[b_2\Delta + d_2(t-1)\Delta - \varepsilon]. \end{aligned}$$

$V_1^{t_2^*+1}$, $V_2^{t_2^*+1}(L)$, and $V_2^{t_2^*+1}(R)$ were defined earlier, completing the recursive definitions. For future use, we note that if we define $V_2^{t_2^*}(R) = V_2^{t_2^*}$ and $V_2^{t_2^*}(L) = a_2t_2^*\Delta$, then the equations for 2's value functions also hold for $t = t_2^* + 1$.

We now show by induction that (7), (8), and (9) must hold for all $t \geq t_2^* + 1$ if Δ is sufficiently small. We have already shown the conclusion for $t = t_2^* + 1$, establishing the basis for the induction. So fix any $t \geq t_2^* + 2$ and suppose that we have shown the result for all smaller t .

To see that (7) holds, substitute for V_1^{t-1} using

$$V_1^{t-1} = c_1(t-1)\Delta + (1 - \beta_{t-1})(b_1 - c_1)\Delta - \varepsilon$$

into (7) and rearrange to obtain

$$0 < (a_1 - c_1)\Delta + (1 - \beta_{t-1})(b_1 - c_1)\Delta < 2\varepsilon.$$

The latter inequality obviously holds if $\Delta < 2\varepsilon/[a_1 - c_1 + \max(0, b_1 - c_1)]$. Hence we only need to show the former inequality. This inequality is

$$a_1 - c_1 > (1 - \beta_{t-1})(c_1 - b_1).$$

Since $a_1 > c_1$, this holds if $b_1 \geq c_1$. So suppose $c_1 > b_1$. Then a sufficient condition for this is $a_1 - c_1 > c_1 - b_1$ or $a_1 + b_1 > 2c_1$, which holds by assumption.

To see that (8) holds, substitute for V_1^{t-1} as above and rearrange to obtain

$$d_1\Delta + (1 - \beta_{t-1})(b_1 - c_1)\Delta < b_1\Delta$$

or

$$c_1 - d_1 > \beta_{t-1}(c_1 - b_1).$$

Recall that $c_1 > d_1$. Hence this certainly holds if $b_1 \geq c_1$. So suppose $c_1 > b_1$. Then a sufficient condition is $c_1 - d_1 > c_1 - b_1$ or $b_1 > d_1$ which holds, implying (8).

The proof of (9) is more involved. For $t \geq t_2^*$, let

$$Z_t = V_2^t(R) - V_2^t(L) - (a_2 - c_2)\Delta - \varepsilon.$$

It is not hard to show by substituting from the definitions of $V_2^t(R)$ and $V_2^t(L)$ for $t \geq t_2^* + 1$ that

$$Z_t = [\alpha_t(L) - \alpha_t(R)] [(b_2 - c_2)\Delta + d_2(t - 1)\Delta - \varepsilon - V_2^{t-1}(R)] - (a_2 - c_2)\Delta. \quad (10)$$

Solving from definitions of the α 's, we obtain

$$\alpha_t(L) = \frac{(d_2 - b_2)\Delta + 2\varepsilon}{Z_{t-1} + (d_2 - b_2)\Delta + 2\varepsilon}$$

and

$$\alpha_t(R) = \frac{(d_2 - b_2)\Delta}{Z_{t-1} + (d_2 - b_2)\Delta + 2\varepsilon}.$$

Hence for $t \geq t_2^* + 1$,

$$Z_t = \frac{2\varepsilon}{Z_{t-1} + (d_2 - b_2)\Delta + 2\varepsilon} [(b_2 - c_2)\Delta + d_2(t - 1)\Delta - \varepsilon - V_2^{t-1}(R)] - (a_2 - c_2)\Delta. \quad (11)$$

We show that there is a $\bar{\Delta} > 0$ such that for all $\Delta \in (0, \bar{\Delta})$, $Z_t > 0$ for all $t \geq t_2^* + 1$.

We first develop a lower bound on Z_t by deriving an upper bound on $V_2^{t-1}(R)$. We claim that if Δ is sufficiently small, then

$$V_2^t(R) \leq b_2\Delta + d_2(t-1) - \varepsilon, \quad (12)$$

for all $t \geq t_2^*$. This proof is by induction. For $t = t_2^*$, the claim follows if

$$c_2 t_2^* \Delta < d_2 t_2^* \Delta + (b_2 - d_2) \Delta.$$

This holds if

$$t_2^* \Delta > \frac{(d_2 - b_2) \Delta}{d_2 - c_2}.$$

As shown above, for Δ small, $t_2^* \Delta$ is approximately $\varepsilon / (c_1 - d_1) > 0$. Hence this must hold for sufficiently small Δ .

So consider $V_2^t(R)$ for $t \geq t_2^* + 1$ and suppose we have demonstrated the claim for all smaller t . By definition of $V_2^t(R)$, the claim holds if

$$b_2 \Delta + d_2(t-1) - \varepsilon \geq c_2 \Delta + V_2^{t-1}(R).$$

or

$$d_2 \Delta + [b_2 \Delta + d_2(t-2) \Delta - \varepsilon] \geq c_2 \Delta + V_2^{t-1}(R).$$

This follows immediately from $d_2 > c_2$ and the induction hypothesis. Hence (12) holds.

Substituting this upper bound on $V_2^t(R)$ into equation (11), we see that for $t \geq t_2^* + 1$,

$$Z_t \geq \frac{2\varepsilon}{Z_{t-1} + (d_2 - b_2) \Delta + 2\varepsilon} (d_2 - c_2) \Delta - (a_2 - c_2) \Delta.$$

Hence $Z_t > 0$ if

$$Z_{t-1} < 2\varepsilon \left(\frac{d_2 - a_2}{a_2 - c_2} \right) - (d_2 - b_2) \Delta. \quad (13)$$

Note also from (10) and the definition of $\alpha_t(L)$ that $a_2 > c_2$ and $Z_{t-1} > 0$ (so $\alpha_t(L) \leq 1$),

$$Z_t < [1 - \alpha_t(R)][(b_2 - c_2) \Delta + d_2 t \Delta - \varepsilon - V_2^t(R)] \quad (14)$$

for $t \geq t_2^* + 1$.

We complete the proof by showing by induction that there is a $B > 0$ such that if Δ is sufficiently small, then for all $Z_t > 0$ and

$$[1 - \alpha_t(R)] Y_t \equiv [1 - \alpha_t(R)][(b_2 - c_2) \Delta + d_2 t \Delta - \varepsilon - V_2^t(R)] < B \Delta$$

for all $t \geq t_2^* + 1$. We showed earlier that $Z_{t_2^*+1} > 0$, so to complete the basis, we find a $B > 0$ such that $[1 - \alpha_{t_2^*+1}] Y_{t_2^*+1} < B \Delta$. We also showed earlier that $\alpha_{t_2^*+1} \in (0, 1)$, so it

is sufficient to show that $Y_{t_2^*+1} < B\Delta$. Substituting the definition of $V_2^{t_2^*+1}(R)$ into the definition of $Y_{t_2^*+1}$ and rearranging,

$$Y_{t_2^*+1} = (d_2 - c_2)\Delta + \alpha_{t_2^*+1}(R)[(b_2 - c_2)\Delta + d_2 t_2^* \Delta - V_2^{t_2^*} - \varepsilon].$$

We showed in Subsection F.1 that for Δ small, $V_2^{t_2^*}$ is approximately $d_2 t_2^* \Delta - \varepsilon$. Choose any B strictly larger than the maximum of $d_2 - c_2$ and $d_2 + b_2 - 2c_2$. Using $\alpha_{t_2^*+1}(R) < 1$ again, we see that the result holds.

To complete the induction, fix $t \geq t_2^* + 2$ and suppose we have shown the result for all smaller t . Because $[1 - \alpha_{t-1}(R)]Y_{t-1}(R) < B\Delta$, we see from (14) at $t - 1$ and (13) that $Z_t > 0$ as long as Δ is less than $[2\varepsilon(d_2 - a_2)]/[(B + d_2 - b_2)(a_2 - c_2)]$. To show the bound on $[1 - \alpha_t(R)]Y_t$, substitute the definition of $V_2^t(R)$ into the definition of Y_t to obtain

$$[1 - \alpha_t(R)]Y_t = [1 - \alpha_t(R)] \{(d_2 - c_2)\Delta + \alpha_t(R)Y_{t-1}\}.$$

Using the induction hypothesis (which implies $\alpha_t(R) \leq 1$),

$$[1 - \alpha_t(R)]Y_t \leq (d_2 - c_2)\Delta + \frac{\alpha_t(R)}{1 - \alpha_{t-1}(R)}B\Delta.$$

Recall that

$$\alpha_t(R) = \frac{(d_2 - b_2)\Delta}{Z_{t-1} + (d_2 - b_2)\Delta + 2\varepsilon}.$$

Since $Z_{t-1} > 0$ by hypothesis,

$$\alpha_t(R) < \frac{(d_2 - b_2)\Delta}{(d_2 - b_2)\Delta + 2\varepsilon}.$$

Hence it is sufficient to show that

$$(d_2 - c_2)\Delta + \frac{(d_2 - b_2)\Delta}{(d_2 - b_2)\Delta + 2\varepsilon} \frac{B\Delta}{1 - \alpha_{t-1}(R)} < B\Delta.$$

Substituting for $\alpha_{t-1}(R)$ from the definition and rearranging gives

$$d_2 - c_2 + \frac{(d_2 - b_2)\Delta}{(d_2 - b_2)\Delta + 2\varepsilon} \left[1 + \frac{(d_2 - b_2)\Delta}{Z_{t-2} + 2\varepsilon} \right] B < B.$$

From the induction hypothesis,⁷ $Z_{t-2} > 0$, so it is sufficient to show

$$d_2 - c_2 + \frac{(d_2 - b_2)\Delta}{(d_2 - b_2)\Delta + 2\varepsilon} \left[1 + \frac{(d_2 - b_2)\Delta}{2\varepsilon} \right] B < B.$$

⁷In the case of $t = t_2^* + 2$, this refers to $Z_{t_2^*}$ which is not covered by the induction hypothesis. However, it is straightforward to show by direct calculation that $Z_{t_2^*} > 0$.

Recall that B was chosen to be strictly larger than $d_2 - c_2$. Since the second term on the left converges to 0 as $\Delta \downarrow 0$, we see that this must hold for Δ sufficiently small, completing the induction. Hence $Z_t > 0$ for all $t \geq t_2^* + 2$.

To complete the proof, then, we see that if Δ is sufficiently small, then in all periods after the first, we move to (D, L) if 1 played D in the previous period and both players randomize otherwise. So consider player 1's payoffs for the game as a function of the first period actions. They are

$$\begin{array}{cc}
 & \begin{array}{c} L \\ R \end{array} \\
 \begin{array}{c} U \\ D \end{array} & \begin{array}{cc} a_1\Delta + V_1^{T-1} & d_1\Delta + V_1^{T-1} \\ c_1T\Delta & b_1\Delta + c_1(T-1)\Delta \end{array}
 \end{array}$$

By (7) and (8), 1's best reply to both L and R is to play D . Hence 1 must begin by playing D . 2's best reply to this is to play L , so (D, L) is played in the first and therefore in all subsequent periods. ■

References

- [1] Beggs, A., and P. Klemperer, “Multi-Period Competition with Switching Costs,” *Econometrica*, 60, 1992, 651–666.
- [2] Benoit, J.-P., and V. Krishna, “Finitely Repeated Games,” *Econometrica*, 53, 1985, 905–922.
- [3] Burdzy, K., D. Frankel, and A. Pauzner, “Fast Equilibrium Selection by Rational Players Living in a Changing World,” University of Tel Aviv working paper, 1996.
- [4] Carlsson, H., and E. van Damme, “Global Games and Equilibrium Selection,” *Econometrica*, 61, September 1993, 989–1018.
- [5] Chakrabarti, S., “Characterizations of the Equilibrium Payoffs of Inertia Supergames,” *Journal of Economic Theory*, 51, June 1990, 171–83.
- [6] Dutta, P., “A Folk Theorem for Stochastic Games,” *Journal of Economic Theory*, 66, June 1995, 1–32.
- [7] Dutta, P., “Collusion, Discounting, and Dynamic Games,” *Journal of Economic Theory*, 66, June 1995, 289–306.
- [8] Fudenberg, D., and D. Levine, “Reputation and Equilibrium Selection in Games with a Patient Player,” *Econometrica*, 57, 1989, 759–778.
- [9] Gale, D., “Dynamic Coordination Games,” *Economic Theory*, 5, 1995, 1–18.
- [10] Gul, F., and H. Sonnenschein, “On Delay in Bargaining with One-Sided Uncertainty,” *Econometrica*, 56, 1988, 81–95.
- [11] Kandori, M., G. Mailath, and R. Rob, “Learning, Mutation, and Long-Run Equilibria in Games,” *Econometrica*, 61, January 1993, 29–56.
- [12] Lagunoff, R., and A. Matsui, “Asynchronous Choice in Repeated Coordination Games,” *Econometrica*, 65, November 1997, 1467–1477.
- [13] Lipman, B., and R. Wang, “Switching Costs in Infinitely Repeated Games,” working paper, June 1999a.
- [14] Lipman, B., and R. Wang, “Investment Dynamics with Switching Costs,” working paper, July 1999b.
- [15] Padilla, A. Jorge, “Revisiting Dynamic Duopoly with Consumer Switching Costs,” *Journal of Economic Theory*, 67, December 1995, 520–530.

- [16] Robson, A., “The Adam and Eve Effect and Fast Evolution of Efficient Equilibria in Coordination Games,” University of Western Ontario working paper, 1994.
- [17] Wang, R., and Q. Wen, “Strategic Invasion in a Market with Switching Costs,” *Journal of Economics and Management Strategy*, 7, Winter 1998, 521–549.
- [18] Young, P., “The Evolution of Conventions,” *Econometrica*, 61, January 1993, 57–84.