Discussion Paper No. 1190

Switching Costs In Frequently Repeated Games*

by

Barton L. Lipman **

and

Ruqu Wang***

Northwestern University

July 1997

* The authors thank Eddie Dekel, Debraj Ray, John Riley, and seminar participants at Boston University, Hong Kong University of Science and Technology, Northwestern, and Waterloo for helpful suggestions, Northwestern University for enjoyable visits during which research was begun, and SSHRCC for financial support.

** Barton Lipman: Department of Economics, University of Western Ontario, London, Ontario, N6A 5B7.

*** Ruqu Wang: Queens's University *and* MEDS, Kellogg Graduate School of Management, Northwestern University, Evanston, IL 60208.

Abstract

We show that the standard results for finitely repeated games do not survive the combination of two simple variations on the usual 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 show that this can yield multiple subgame perfect equilibria in games like the Prisoners' Dilemma which normally have a unique equilibrium. Also, it can yield uniqueness in games which normally have multiple equilibria. For example, in a two by two coordination game, if the Pareto dominant and risk dominant outcomes coincide, the unique subgame perfect equilibrium for small switching costs and frequent repetition is to repeat this outcome every period. Also, in a generic Battle of the Sexes game, there is a unique subgame perfect equilibrium for small switching costs.

1 Introduction

The basic facts about subgame perfect equilibria in finitely repeated games are well known. If, as in the Prisoners' Dilemma, the stage game has a unique Nash equilibrium, then the unique subgame perfect equilibrium is to repeat this stage game equilibrium every period. If, as in coordination or Battle of the Sexes games, the stage game has multiple equilibria, then it is possible to have periods in which the play does not constitute a Nash equilibrium of the stage game. In particular, if the stage game equilibria are Pareto ranked, one obtains a folk theorem as the number of repetitions grows large.[1] Also, given the usual focus on total (or average), not discounted, payoffs in such games, these results depend only on how many times the game is played, not the length of time over which these repetitions occur.

In this paper, we consider two seemingly small changes from the usual framework and show that these changes overturn both of these results. First, we add a small cost to changing actions from one period to another. To keep the analysis as close as possible to the standard repeated game model, we treat this cost 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 one of the games we consider, firms choose between investing and not investing in each period. It seems quite reasonable to believe to switching from not investing to investing requires a certain fixed set up cost. Similarly, shutting down an investment may also incur costs. In short, the existence of such costs seems plausible for many economic settings, so the inclusion of small switching costs seems to be a very 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. 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

[1] See Benoit and Krishna [1985] for details.

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.

We show that small switching costs in frequently repeated games overturn both of the standard results for finitely repeated games. In particular, we can have multiple subgame perfect equilibria in games like the Prisoners' Dilemma and unique equilibria in games like coordination or Battle of the Sexes games.

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.

In the Prisoners' Dilemma, the result we get is that under certain conditions on the payoffs, there are multiple subgame perfect equilibria. In particular, cooperation in each of the finitely many periods is possible in a subgame perfect equilibrium. In two by two coordination games, we show that if the Pareto dominant and risk dominant outcomes coincide, then the unique subgame perfect equilibrium outcome consists of repeating this one–shot outcome in every period. If the Pareto and risk dominant outcomes differ, this result is no longer true — there will necessarily be multiple equilibria. In generic Battle of the Sexes games, we obtain a unique subgame perfect equilibrium outcome. While we give the precise statement later, loosely, the outcome here is that the player whose less preferred equilibrium is worse gets his favorite equilibrium every period. Finally, we show that the coordination game results generalize to an n player coordination game, the Investment Game, studied by Gale [1995].

To see the intuition, note that for any strictly positive switching cost, there will be a point near enough to the end of the game that it is not optimal to change actions regardless of what actions are being played. In this sense, all action profiles become *frozen*. As we move further away from the end, certain changes of action become optimal and so some profiles "melt," while others remain frozen. In this phase of the game, the players have an incentive to push play toward frozen profiles which benefit them. In the repeated Prisoners' Dilemma, mutual cooperation can be such a profile, while in coordination games, the Pareto dominant outcome can play this role. In Battle of the Sexes games, of course, the players will disagree about which profile they favor. The key is whether such "good" profiles are frozen earlier in the game than less desirable profiles. If the good profiles are frozen sooner, then the less desirable profiles can be abandoned in favor of the better payoffs, regardless of whether these better payoffs would be a Nash equilibrium without switching costs. In the Prisoners' Dilemma, this means that the usual backward induction argument breaks down, allowing cooperation throughout the game. In coordination games, we get uniqueness because the players anticipate that the "good" equilibrium will eventually dominate and so have an incentive to start there and avoid switching costs later. Similarly, in generic Battle of the Sexes games, one player will necessarily be able to push the outcome to his favorite equilibrium.

In addition to the papers mentioned above, there are several strands of the literature related to our work. 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 [1996]. 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). Similarly, 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 [1995] 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 [1995], 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.[2]

The next section contains the model. Section 3 gives our results on obtaining unique equilibria in games which have multiple equilibria in the absence of switching costs. We give some reasonably general results for two by two games and one result for a more

[2]Burdzy, 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.

"economic" example, namely Gale's [1995] Investment Game. Section 4 gives our results on obtaining multiple equilibria in games with a unique equilibrium in the absence of switching costs. Concluding remarks are offered in Section 5.

2 Model

Fix a finite normal form game, $G = (N, A, u)$ where $N = \{1, \ldots, n\}$ is the set of players, A_i the set of pure strategies for i, and $u_i : A \to \mathbf{R}$ the payoff function for i. 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, ℓ a length of time, and τ for the length of time remaining in the game. In particular, note that $t\Delta$ is the length of time it takes for t periods to pass while at period t, the length of time remaining, τ, is $t\Delta$.

3 Uniqueness Results

In the first subsection, we consider two by two games with two strict Nash equilibria and give sufficient and almost necessary conditions for a generic such game to have a unique equilibrium outcome for small switching costs and frequent repetition. In the following subsection, we turn to Gale's [1995] Investment Game as a more economic example.

3.1 Two by Two Games

Let Here $n = 2$ and $A_i = \{L, R\}$. We denote the payoffs in the stage game G by

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

Throughout, we wish to characterize equilibria of "generic" games and so will rule out a variety of linear relationships among these payoffs. Also, because we are interested in games which have multiple equilibria in the absence of switching costs, we assume G has two strict Nash equilibria, taken to be (L, L) and (R, R) without loss of generality. Hence we assume $a_i > c_i$ and $b_i > d_i$ for $i = 1, 2$. For generic payoffs, the players cannot be indifferent between these two equilibria. Without loss of generality, we assume that player 1 prefers (L, L) to (R, R), so $a_1 > b_1$. Finally, for reasons that will be clear shortly, we make the further genericity assumption that $a_i - c_i \neq b_j - c_j$ for any i, j.

We will call this a *coordination game* if player 2 also prefers (L, L) to (R, R) — that is, if $a_2 > b_2$ — and a *Battle of the Sexes game* if $a_2 < b_2$.

Finally, say that L is *risk dominant* for player i if L is the best reply to a 50–50 mixture by the opponent — that is, if $a_i - c_i > b_i - d_i$. Say that R is risk dominant for i if the opposite strict inequality holds. By our genericity assumption, one action must be risk dominant for each player.

It is easy to see that 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 1 *Assume G is a coordination game. If L is risk dominant for each player, then there is a $\bar{\varepsilon} > 0$ 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^ε_Δ is (L, L) every period.*

Also:

Theorem 2 *Assume G is a Battle of the Sexes game. Then there is a $\bar{\varepsilon} > 0$ and $K > 0$ such that for almost every $\varepsilon \in (0, \bar{\varepsilon})$, for all $\Delta \in (0, K\varepsilon)$, G^ε_Δ has a unique subgame perfect equilibrium outcome. If $a_2 - c_2 > b_1 - d_1$, this unique outcome is (L, L) every period. If $a_2 - c_2 < b_1 - d_1$, it is (R, R) every period.*

Before explaining the proofs, we offer a few comments. First, the restriction to "almost all" ε is used to avoid certain ties in payoffs. We use this only to ensure that we do not have a player indifferent between paying the switching cost and not doing so at a certain key juncture. This indifference can create new equilibria under certain conditions and we simplify matters by eliminating them. To understand this restriction, recall that the game is played over the time interval $[0, 1]$ 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.

Second, the theorems do not imply that there is a unique subgame perfect *equilibrium*, only a unique *outcome*. It turns out that we do not need to work out equilibrium strategies for every possible subgame to characterize the unique outcomes.

Finally, note that our constraint on Δ is that it is less than some multiple of ε. The proof we give actually shows that given ε, if Δ is sufficiently small, then the stated result holds. 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 the argument is 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.

Full proofs of these results are contained in the Appendix, but it is not hard to see the basic intuition. First, consider the coordination game result. For simplicity, suppose $c_i = d_i = 0$, $a_i = 2$, and $b_i = 1$ for $i = 1, 2$. It is easy to see that L is risk dominant for both players. Suppose Δ is very small so that players can change actions very frequently. Consider a period near the end of the game for which both players used action L in the previous period. Clearly, if we are close enough to the end of the game, then even if i expects the other player to switch actions, he will not change his action. If the length of time remaining is τ, the maximum payoff gain i can expect for changing actions is τ, so if $\tau < \varepsilon$, he will not change actions. Similarly, suppose both players used action R in the previous period. Again, if the length of time remaining is sufficiently small, i will not change actions even if he expects his opponent to change. Notice, though, that if τ is the length of time remaining, the maximum payoff gain is now 2τ since switching to match the opponent would earn 2τ while being miscoordinated earns 0. Hence we can only be sure that neither would switch if $\tau < \varepsilon/2$. In this sense, both the (L, L) and (R, R) profiles are eventually *frozen* in the sense that it cannot be optimal for players to switch away from either profile, but the (L, L) profile is frozen earlier.

So suppose the length of time remaining is between ε and $\varepsilon/2$. As shown above, if both played L the previous period, neither will change actions for the rest of the game. Suppose instead that one player played L the previous period while the other played R.

The argument above shows that $\tau < \varepsilon$ implies the player who used L will not find it optimal to change actions, regardless of his beliefs about the opponent. Because of this, $\tau > \varepsilon/2$ implies that the player who used R in the previous period must switch actions.

To complete the argument, suppose the length of time remaining τ is slightly more than ε. Suppose one player used L in the previous period. The above establishes that if he continues with L just a little while longer, till the time remaining falls below ε, then the outcome will be (L, L) from then on. This may cause the player to earn zero until the time left is ε, but if the length of time he earns zero is small enough, the saving of switching costs and the ability to earn 2 for the last part of the game must make this worthwhile. As a result, we see that the "commitment" to L actually extends earlier in the game to some point where more time than ε is left. Because of this, the conclusion that when one played L and the other played R the previous period, the latter switches also extends to earlier in the game. It is not hard to show that this works backward to the very beginning of the game. That is, a player who plays L in the first period *always* plays it thereafter. Consequently, if only one player uses action L in the first period, the other will certainly match this in the second period and from then on. It is not hard to show that in light of this, both players must start with action L.

The result for Battle of the Sexes is based on a similar intuition. To understand the differences, consider the following Battle of the Sexes game:

	L	R
L	3, 2	0, 0
R	0, 0	1, 5

The analysis is more complex than the above, so we give a less detailed explanation. Intuitively, the outcome hinges on what happens when the profile (L, R) is reached at a point late in the game but before it is frozen. This profile is critical because it is reached when each player uses the action associated with his preferred equilibrium and so reactions here will tell us which player can force the other to his favorite outcome. Note that if the time remaining is less than ε, then it will not be worthwhile for player 1 to switch actions even if he knows 2 will not change. On the other hand, if the time remaining is greater than $\varepsilon/2$, it will be worthwhile for player 2 to change actions. Hence if the time remaining is between these two, player 1 can force the outcome to (L, L) by playing L. Just as above, this effect works backward, enabling him to force this outcome from outset.

Given that the Battle of the Sexes result needs no additional condition on payoffs, it is natural to wonder if the risk dominance of L is needed for the coordination game result. We suspect it is not necessary, but some condition along this line certainly is necessary. To see the point, consider a symmetric coordination game (that is, $a_1 = a_2 = a$, etc.) and suppose R is risk dominant for both players, so $b - d > a - c$. It is not difficult to construct two subgame perfect equilibria, one with outcome (R, R) every period and the other with outcome (L, L) in every period.[3] The key to the risk dominance condition is that the risk dominant profile is the one which is frozen first in the sense that no player would have an incentive to switch away regardless of what the opponent does. 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, 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.

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 some profiles are frozen earlier than 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 result 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 possible (0) to the best possible (2) but does so only for a single period is not worthwhile. More generally, as both theorems indicate, we only need ε/Δ to be sufficiently large.

3.2 The Investment Game

In this subsection, we turn to a more economic example, an n player coordination game studied by Gale [1995]. Now there are n players and two actions each, called invest (I) and wait (W). The payoff to waiting is always 0, while the payoff to investing depends on the total number investing. Let $a(k)$ be the payoff to an investor when k agents (including himself) invest. Assume that $a(n) > 0 > a(1)$ and that $a(n) \geq a(k) \geq a(1)$

[3] In both cases, the strategies are such that each player plays his part of the equilibrium, switching to this action if need be, for every history which is not too late in the game.

for all $k \leq n$.

This game is very similar to that considered by Gale [1995]. Gale's model differs in four ways. First, he considers an infinite horizon with discounting. Second, he has stronger assumptions on $a(\cdot)$.[4] Third, he has no switching costs. Finally, he makes the assumption that once a player invests, he must always invest thereafter. In effect, Gale assumes that the cost of switching from waiting to investing is zero, while the cost of switching in the other direction is infinite.

For our purposes, Gale's main result is the following. Fix any $\ell > 0$ and let $\eta_\Delta(\ell)$ be the supremum over the set of subgame perfect equilibria of the probability that the length of time till all agents invest is greater than ℓ. Then for any $\eta \in (0, 1)$, for all Δ sufficiently small, $\eta_\Delta(\ell) < \eta$. That is, if Δ is sufficiently small, all agents invest almost immediately almost for sure.

It is important to note that even for very small Δ, there are equilibria in Gale's model with delay. To see this, suppose that the strategies are that every agent invests starting at the second period regardless of what happens in the first period. Clearly, no agent has any incentive to invest in the first period since he earns $a(1) < 0$ in that period, while simply waiting till the second period avoids this negative return in the first period and has no effect on subsequent payoffs. Of course, as the length of the period goes to zero, the length of delay in this equilibrium goes to zero.

When $n = 2$, the Investment Game is a coordination game as defined in the previous subsection. It is easy to see that investing is risk dominant if and only if $a(2) + a(1) > 0$. Hence in this case, Theorem 1 tells us that if Δ is sufficiently small relative to ε, then the unique subgame perfect equilibrium outcome is for both players to invest in every period, a result stronger than Gale's for his game.

We get a stronger result because of the assumption that there is no switching cost associated with the first choice of an action. Intuitively, our result, like his, shows that delay must be small in any subgame perfect equilibrium. In our case, we can then show that both players must invest immediately because the players will invest soon and so it is better to avoid the switching cost later by investing now. If we assume that the game begins with a "default" of not investing, so that there is a cost ε of investing even in the very first period (and no cost to waiting in the first period), then we do not get this strengthening.

The following theorem, proved in the Appendix, generalizes this result to more than two agents.

[4]These assumptions are used only for results unrelated to the issues considered here.

Theorem 3 *If $a(n) + a(1) > 0$, then for almost every sufficiently small $\varepsilon > 0$, there is a $\bar{\Delta}_\varepsilon > 0$ such that for all $\Delta < \bar{\Delta}_\varepsilon$, the unique subgame perfect equilibrium outcome of G^ε_Δ is that every player invests in every period.*

The intuition of the result is very simple. Gale's result does not really require the infinite horizon (since, after all, the result shows that investment must occur almost immediately), only the assumption that if a player ever invests, he must invest from then on. In our model, it is clear that once we are close enough to the end of the game, the switching costs imply that no investor would have an incentive to stop investing, so we replicate this feature of Gale's model late in the game. By itself, this is not enough. If $a(n)$ is not large enough, then by the time we reach this point, there may be too little time left for noninvestors to be willing to pay the costs to start investing, even if they believe everyone else will invest. In particular, the key turns out to be whether $a(n)+a(1)$ is positive or negative. If it is positive, then when the length of time left is such that a lone investor is indifferent between shutting down and not, a lone noninvestor would strictly prefer to start investing. In other words, if $a(n) + a(1) > 0$, we have a window of opportunity where investors are frozen but noninvestors are not. In this window, Gale's reasoning shows that the noninvestors must start investing almost immediately with probability close to one. But this means that just before this window, investors won't stop investing because they would start investing again soon anyway, so it is not worth paying the switching cost twice. Hence we can push back the date at which investors become frozen. This enables us to push back the date at which noninvestors must start investing, giving us an induction which brings us back to the beginning of the game. But then players will invest at the very beginning, rather than wait and incur the switching cost when they begin investing.

As this intuition also suggests, if $a(n) + a(1) < 0$, we do not get this result. One can adapt the examples of the preceding section to show that if $a(n) + a(1) < 0$, then no player ever investing is a subgame perfect equilibrium outcome.

As noted above, this last statement relies on the assumption that there are no switching costs in the first period. If we assume that waiting is the "default" action at the beginning of time, so that "switching" costs are incurred from investing in the first period, we do not get the strong result that all agents invest immediately for sure. We do still get the analog of Gale's result that the length of the lag till all agents invest is arbitrarily small with arbitrarily high probability.

4 Multiplicity Results

In this section, we show how switching costs can create equilibria, illustrating the point with the finitely repeated Prisoners' Dilemma. As we show, the possibility of attaining multiplicity in this context hinges on an unusual payoff condition similar to that of risk dominance. So consider the Prisoners' Dilemma with payoffs

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

where $c > a > b > d$.[5] We say that *cooperation freezes first* if $b - d > c - a$. Intuitively, when this holds, the incentive to defect is larger when the opponent is defecting than when the opponent is cooperating. In this sense, mutual cooperation can be stable at a point in the game where one player cooperating and one defecting is not.

In this section, we show that mutual cooperation can be sustained in a subgame perfect equilibrium for small ε and Δ if cooperation freezes first. On the other hand, regardless of whether cooperation freezes first or not, mutual defection is always an equilibrium. When cooperation does not freeze first, we get the standard conclusion that mutual defection is the unique equilibrium outcome. Hence cooperation can be sustained in equilibrium if and only if cooperation freezes first.

It is useful to define some notation for these results. Let $s_C = \varepsilon/(c - a)$ and $s_D = \varepsilon/(b - d)$. Note that $s_C > s_D$ iff cooperation freezes first.

Theorem 4 *If cooperation freezes first, then there is a $K > 0$ such that for all $\varepsilon \in (0, M(c - a))$, for all $\Delta \in (0, K\varepsilon)$, there is a subgame perfect equilibrium where both players cooperate in every period.*

Proof. Assume $\varepsilon < M(c - a)$, so $s_C < M$. By the assumption that cooperation freezes first, $s_C > s_D$. Assume that Δ is sufficiently small that there are periods t satisfying $s_D < t\Delta < s_C$.

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 $t\Delta \geq s_D$, i defects, switching to this action if need be. Finally, for any later period, i does whatever he did the previous period.

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

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 $t\Delta < s_D$. 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 $dt\Delta > bt\Delta - \varepsilon$, depending on whether the opponent is cooperating. But by assumption, $t\Delta < \varepsilon/(b-d) < \varepsilon/(a-c)$, so both inequalities hold. Hence it is never optimal to change actions at such a period.

Now suppose $t\Delta \geq s_D$. 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 $t\Delta \geq \varepsilon/(b-d)$. Hence the specified strategy is optimal.

Second, suppose both players cooperated in the previous period, that $t\Delta \geq s_D$, and $(t-1)\Delta < s_D$. Player i expects the opponent to cooperate from now on regardless of what he does because actions will be frozen beginning in the next period. 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 $t\Delta \leq s_C$. But since t must be the last period such that $t\Delta \geq s_D$, we must have $t\Delta < s_C$, so this holds.

Finally, suppose we are at a period t such that $(t-1)\Delta \geq s_D$ and both players cooperated in the previous period. 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. Hence for small Δ, these strategies form a subgame perfect equilibrium. ▌

So mutual cooperation can be supported as a subgame perfect equilibrium outcome if cooperation freezes first. On the other hand, it is easy to see that mutual defection can *always* be supported also. To see this, construct an equilibrium as follows. Player i defects in every period *unless* both agents cooperated in the previous period and $t\Delta < s_C$ or i alone cooperated in the previous period and $t\Delta < s_D$. 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 cooperation does not freeze first, then mutual defection is the unique subgame perfect equilibrium outcome. More specifically,

Theorem 5 *If cooperation does not freeze first, 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 cooperation freezes first or not, 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 it is clear that the usual backward induction arguments do not apply. This does not explain, however, the role played by whether cooperation freezes first or not. To understand this, note that if cooperation freezes first, then at the point where cooperation first freezes, a player is willing to switch to defection if the opponent is defecting. This means that a deviation from mutual cooperation will lead the opponent to switch to defection. On the other hand, if cooperation does not freeze first, then at the time cooperation freezes, deviations from cooperation *will not* be punished. As a result, each player has an incentive to switch to defection just before cooperation freezes. 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 Conclusion

In summary, two seemingly minor variations on finitely repeated games overturn the standard results. Small switching costs in games that are repeated sufficiently frequently can lead to multiple equilibrium outcomes in games which usually have a unique equilibrium or uniqueness in games which usually have many equilibria.

An interesting open question concerns games where the equilibrium outcome set with small switching costs and frequent repetition is actually disjoint from the usual equilibrium outcome set. While we show by example below that such games exist, we have no real characterization of such games or their equilibria.

For such an example, suppose the stage game G is given by

	L	R
L	10, 2	1, 0
R	9, 10	0, 4

Note that each player has a dominant strategy of L. Hence the unique subgame perfect equilibrium of the usual finitely repeated game has (L, L) played in every period. It is not hard to show that this is *not* an equilibrium outcome of the game with small ε and Δ. To see this, fix a small ε. If Δ is sufficiently small, there must be a point late in the game where there are no further changes of actions. Given that the largest payoff gain to either player from a unilateral change of action is 6 and is generated by player 2 switching to L from (R, R), we see that when $t\Delta < \varepsilon/6$, no one changes actions. Let t^* be the smallest t such that $t\Delta > \varepsilon/6$ and assume that

$$(t^* - 1)\Delta < \varepsilon/6.$$

Assume Δ is small, so that $(t^* + 3)\Delta$ is very close to $\varepsilon/6$. Then at period t^*, no one changes actions unless (R, R) were played at $t^* + 1$ and only player 2 changes actions in this case.[6]

In light of this, consider period $t^* + 1$. It is easy to show that no one changes actions if (L, L) or (R, L) were played the previous period and that, again, only 2 changes actions if (R, R) were played. Suppose, though, that (L, R) were played. If neither player changes actions, then they will remain at this point for the rest of the game. This cannot be an equilibrium in the subgame as player 1 would prefer to switch to R to induce player 2 to switch in the subsequent period. This change gives player 1 a gain of 9 per remaining period, which is certainly worthwhile. It is also not an equilibrium for both players to switch actions at this point. If 2 does change, 1 prefers to not change his action in order to get a payoff of 10 per remaining period and avoid the switching cost. It is not an equilibrium for 1 only to change action since if 1 is changing, 2 will change simultaneously to begin earning 10 sooner. Finally, it is not an equilibrium for 2 only to change action since there would be no further changes and he would only gain 2 per remaining period. In short, the players must randomize at this point.

It is tedious but not difficult to calculate the unique mixed strategies which must be used at this point. One can use this to show that at $t^* + 2$, if (L, L) were played in the previous period, it is *not* an equilibrium for the subgame for no one to change actions. If 1 does not change at this point, it is optimal for 2 to unilaterally change to induce the mixed strategy continuation from (L, R) at $t^* + 1$. Roughly, if Δ is small, these mixed strategies have player 1 changing action with very high probability so 2 has a very good chance of obtaining 10 per period from $t^* + 1$ onward, well worth the switching cost.

We do not have a complete characterization of the equilibria of this game as the randomization makes the analysis quite complex. However, this example clearly shows that the equilibrium set under small switching costs and frequent repetition can be completely disjoint from the equilibrium without switching costs, further underscoring our message.

[6] It is not optimal for 1 to also change actions because given player 2's behavior, the change only increases 1's payoff per remaining period from 9 to 10.

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 1. 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 $M \to \infty$. 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. On the other hand, we do not know whether the results carry over in any natural way to infinitely repeated games.

Appendix

A Proof of Theorem 1

Because L is risk dominant for each i, we have $a_i - c_i > b_i - d_i$ for $i = 1, 2$. Fix any $\varepsilon \in (0, M \min\{b_1 - d_1, b_2 - d_2\})$. Clearly, if Δ is sufficiently small, then no player changes his action in the last period, $t = 1$, since the gain to doing so is proportional to Δ. Fix any period t and suppose that whatever action either player uses at t, he will play that action from then onward. (We know this is true at $t = 2$ for Δ sufficiently small so such values of t exist.) Then i will play the same action at t that he played at the previous period $t + 1$ iff the change in payoff is strictly less than ε. To be more precise, suppose i used L at $t + 1$. Let q be the probability that j plays L at t. Then i's unique best strategy is L if

$$qa_i t\Delta + (1 - q)d_i t\Delta > qc_i t\Delta + (1 - q)b_i t\Delta - \varepsilon$$

or

$$\varepsilon > [q(c_i - a_i) + (1 - q)(b_i - d_i)]t\Delta.$$

By risk dominance, the right-hand side is strictly decreasing in q. Hence this holds for all q if

$$t\Delta < \frac{\varepsilon}{b_i - d_i}.$$

Let $s_L^i = \varepsilon/(b_i - d_i)$.

Similarly, again letting q be the probability that j plays L at t, suppose that whatever either player plays at t is what he will play from then on. Suppose i played R at $t + 1$. Then his unique best strategy is R at t if

$$qc_i t\Delta + (1 - q)b_i t\Delta > qa_i t\Delta + (1 - q)d_i t\Delta - \varepsilon$$

or

$$\varepsilon > [q(a_i - c_i) + (1 - q)(d_i - b_i)]t\Delta.$$

Since the right-hand side is strictly increasing in q, this holds if

$$t\Delta < \frac{\varepsilon}{a_i - c_i}.$$

Let $s_R^i = \varepsilon/(a_i - c_i)$. Summarizing this discussion, then, we see that whenever $t\Delta < \min\{s_L^1, s_L^2, s_R^1, s_R^2\}$, then whatever action player i used at $t + 1$ is played at t in every subgame perfect equilibrium.

Recall that we are assuming that $a_i - c_i > b_i - d_i$, so $s_R^i < s_L^i$ for $i = 1, 2$. So for every $t\Delta < \min\{s_R^1, s_R^2\}$, both players use the same actions they used at $t + 1$. In other words, for all such t, all action profiles are frozen in the sense that no player changes actions under any circumstances. Recall that we have assumed $\varepsilon < M(b_i - d_i)$, so $M > s_L^i > s_R^i$ for $i = 1, 2$.[7] For concreteness, assume $s_R^2 \geq s_R^1$, though it is straightforward to rewrite the proof for the opposite inequality. Given this, $s_L^i > s_R^i$ for $i = 1, 2$ implies $s_L^i > s_R^1$ for $i = 1, 2$.

Assume for the remainder of the proof that Δ is sufficiently small that there are values of t satisfying

$$s_L^1, s_L^2 > t\Delta > s_R^1.$$

Also, assume that ε is such that there is no integer t satisfying $t\Delta = s_R^1$. Obviously, almost all ε satisfy this requirement. So consider any period t such that $t\Delta > s_R^1$. We claim that three facts are true for any such t. First, if player 2 used L in period $t+1$, he never changes actions again. Second, because of this, if both players 1 and 2 used L in period $t+1$, neither ever changes actions again. In this sense, the action profile (L, L) is *frozen*. Finally, if (R, L) were played at $t+1$, then 1 switches to L at t and the two play (L, L) from then on.

We show all three facts by induction. Fix the smallest t such that $t\Delta > s_R^1$. By our choice of ε, then, $(t-1)\Delta < s_R^1$, so we know that whatever actions are used at t are used in every later period. Also, by assumption, $t\Delta < s_L^i$ for $i = 1, 2$. Hence precisely the argument above shows that if player i used L at period $t+1$, he must play L at period t. This establishes the first two facts for this initial value of t. To show the third, suppose player 2 used L and player 1 used R at period $t+1$. From the previous argument, 2 will not change actions ever again. Hence 1 will switch to L if $a_1 t\Delta - \varepsilon > c_1 t\Delta$ or $t\Delta > s_R^1$ which is true by assumption.

To complete the induction, suppose our claim is true when the length of time remaining is τ or smaller. We now show that there is a length of time, $\bar{\ell} > 0$, such that if the time remaining is less than $\bar{\ell} + \tau$, then if player 2 played L at the previous date, then he must play L from then on. To see this, note that the worst payoff 2 could get from playing L for the rest of the game is approximately $d_2\bar{\ell} + a_2\tau$ since 1 will necessarily switch to R at the first date such that the time remaining is less than τ. This is an approximation because $\bar{\ell} + \tau$ remaining may not be a point at which a change of actions is possible and τ remaining may not be either. However, this approximation becomes exact as $\Delta \downarrow 0$. Suppose instead that 2 uses R at this period. His payoff certainly cannot exceed $\max\{c_2, b_2\}\Delta + a_2(\bar{\ell} + \tau - \Delta) - \varepsilon$ 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_2!). For Δ small enough, this is strictly worse if

$$d_2\bar{\ell} + a_2\tau > a_2(\bar{\ell} + \tau) - \varepsilon,$$

[7] One can weaken our assumption on ε to $\varepsilon < M \min\{a_1 - c_1, a_2 - c_2\}$ at the cost of some notation but with no change in the proof technique.

or $\bar{\ell} < (a_2 - d_2)/\varepsilon$. For Δ sufficiently small, there will necessarily be periods between the point where $\bar{\ell} + \tau$ is left and where the time remaining is τ.

To show the second fact, then, is simple. Consider any period with strictly less time remaining than $\bar{\ell} + \tau$. From the above, we see that if player 2 played L in the previous period, he plays L from then on. Clearly, then, if both players played L in the previous period, it cannot be optimal for player 1 to ever change actions. Hence both will play (L, L) for the rest of the game. Establishing the third fact is also easy. From the above, we see that if player 2 used L in the previous period, he plays L from then on. Because the time remaining strictly exceeds s_R^1, if 2 played L and 1 played R in the previous period, 1 will switch to L and (L, L) is played from then on.

By induction, then, we see that the three facts above are true for the entire game.

Finally, consider the first period of the game. If 2 plays L, his payoff must be at least $d_2\Delta + a_2(M - \Delta)$. For Δ close to zero, this is close to Ma_2. If 2 plays R instead, his payoff cannot be larger than

$$\max\{b_2, b_2\Delta + a_2(M - \Delta) - \varepsilon, c_2\Delta + a_2(M - \Delta) - \varepsilon\},$$

which converges to $M \max\{b_2, a_2 - \varepsilon\}$ as $\Delta \downarrow 0$. Clearly, Ma_2 is strictly larger than this, so for Δ sufficiently small, 2 must play L in the first period. Since the unique best reply to this is to play L every period, we see that the outcome must be (L, L) every period if Δ is sufficiently small. ▌

B Proof of Theorem 2

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 $t\Delta < \varepsilon/(a_2 - c_2)$, 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 $t + 1$. If he uses L again at t, his payoff against L by 2 is $t\Delta a_1$, which is the largest continuation payoff he can get. Hence L at t is a best reply to L by player 2. 1's payoff to L if 2 plays R is $t\Delta d_1$, while his payoff to R is $t\Delta b_1 - \varepsilon$. Hence 1 must play L if $t\Delta d_1 > t\Delta b_1 - \varepsilon$ or $t\Delta < \varepsilon/(b_1 - d_1)$. By assumption, $t\Delta < \varepsilon/(a_2 - c_2) < \varepsilon/(b_1 - d_1)$, so this holds. Hence 1 certainly plays L at t if he played it at $t + 1$. Given this, if (L, L) was played at $t + 1$, 2's best reply is to play L at $t + 1$.

Next, suppose player 2 used R at $t+1$. Clearly, his best reply to R by player 1 at t is to play R since this yields his highest possible continuation payoff of $b_2 t\Delta$. His best reply to L is also R if $t\Delta c_2 > t\Delta a_2 - \varepsilon$ or $t\Delta < \varepsilon/(a_2 - c_2)$ which holds by assumption. Hence player 2 plays R if he played it at $t+1$. This implies that if (R, R) were played at $t+1$, then 1's equilibrium strategy is to play R at t. In short, no one changes actions at t if (L, L), (R, R), or (L, R) were played at $t+1$.

Assume ε is such that there is no t such that $t\Delta = \varepsilon/(a_2 - c_2)$. (Obviously, almost all ε will satisfy this condition.) Assume Δ is small enough that there are values of t satisfying

$$\min\left\{\frac{\varepsilon}{a_1 - c_1}, \frac{\varepsilon}{b_2 - d_2}\right\} > t\Delta > \frac{\varepsilon}{a_2 - c_2}.$$

We claim that if Δ is sufficiently small, then for all t such that $t\Delta > \varepsilon/(a_2 - c_2)$, 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. By the above assumptions, we know that $(t-1)\Delta < \varepsilon/(a_2 - c_2)$, so no one will ever change actions again 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 $t\Delta > \varepsilon/(a_2 - c_2)$ 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, $\Delta d_1 + (t-1)\Delta a_1$. If he plays R instead, his payoff certainly cannot be larger than $\Delta \max\{b_1, c_1\} + \max\{(t-1)\Delta b_1, (t-1)\Delta c_1, (t-1)\Delta a_1 - \varepsilon\} - \varepsilon$. (Recall that $a_1 > b_1 > d_1$.) Hence L is certainly optimal if

$$\Delta d_1 > \Delta \max\{b_1, c_1\} - 2\varepsilon$$

and

$$\Delta d_1 + (t-1)\Delta a_1 > t\Delta \max\{b_1, c_1\} - \varepsilon.$$

The former holds iff $\Delta < 2\varepsilon/[\max\{b_1, c_1\} - d_1]$. Note that the denominator of this expresssion is strictly positive. The latter holds trivially if $d_1 + (t-1)a_1 \geq t\max\{b_1, c_1\}$. Otherwise, it holds iff $\Delta < \varepsilon/[t\max\{b_1, c_1\} - d_1 - (t-1)a_1]$. Note that $a_1 > \max\{b_1, c_1\}$ implies

$$t\max\{b_1, c_1\} - d_1 - (t-1)a_1 > \max\{b_c, c_1\} - d_1.$$

Hence L is certainly optimal if $\Delta < \varepsilon/[\max\{b_c, c_1\} - d_1]$. (Note that this condition is independent of t.) Clearly, given that 1 plays L at t, 2 will certainly play L at t if he played it at $t+1$ since otherwise he will switch back to L at $t-1$ anyway. L must be better since it saves the 2ε in switching costs and earns a higher payoff in period t. 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 difference in his payoff from not switching at t is that he earns a lower payoff at

t if he does not switch. 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 at least $\Delta d_1 + (M - \Delta)a_1$, while playing R at the first period cannot give a higher payoff than $\max\{Mb_1, Mc_1, \Delta b_1 + (M - \Delta)a_1 - \varepsilon, \Delta c_1 + (M - \Delta)a_1 - \varepsilon\}$. For Δ close to zero, his payoff to L must be close to Ma_1, while his payoff to R 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. ▌

C Proof of Theorem 3

Clearly, if Δ is sufficiently small, no player will change actions in the last period since the gain cannot be worth the switching cost. So fix any t such that $t\Delta < \varepsilon/a(n)$. Suppose that every player expects no player to change strategies from t onward. Then it is easy to see that it cannot be optimal for i to change strategies. If i is investing and $k-1$ other agents are, then it is certainly not optimal for i to stop investing if $a(k) > 0$. If $a(k) < 0$, it is optimal for i to stop investing if and only if $-a(k)t\Delta > \varepsilon$ or $t\Delta > -\varepsilon/a(k)$. By assumption, $a(1) < a(k)$, so $-a(1) > -a(k)$. But $a(n) > -a(1) > -a(k)$. Hence $\varepsilon/a(n) < -\varepsilon/a(k)$. So the fact that $t\Delta < \varepsilon/a(n)$ implies that it cannot be optimal to stop investing. Similarly, suppose i did not invest at the previous period and suppose $k-1$ others did invest. Then if $a(k) < 0$, it cannot be optimal for i to invest. If $a(k) > 0$, it is only optimal if $t\Delta > \varepsilon/a(k)$, but this implies $t\Delta > \varepsilon/a(n)$, a contradiction.

Hence in every subgame perfect equilibrium, for all t such that $t\Delta < \varepsilon/a(n)$, all action profiles are frozen in the sense that no player changes his action. Let $s_A = \varepsilon/a(n)$. Throughout, restriction attention to values of Δ such that there is no integer t such that $t\Delta = s_A$. Almost all ε satisfy this requirement.

We now show that a player who invested at $t+1$ always invests at t for all t such that

$$t\Delta < \varepsilon \min\left\{\frac{2}{a(n) - a(1)}, -\frac{1}{a(1)}\right\} \equiv s_I.$$

To see this, consider an agent who invested at $t+1$. Clearly, if he does not invest at t, his payoff cannot be larger than $-\varepsilon + \max\{0, t\Delta a(n) - \varepsilon\}$. Hence if $t\Delta a(1)$ is larger than this, he must invest at t and all subsequent dates. Rearranging yields the inequality above. Just as before, there can be no t such that we have equality in the equation above. Again, assume ε is such that there is no integer t with $t\Delta = s_I$. Assume ε is sufficiently small that $s_I < 1$ and $s_A < 1$.

The critical implication of $a(n)+a(1)>0$ is that $s_I > s_A$. To see that this holds, note that if the minimum defining s_I is the first term, this is true if $2a(n) > a(n) - a(1)$ or $a(n)+a(1) > 0$. If the minimum is the second term, this holds iff $-a(n) < a(1)$ or, again, $a(n)+a(1) > 0$. Hence for any t such that $s_I > t\Delta > s_A$, any player who invested at $t+1$ invests thereafter, but other changes of action might occur. Assume that Δ is sufficiently small that there are values of t in this range.

For any length of time till the end of the game τ, any number of agents k, and any length of time $\ell \leq \tau$, let $\eta(\tau, k, \ell, \Delta)$ denote the supremum probability that the length of time till all agents invest exceeds ℓ where this supremum is taken over the set of subgame perfect equilibria and the set of histories such that $n-k$ agents invested at the previous date.

We now show that for all τ between s_I and s_A, all $\ell \in (0,\tau)$, and all k, $\eta(\tau, k, \ell, \Delta) \to 0$ as $\Delta \to 0$.

To see this, fix any τ and ℓ in this range. We know from the above that whatever agents invested at the previous period will invest in all subsequent periods. Clearly, then, for any history such that $n-1$ agents invested at the previous period, then $\eta(\tau, 1, \ell, \Delta) = 0$ for all $\ell \geq \Delta$ — that is, with probability 1, the remaining agent will invest at his first possible opportunity (which must come within a length of time equal to Δ), so the result is clearly true for $k=1$.

So suppose we have shown that the result is true for any $k = 1, \ldots, \bar{k}-1$. We complete the proof by showing that the result is true for $k = \bar{k}$. So consider any history such that the number of agents who invested in the previous period is $n - \bar{k}$. Suppose, contrary to the claim, that $\eta(\tau, \bar{k}, \ell, \Delta) \geq \eta^* > 0$ for all Δ. The payoff to any agent who has not yet invested in such an equilibrium must be no greater than $(1-\eta^*)\tau a(n) + \eta^*(\tau-\ell)a(n) - \varepsilon$. If, instead, the agent invests at the next possible opportunity, his payoff must be no worse than

$$\ell a(1) + \eta(\tau, \bar{k}-1, \ell, \Delta)(\tau-\ell)a(1) + [1-\eta(\tau, \bar{k}-1, \ell, \Delta)](\tau-\ell)a(n) - \varepsilon.$$

To see this, note that if all agents invest with a length of time less than or equal to ℓ, the worst case is that *no* other agents invest until after a lag of exactly ℓ. Furthermore, this calculation gives the largest possible probability to the lag exceeding ℓ and assumes that no one else ever invests if the lag is greater than ℓ. Clearly, then, the equilibrium payoff to an agent who has not yet invested must exceed this. But the induction hypothesis tells us that for any $\ell > 0$, we can choose Δ sufficiently small that this is arbitrarily close to

$$\tau a(n) - \varepsilon > (1-\eta^*)\tau a(n) + \eta^*(\tau-\ell)a(n) - \varepsilon,$$

a contradiction.

Hence the claim is proved. Thus we see that for every subgame perfect equilibrium, for every $\ell > 0$ and $\eta \in (0, 1)$, we can choose Δ small enough to make the probability that every agent invests by the time there is $s_I - \ell$ remaining in the game at least $1 - \eta$. Note that the only important aspect of s_I was the fact that it is strictly earlier than s_A and that for every later date t, no agent who invested at $t + 1$ ever stops investing. Hence we have shown that if any length of time remaining, τ, has this property, then for an appropriate choice of Δ, every agent invests by time there is $\tau - \ell$ remaining with probability at least $1 - \eta$ and always invests thereafter.

We now use this to show that every length of time remaining $\tau < 1$ has this property. To see this, note that we have seen that all agents who invest when the length of time remaining is less than s_I, necessarily invest in all later periods and that, as a result, for any $\ell > 0$, if Δ is sufficently small, then every agent invests from time remaining of $s_I - \ell$ onward with arbitrarily high probability.

So consider a length of time ℓ' before s_I remaining and an agent who invested in the previous period. If this agent does not invest at $s_I + \ell'$ remaining, the largest his payoff can be is if he and all other agents immediately begin investing in the following period. Hence his payoff to not investing must be smaller than

$$-\varepsilon + [s_I + \ell']a(n) - \varepsilon.$$

If instead he invests through time remaining of $s_I - \ell$, his payoff must be at least (approximately[8])

$$(\ell' + \ell)a(1) + [s_I - \ell][\eta a(1) + (1 - \eta)a(n)]$$

where η is the probability that all agents invest by time $s_I - \ell$ remaining. Since we can choose Δ to make ℓ and η arbitrarily close to zero, this must hold for Δ sufficiently small if

$$\ell' < \frac{2\varepsilon}{a(n) - a(1)}.$$

So choose any $\ell' \in (0, 2\varepsilon/[a(n) - a(1)])$ and any Δ small enough. Then we know that from the point where the time remaining is $\tau_1 \equiv s_I + \ell'$ onward, any agent who invests once invests in every subsequent period. Using this, we know that all agents must invest by the time there is $\tau_1 - \ell$ remaining with probability at least $1 - \eta$. Hence from the point where the time remaining is $\tau_2 = \tau_1 - \ell'$ onward, any agent who invests must always invest thereafter, etc. Hence from the beginning of the game, any agent who invests in any period always invests thereafter. Hence every agent invests no later than ℓ after the beginning of the game with probability at least $1 - \eta$.

So consider any agent at the beginning of the game. If he does not invest right away, his payoff certainly cannot be larger than $Ma(n) - \varepsilon$, while if he does invest right away,

[8]As in the proof of Theorem 1, there is an approximation here in that we are not being careful about when the periods are. As $\Delta \downarrow 0$, the approximations become exact.

his payoff must be greater than $(1-\ell)[\eta a(1)+(M-\eta)a(n)]$. Clearly, for ℓ and η close enough to zero, it must be optimal to invest from the beginning. ▌

D Proof of Theorem 5

Clearly, $c-a \geq b-d$ implies that $s_D \geq s_C$. Also, it is easy to show that for any t such that $t\Delta < s_C$, every player uses whatever action he used at the previous period $t+1$. Finally, restrict attention to ε such that there is no integer t satisfying $t\Delta = s_D$.

The first claim to establish is that for the last t such that $t\Delta > s_D$, unless t is the first period, both players defect at t and thereafter. Let t^* denote this period. There are two cases to consider. First, suppose Δ is such that there are no values of t satisfying $s_D > t\Delta \geq s_C$ (including, of course, the possibility that $s_D = s_C$). In this case, we must have $(t^*-1)\Delta < s_C$. 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 $t^*\Delta > s_C$ which holds by assumption. Similarly, if his opponent defects at t^*, i should defect at t^* as long as $t^*\Delta > s_D$ 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 $s_D > t\Delta \geq s_C$. 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 $s_D > t\Delta \geq s_C$, any player who defected at the previous period $t+1$ must defect at t and from then on. Furthermore, if i cooperated at the previous period 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, $(t-1)\Delta < s_C$, so whatever actions are played at t will be played in all succeeding periods. Hence i's action at t *cannot* affect his opponent's action in 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 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 $dt\Delta > bt\Delta - \varepsilon$ or $t\Delta < s_D$ which holds by assumption.

To complete the induction, suppose we are at a period t such that $s_D > t\Delta > s_C$ 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. Precisely the same reasoning as above shows that i will continue cooperating.

We now use this to show our claim that if t^* is the last period t satisfying $t\Delta > s_D$ and is not the first period, then both players defect at t^* and thereafter in any subgame perfect equilibrium. The restriction to small values of Δ such that there is no t with $t\Delta = s_D$ implies that $s_D > (t^*-1)\Delta \geq s_C$. 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 $t^*\Delta > s_D$ 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 is implied by $t^*\Delta > s_C$) 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 $t\Delta > s_D$, 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 $t\Delta < s_C$. 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. ▌

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] Dutta, P., "A Folk Theorem for Stochastic Games," *Journal of Economic Theory*, 66, June 1995, 1–32.

[6] Dutta, P., "Collusion, Discounting, and Dynamic Games," *Journal of Economic Theory*, 66, June 1995, 289–306.

[7] Fudenberg, D., and D. Levine, "Reputation and Equilibrium Selection in Games with a Patient Player," *Econometrica*, 57, 1989, 759–778.

[8] Gale, D., "Dynamic Coordination Games," *Economic Theory*, 5, 1995, 1–18.

[9] Gul, F., and H. Sonnenschein, "On Delay in Bargaining with One-Sided Uncertainty," *Econometrica*, 56, 1988, 81–95.

[10] Kandori, M., G. Mailath, and R. Rob, "Learning, Mutation, and Long-Run Equilibria in Games," *Econometrica*, 61, January 1993, 29–56.

[11] Lagunoff, R., and A. Matsui, "An 'Anti-Folk Theorem' for a Class of Asychronously Repeated Games," University of Pennsylvania working paper, 1995.

[12] Padilla, A. Jorge, "Revisiting Dynamic Duopoly with Consumer Switching Costs," *Journal of Economic Theory*, 67, December 1995, 520–530.

[13] Robson, A., "The Adam and Eve Effect and Fast Evolution of Efficient Equilibria in Coordination Games," University of Western Ontario working paper, 1994.

[14] Wang, R., and Q. Wen, "Strategic Invasion in a Market with Switching Costs," Queen's University working paper, 1996.

[15] Young, P., "The Evolution of Conventions," *Econometrica*, 61, January 1993, 57–84.