



**THE PINHAS SAPIR CENTER FOR DEVELOPMENT
TEL AVIV UNIVERSITY**

Agility in Repeated Games: An Example

Ran Spiegler¹

Discussion Paper No. 2-14

February 2014

The paper can be downloaded from: <http://sapir.tau.ac.il>

Thanks to The Pinhas Sapir Center for Development at Tel Aviv University for their financial support.

¹ Ran Spiegler - The Eitan Berglas School of Economics, Tel Aviv University, University College London and CFM. Email: rani@post.tau.ac.il

Abstract

I study a repeated matching-pennies game between players having limited "agility": when player i decides to switch his action, it takes (geometrically distributed) time for the decision to be implemented. I characterize the unique Nash equilibrium in this game. Players obtain the same equilibrium payoff as in the benchmark game with unlimited agility. However, players' equilibrium behavior displays endogenous hysteresis, which is more pronounced for less agile players.

Keywords: guerilla, repeated games, imperfect monitoring, agility, organizational behavior, hysteresis

1 Introduction

Conventional economic theory models long-run interaction between individual players and organizations using essentially the same tools. Yet there are important differences in the decision processes that characterize individuals and organizations in such settings. For example, when a military organization chooses to redeploy its forces or change its fighting strategy, implementing this decision takes time. Similarly, when the leadership of a large business organization reaches a strategic product-design decision, it takes time for the decision to trickle down the organization's hierarchy. The larger the organization, the slower this process of implementing decisions may be.

Two important real-life examples come to mind. Imagine a large regular army fighting against a guerilla military organization. When the two armies fight in the same arena, the large army is likely to win thanks to its superior means. However, the same asymmetry also implies that if the guerilla organization switched to another arena, the regular army could only slowly adapt to the change. The guerilla army's smallness is a disadvantage in head-to-head conflicts, but it is an advantage in terms of agility. Similarly, imagine a software giant competing against a small start-up for a certain market segment. When market demand is stable, the giant's superior infrastructure and marketing capabilities will enable it to smash the start-up. However, if features of market demand change, the start-up may be able to adapt its product more quickly.

To my knowledge, the literature on repeated games has not address such aspects, which are peculiar to organizational decision making (although the model of repeated games with endogenous strategy complexity, due to Rubinstein (1986) and Abreu and Rubinstein (1988), can be interpreted along these lines). This note is a step in this direction. I construct a model of repeated interaction between a strong, heavy organization and a weak, agile organization. As a first step, I focus on repeated matching pennies. The modeling innovation is that I introduce a new parameter for each player, referred to as his agility. Whenever a player chooses to switch his action, implementing this change takes time, according to a geometric distribution. Player i 's agility is thus captured by the single parameter that characterizes this geometric distribution, namely the per-period implementation probability β_i . The more agile player is assumed to be the one who benefits from miscoordination. This captures the idea that strength and heaviness are positively correlated, such that if two organizations fight in the same "arena", the heavier player will prevail thanks to his superior strength. Finally, I assume that each player can perfectly monitor its opponent's past actions, but it does not observe the opponent's decision to change his course of action. Thus, the model

falls into the general category of games with imperfect private monitoring.

I characterize Nash equilibrium in this game, under the assumption that $\beta_1, \beta_2 \geq \frac{1}{2}$. I show that in this case each player has a strategy that induces the stationary observed switching rate of $\frac{1}{2}$. Therefore, the value of the game is exactly as in the benchmark model - i.e., superior agility does not lead to a payoff advantage - and the equilibrium behavior induces a stationary observed switching probability of $\frac{1}{2}$. However, the equilibrium probability of player i 's *decision* to switch is not stationary, but strictly *decreasing* in the number of periods that elapsed since the last realized switch, at a constant rate of β_i per period. Thus, the players exhibit hysteresis in equilibrium - their tendency to stick to their current action becomes stronger over time, and in the limit it becomes overwhelming. Moreover, the less agile player exhibits greater hysteresis. This effect is consistent with the intuition that larger (i.e. heavier) organizations are also more "conservative", in the sense that they are less likely to change their behavior over time.

2 A Model

Two players, 1 and 2, play an infinitely repeated zero-sum game with a common discount factor δ . The set of actions for each player is $\{A, B\}$. Let a_i^t denote player i 's chosen action at period $t = 1, 2, \dots$. Player 1's stage-game payoffs are as follows: $u_1(a_1, a_2) = 1$ (0) when $a_1 = a_2$ ($a_1 \neq a_2$). The departure from the conventional repeated-game model is the following. Suppose that at period t , player i chooses to take a different action than a_i^{t-1} . Then, this decision is not implemented instantaneously. Rather, for every $k = 0, 1, 2, \dots$, the switch is implemented at period $t + k$ with $\beta_i(1 - \beta_i)^k$, where $\beta_i \geq \frac{1}{2}$. I assume that the player is unable to take any action from the moment a decision to switch has been made until the moment the decision has been implemented. If the switch is implemented at some period $t' > t$, the player is free to choose to switch again at $t' + 1$. Fix an arbitrary initial condition (a_1^0, a_2^0) .

I assume that each player perfectly monitors the opponent's past actions and his own entire private history. Thus, when player i observes that player j has played the same action in a number of consecutive periods, he does not know whether this reflects the slow implementation of a decision to switch, or the fact that a decision to switch has not been made at all. In this sense, the game falls into the general class of dynamic strategic models with imperfect private monitoring.

The parameter β_i captures player i 's "agility". A more agile player is represented by a higher value of β . The interpretation is that players are not individuals but *orga-*

nizations. One of the characteristics of a large, heavy organization is that it takes time from the moment a decision to change a course of action is made until the moment it is finally implemented. My objective is to explore the implications of such "heaviness" for repeated interaction between organizations. I assume that $\beta_1 \leq \beta_2$. The interpretation is that agility and strength are negatively correlated. Thus, when both organizations choose the same arena at period t (i.e., $a_1^t = a_2^t$), the heavier organization will beat the more agile organization because it is stronger. When the two organizations choose different arenas (i.e., $a_1^t \neq a_2^t$), the conflict is avoided and thus the outcome is a draw. Player 2's superior agility enables him to avoid head-to-head conflict with his stronger, heavier rival.

The model abstracts from other aspects of limited agility. For example, I stick to the conventional assumption that monitoring the opponent's action is quick, such that every switch in his behavior is immediately detected. In addition, the assumption that players are unable to make any choice while their decision to switch is implemented means that they cannot issue a retraction. One could imagine a model in which the organization chooses to reverse a decision to switch while it is being implemented, and implementing a retraction may also take time. These examples demonstrate that once we open the door for delays in the decision process, many interesting factors could be integrated into the repeated-game model. Hopefully this paper will be a first step in this direction.

3 Analysis

The analysis in this section holds for an arbitrary discount factor. Recall that when $\beta_1 = \beta_2 = 1$, we are back with standard repeated matching pennies. In this case, there is a unique Nash equilibrium, in which each player chooses each action with equal probability at every history. The Minimax Theorem holds, such that this is also the unique max-min strategy for each player.

At any period- t history h^t in which player i gets to act, let $k(i, h^t)$ denote the number of periods that elapsed since his latest observed switched action. That is, player i gets to act at period t , and $a_i^{t-1} = \dots = a_i^{t-1-k(i, h^t)} \neq a_i^{t-2-k(i, h^t)}$. In particular, if $a_i^{t-1} \neq a_i^{t-2}$, then $k(i, h^t) = 0$.

Proposition 1 *There is a unique Nash equilibrium. At any history h in which player i gets to act, he chooses to switch his action with probability $\frac{1}{2}\beta^{k(i, h)-2}$.*

Proof. The method of proof is as follows. I first show that each player has a unique strategy that induces an actual (i.e., observed) switching probability of $\frac{1}{2}$ after every history. Using the Minimax Theorem, this will imply that this must be the player's unique Nash equilibrium strategy.

Let $p_i(k)$ be the probability that player i decides to switch, conditional on the event that k periods elapsed since his last switch and that he chose not to switch in each one of them. Then, for any k , the probability that the player chooses not to switch at all periods $1, \dots, k$ since the last switch is

$$\prod_{k'=1}^k (1 - p_i(k'))$$

The probability that the player does not actually switch at all periods $1, \dots, k$ is

$$\prod_{k'=1}^k (1 - \beta_i p_i(k'))$$

Thus, the probability that player i chose not to switch at all periods $1, \dots, k$ conditional on the observation that he did not actually switch in those periods is

$$x_i(t) = \frac{\prod_{k'=1}^k (1 - p_i(k'))}{\prod_{k'=1}^k (1 - \beta_i p_i(k'))}$$

Note that by definition, this function weakly decreases in k . Our objective is to construct a function $p_i(k)$ such that

$$x_i(k-1) \cdot \beta_i p_i(k) + (1 - x_i(k-1)) \cdot \beta_i = \frac{1}{2}$$

because the L.H.S of this equation is the probability that there is an observed switch at period k conditional on having no observed switch in periods $1, \dots, k-1$. If we can find a function that satisfies this equation for every k , then we will have generated a stationary observed switching rate of $\frac{1}{2}$. Note that $x_i(0) = 0$ by definition. Thus, we

can solve this equation by induction and obtain the unique solution

$$\begin{aligned}
 p_i(k) &= \frac{1}{2}(\beta_i)^{k-2} \\
 x_i(k) &= \frac{2\beta_i - 1}{\beta_i(2 - (\beta_i)^{k-1})}
 \end{aligned}$$

The function $x_i(k)$ is strictly decreasing, and asymptotes at $1 - \frac{1}{2\beta_i}$ when $k \rightarrow \infty$.

A stationary observed switching rate of $\frac{1}{2}$ coincides with the max-min strategy in the case of $\beta_1 = \beta_2 = 1$. Clearly, the set of feasible strategies available to player i when $\beta_i < 1$ is contained in the feasible set for $\beta_i = 1$, because the difference is that there are histories in which he is unable to act and restricted to a fixed switching rate of β_i . Therefore, the max-min payoff for each player is weakly below the benchmark case. Since each player i can generate the same realized switching rate as in the benchmark, using a unique defined strategy, by the Minimax Theorem it follows that this is also the unique Nash equilibrium strategy for the player. ■

Thus, players' equilibrium payoff is exactly as in the benchmark case with unlimited agility. In other words, the heavy player 1 is not harmed by his inferior agility. The players' observed behavior is indistinguishable from the benchmark, too. However, the players' behavior differs markedly from the benchmark. Players become less likely to choose to switch as time goes on, and in the limit they are completely inert. In other words, players display hysteresis. Moreover, the less agile the player, the greater the tendency for hysteresis. This is consistent with the intuition that stronger, heavier organizations are also slower to change their behavior. In other words, greater *exogenous* slowness in *implementing* changes leads to greater *endogenous* slowness in *initiating* them.

4 Concluding Remarks

This note presented a simple example that demonstrates the implications of limited agility in repeated interaction. The example captured in a stylized manner an ongoing military conflict between a "regular" army and a guerilla army, where the guerilla army's advantage in terms of agility enables it to escape head-to-head conflict with the stronger regular army. The next step in this research program is to expand the analysis to a larger class of games, thus capturing other relevant situations. In particular, if we wish to capture the effects of limited agility on organizations' adaptation to changing circumstances (as in the software giant vs. start-up story mentioned in the

Introduction), we will need to introduce an evolving state of Nature. Finally, it would be interesting to incorporate other aspects of agility in organizational decision making, such as the possibility to retract or slowness in monitoring changes in other players' behavior.

References

Abreu, D. and A. Rubinstein (1988), "The Structure of Nash Equilibrium in Repeated Games with Finite Automata," *Econometrica* 56, 1259-1282.

Rubinstein, A. (1986), "Finite Automata Play the Repeated Prisoner's Dilemma," *Journal of Economic Theory* 39, 83-96.