AgEcon SEARCH
RESEARCH IN AGRICULTURAL & APPLIED ECONOMICS

The World's Largest Open Access Agricultural & Applied Economics Digital Library

This document is discoverable and free to researchers across the globe due to the work of AgEcon Search.

Help ensure our sustainability.

Give to AgEcon Search

AgEcon Search
<http://ageconsearch.umn.edu>
aesearch@umn.edu

*Papers downloaded from **AgEcon Search** may be used for non-commercial purposes and personal study only. No other use, including posting to another Internet site, is permitted without permission from the copyright owner (not AgEcon Search), or as allowed under the provisions of Fair Use, U.S. Copyright Act, Title 17 U.S.C.*

History as a Coordination Device

By

Itzhak Gilboa and Rossella Argenziano

Working Paper No. 8-2006

July, 2006

*The Foerder Institute for Economic Research
and
The Sackler Institute of Economic Studies*

History as a Coordination Device*

Rossella Argenziano[†] and Itzhak Gilboa[‡]

February 2006

Abstract

Coordination games often have multiple equilibria. The selection of equilibrium raises the question of belief formation: how do players generate beliefs about the behavior of other players? This paper takes the view that the answer lies in history, that is, in the outcomes of similar coordination games played in the past, possibly by other players. We analyze a simple model in which a large population has to make a simultaneous decision regarding participation in a coup attempt. We assume a dynamic process that faces different populations with such games for randomly selected values of a parameter. We show that history serves as a coordination device, and determines for which values of the parameter a revolution would succeed. We also show that, for certain values of the parameter in question, the limit behavior depends on the way history unfolds, and cannot be determined from a-priori considerations.

1 Introduction

Consider a population of identical individuals who have to make a simultaneous decision regarding participation in a coup attempt. The probability

*We are grateful to Dirk Bergemann, Don Brown, Dino Gerardi, Stephen Morris, Barry O'Neill, Ady Pauzner, Alessandro Pavan and seminar participants at the 2004 Winter Meeting of the Ec PAC Project and the 2005 Cowles workshop on coordination games for comments and references. Gilboa gratefully acknowledges support from the Israel Science Foundation (Grant No. 975/03). This paper is part of the Polarization and Conflict Project CIT-2-CT-2004-506084 funded by the European Commission-DG Research Sixth Framework Programme.

[†]Department of Economics, University of Essex. rargenz@essex.ac.uk

[‡]Berglas School of Economics, Tel-Aviv University, and Cowles Foundation, Yale University. igilboa@tau.ac.il

of success increases with the proportion of individuals who decide to participate. The nature of the problem could be explained in the context of a two-player stag-hunt game. Assume that each player has to decide whether to join the rebellion (R) or to opt out (O).

	O	R
O	7, 7	6, 0
R	0, 6	9, 9

In this two-person example, the rebellion will succeed if and only if both players choose R . In this case, both players will be better off than in the status quo ($9 > 7$). That is, the equilibrium (R, R) Pareto dominates the equilibrium (O, O) . Yet, strategy O guarantees a higher minimal payoff than does R . Indeed, the equilibrium (O, O) risk dominates (R, R) (Harsanyi and Selten (1988)), and may be a reasonable prediction of the outcome of the game even though (R, R) is a Pareto dominant equilibrium.

In this paper we analyze a large population version of this revolution game. In this game, a continuum of players have to make a simultaneous decision regarding their participation in a coup attempt. The probability that the revolution succeeds depends on the proportion of players who decided to join the attempt. As in the game above, everyone will be better off if the revolution succeeds, but in case of failure the attempting rebels will be punished. We will also assume that, should the revolution succeed, every individual has an incentive to be among the rebels rather than to remain obedient.

Consider the decision of a single player in this game. Imagine that rumors have been spreading that the revolution would start tonight. She can ignore the rumors and go to sleep, or take to the streets. For simplicity, assume that this is a one-shot, binary decision. The potential rebel sits at home and attempts to assess the probability that the revolution would succeed. How would she do that?

We maintain that the assessment of this probability would and should be

based on the results of past coup attempts in similar games. These games may have been played by the same population or by others. They may have been more or less similar. Both the nature of the game and the identity of the population playing it should be taken into account in the evaluation of the similarity of past games to the present one. But ignoring these past games would hardly seem a rational way of generating beliefs.¹

In this paper we are interested in a dynamic process, according to which large populations are called upon to play a simple “revolution” game, where the games differ from each other by one parameter at most. This parameter designates the status quo, and the lower it is, the more do the people have to gain from a successful revolution. We assume that players generate beliefs regarding the success of a revolution based on the similarity-weighted relative frequency of successes in the past: they calculate relative frequencies, but each past case is assigned a weight that is proportional to its similarity to the game at hand.

Our dynamic process may explain a “Domino” effect. Consider, for example, the revolutions in the Soviet block in the late ’80s and early ’90s. A successful revolution in one country renders the success of a revolution in another country more likely, and vice versa. This process will also exhibit path-dependence. Assume, first, that the first attempted revolution occurs in a country in which a revolution is almost inevitable, because the present conditions leave people with nothing to lose. A successful attempt in this country would make it more likely that a revolution would succeed in a similar country, even if the conditions in the latter are not as dire. Continuing in this way, one may generate a sequence of successful revolutions.

If, on the other hand, the first revolutionary attempt occurs in a country in which most people have little to gain from a revolution, this attempt might

¹The belief formation process may be embedded in a meta-game, which will also have a flavor of a coordination game. We assume, however, that people have a fundamental tendency to expect the future to be similar to the past. To quote Hume (1748), “From similar causes we expect similar effects.”

fail. If it is then followed by a revolution attempt in a similar country, where conditions are worse yet similar, a revolution in the latter may also fail, and, continuing in this way one may generate a sequence of failed revolutions.

We conclude that history serves as a coordination device. It informs the belief formation process of all individuals, and, being commonly known, it coordinates among them. However, beliefs that are history-dependent may lead to different behavior, depending on the way history unfolds.

The rest of this paper is organized as follows. We first discuss related literature. Section 2 describes the stage game. We devote Section 3 to modeling the way players generate beliefs given history. Finally, Section 4 describes the dynamic process and provides the main result of the paper.

1.1 Related Literature

The game theoretic literature has witnessed many attempts to select equilibria based on the parameters of the game. The equilibrium selection literature includes many notions that are defined by the game itself (see van Damme (1983)), such as the risk-dominance criterion mentioned above. Other types of considerations attempted to embed the game in a dynamic process (Young (1993), Kandori, Mailath, and Rob (1993), Burdzi, Frankel and Pauzner (2001)) or in incomplete information set-up (Carlsson and van Damme (1994)).

It is noteworthy that risk dominance has emerged as the preferred selection criterion based on quite different types of considerations. On the other hand, the literature on network externalities tends to favor Pareto dominant equilibria over risk dominant ones (see Katz and Shapiro (1986)). This suggests a more agnostic view, according to which the parameters of the game cannot, in general, predict equilibrium selection. It appears that game theoretic considerations could be used to impose certain restrictions on the possible outcomes, but the actual selection of an equilibrium is often left to history, chance, institutional details, or other unmodeled factors.

The conceptualization of a revolution as a coordination game dates back to Schelling (1960) at the latest. There exist alternative conceptualizations in the political science literature, such as Muller and Opp (1986), who emphasize the public good aspect of a revolution. Yet, the coordination game model of a revolution has been the subject of many studies. Lohmann (1994) studied the weekly demonstrations in Leipzig and the evolution of beliefs along the process. More recently, Edmond (2003) studied the effect of mass media on revolutions, whereas Angeletos, Hellwig, and Pavan (2004) focus on a learning process by which individuals form beliefs. As in Lohmann (1994) and Angeletos, Hellwig, and Pavan (2004), we study the evolution of beliefs in a game that is played repeatedly. However, as opposed to these papers, our game is played by a new population at every stage. Thus, our focus is on the generation of prior beliefs (over other players' actions), based on similar games, rather than on the update of already existing prior beliefs by Bayes's law.

2 The Stage Game

We describe a symmetric two-stage extensive form game G_x depending on a parameter $x \in [0, 1]$. There is a continuum of players $[0, 1]$. In stage 1 all players move simultaneously. The set of moves for each player i is $S_i = \{0, 1\}$, where 1 stands for participation, and 0 – for opting out.

In stage 2, after each player determined her move in $\{0, 1\}$, nature chooses a move in $\{F, S\}$, which stand for *Failure* and for *Success* of the revolution, respectively. Nature's move depends on the set of players choosing 1 in stage 1, $A \subset [0, 1]$. Specifically, if A is Lebesgue-measurable, we assume that nature chooses S with probability $\lambda(A)$, where λ stands for Lebesgue's measure. If A is non-measurable, the probability of nature choosing S can be defined arbitrarily (say, by the inner measure of A). At equilibrium, the set A will be measurable.

After each player determined her choice of participation (0 or 1) and

nature determined the success of the revolution (by the probability $\lambda(A)$), the game is over. The payoff of each player depends only on her own choice of participation, and on nature's move (i.e., on the success of the revolution). The payoff function $u = u_i$ for every $i \in [0, 1]$ is given by the following matrix:

$$\begin{array}{cc}
 & \begin{array}{cc} S(success) & F(failure) \end{array} \\
 \begin{array}{c} 1(YES) \\ 0(NO) \end{array} & \begin{array}{cc} 1 & 0 \\ \frac{x+1}{2} & x \end{array}
 \end{array} \tag{1}$$

where $x \in [0, 1]$ is the parameter of the game.

The interpretation of this matrix is as follows. The worst thing that can happen to an individual in this game is to participate in a failed coup. The result is likely to involve imprisonment, exile, decapitation, and the like. We normalize this worst outcome to 0. The best thing that can happen to an individual is that she participates in a revolution that succeeds. In this case she is a part of a (presumably) better and more just society. We normalize this payoff to 1.

An individual who decides to participate in the revolution therefore decides to bet on its success with the extreme payoff of 0 and 1. Between these extreme payoffs lie the payoffs for an individual who decides to opt out, foregoing the chance of being part of the revolution. The payoff of such an individual still depends on the outcome of the revolutionary attempt. Should this attempt fail, such an individual would get x , which is a measure of the well-being of the people in the status quo. We implicitly assume that such an individual, who did not participate in a failed coup d'etat, will be unaffected by the attempted coup. If, however, the revolution succeeds, even the individuals who were passive will benefit from the new regime. However, not being part of the revolutionary forces, they would not reap the benefits of revolution in its entirety. We choose to set their payoff to the arithmetic average between the full benefit, 1, and the status quo, x .

Observe that if $x = 1$, there is nothing to be gained from a revolution. In this case the well-being of the people in the status quo is just as good

as it could possibly be in the case of a successful revolution, and joining a revolution is a dominated strategy. This is what we would expect the situation to be in a democracy.

If, on the other hand, $x = 0$, the well-being of the people in the status quo is comparable to the well-being of an individual who participated in a failed revolution. This describes a situation in which the people has nothing to lose, as in a situation of starvation. In this case, indeed, joining the revolution is a dominant strategy.

In between, when $0 < x < 1$, lie the cases that are strategically more interesting: in these cases, each individual will be better off if the revolution succeed, but she prefers to be passive if the revolution is doomed to fail. In these situations there is no dominant strategy, and each individual player has to determine her choice depending on her beliefs about the other players's choices.

Assume, then, that an individual i attempts to estimate the expected utility of a player playing 1 (participating in the revolution) versus 0 (opting out). Realizing the nature of the game, the individual knows that the probability of a revolution succeeding is independent of her own choice. Suppose that individual i 's belief over the measure of other individuals who choose 1 is given by a measure μ_i over (the Lebesgue σ -algebra on) $[0, 1]$. That is, for every Lebesgue-measurable set $B \subset [0, 1]$, individual i assigns probability $\mu_i(B)$ to the event that the measure of individuals who eventually choose 1 (with or without herself) lies in B . Individual i 's subjective probability that the revolution would succeed is, therefore,

$$\hat{p}_i = \int_{[0,1]} p d\mu_i(p)$$

That is, individual i is assumed to calculate the overall probability of a successful revolution by Bayes's formula, taking into account the assumption that the probability of success equals the measure of the set of individuals who participate in the coup.

Let $U(k)$ denote the expected payoff for an individual i who chooses strategy $k \in \{0, 1\}$, given beliefs $\hat{p}_i$. That is, $U(1) = \hat{p}_i$ and $U(0) = \frac{x+1}{2}\hat{p}_i + x[1 - \hat{p}_i]$.

It is useful to calculate the *critical belief* $p^*(x)$, that is the minimal value of $\hat{p}_i$ that is necessary for an individual to participate in the revolutionary attempt. If we set $p^*(x) = \frac{2x}{1+x}$, it is easily observed that $U(1) \geq U(0)$ iff $\hat{p}_i \geq p^*(x)$.

At equilibrium, $\hat{p}_i$ is independent of i , and it coincides with the actual probability of a successful revolution. Thus, for every $x \in [0, 1]$ the game has three symmetric Nash equilibria (where $\sigma_i \in \Delta(\{0, 1\})$ denotes player i 's mixed strategy):²

1. $\sigma_i^* = \{1, 0\} \forall i \in [0, 1]$
2. $\sigma_i^* = \{0, 1\} \forall i \in [0, 1]$
3. $\sigma_i^* = (\frac{2x}{1+x}, 1 - \frac{2x}{1+x}) \forall i \in [0, 1]$

²As usual, the mixed equilibrium is rather arbitrary and dynamically unstable. Yet, at this point we do not rule out mixed or asymmetric equilibria.

These equilibria are easy to describe. Revolution will succeed in equilibrium if $\hat{p}_i \geq p^*(x)$, that is, if the subjective belief of each agent about the proportion of individuals who would participate in the coup, $\hat{p}_i$, exceeds the critical belief $p^*(x)$. The revolution would fail otherwise. Simply put, if everyone believes that the revolution is likely to succeed, it will, and if everyone believes that the revolution is unlikely to succeed, it won't. This begs the question, however, of how do individuals form their beliefs.

3 Where Do Expectations Come From?

Our approach to the belief formation question is history- and context-dependent. Specifically, we assume that games of the type G_x above are being played over and over again, by different populations $[0, 1]$, in different countries, at different times, and for different values of x . Yet, the history of similar games played in the past, which is assumed to be common knowledge, determines the beliefs $\hat{p}_i$ of the individuals in question.

More concretely, we assume that time is discrete and that the game G_x is played in every period by a new generation of players. We further assume that at the beginning of each period t nature selects a value for x_t in an i.i.d. manner, according to a known discrete distribution. For simplicity, we assume that the possible values for x are only $\{0, \frac{1}{2}, 1\}$. Thus the process is determined by a probability vector $(p_0, p_{\frac{1}{2}}, p_1)$ where $\Pr(x_t = \alpha) = p_\alpha$ for $\alpha \in \{0, \frac{1}{2}, 1\}$.

In each period, all the players of the current generation have a common memory of all the games that have been played in the past. Before playing, they observe the current state of the world and form an expectation on the probability of a success that is based on the similarity between the current state of the world and the state of the world in previous games that ended, respectively, with a success or a failure.

In particular, denoting by S_t the set of (indices of) past games that ended with a success and with F_t the set of (indices of) past games that ended in

a failure, we assume that for every i and every $t \geq 0$:

$$\hat{p}_{it}(x_1, x_2, \dots, x_t) = \hat{p}_t(x_1, x_2, \dots, x_t) = \frac{\sum_{k \in S_t} s^+(x_k, x_t)}{\sum_{k \in S_t} s^+(x_k, x_t) + \sum_{k \in F_t} s^-(x_k, x_t)} \quad (2)$$

where $s^+(x_k, x_t)$ is a function that measures the degree of support that a success in a game where the state of the world was x_k brings to the possibility that there is a success in the game being currently played and $s^-(x_k, x_t)$ measures the degree of support that a failure in a game where the state of the world was x_k leads to the possibility that there is a failure in the game being currently played. (We normalize the "degree of support" in such a way that a past success only lends support to a prediction of a future success, and a past failure – to a future failure.)

The formula (2) is not well-defined for the first period, $t = 1$. Also, it allows $\hat{p}_t(\cdot)$ to be 0 or 1, if history consists of failures alone, or of successes alone, respectively. We find such extreme beliefs unwarranted. Hence we use Equation (2) only when history contains both successes and failures. Formally, we assume that $t \geq 3$, and that history contains at least one success and at least one failure, so that $\hat{p}_t(\cdot) \in (0, 1)$.

Possible functional forms for $s^+(x_k, x_t)$ and for $s^-(x_k, x_t)$ are

$$s^+(x_k, x_t) = 1 + (x_k - x_t) \quad (3)$$

$$s^-(x_k, x_t) = 1 - (x_k - x_t) \quad (4)$$

Observe that the denominator of $\hat{p}_t(\cdot)$ doesn't vanish for any $x_t \in [0, 1]$ because $\sum_{k \in S_t} s^+(x_k, x_t)$ and $\sum_{k \in F_t} s^-(x_k, x_t)$ are both nonnegative and cannot be simultaneously equal to zero because the first one can be zero only for $x_t = 1$ while the second one can be zero only for $x_t = 0$.

The interpretation of these formulae is as follows. If $x_t = x_k$, the current game t is practically identical to the game that was played in period k in

the past. In this case, the degree of support is normalized to 1: whatever happened in period k (success or failure) lends empirical support to the belief that it is going to occur again in period t .

Assume, now, that $x_t > x_k$. In this case the game played in period k is similar, but not identical to, the game played in period t . The similarity is smaller than in the case $x_t = x_k$, because now (period t) people have a lower incentive to rebel than in the past (period k), since the status quo is more agreeable. Suppose that in period k the revolution succeeded. An individual is expected to reason as follows, "Well, in period k , when people were hungry and had little to lose, they rebelled. But it is still possible that today, when things are better, they won't. Hence, the success of the revolution in period k does lend *some* support to the assumption that people would rebel today, but this support is lower than it would be if the situations were identical." This is captured by the function $s^+(x_k, x_t) = 1 + (x_k - x_t) = 1 - (x_t - x_k)$. The bigger is the difference $x_t - x_k > 0$, the lower is the support that one gets for a successful revolution in period t from a successful revolution in period k . In the extreme case in which $x_t = 1$ and $x_k = 0$, a revolution in period k (when people had nothing to lose) lends no support to the prediction of a successful revolution in period t (when people have nothing to gain).

Next suppose that the revolution in period k failed (retaining the assumption $x_t > x_k$). In this case there is no support for a prediction of a success, but there is support for a prediction of a failure. Past failures make future failures more likely. But to what degree? Here the function $s^-(x_k, x_t) = 1 - (x_k - x_t) = 1 + (x_t - x_k)$ is larger than 1, that is, larger than in the case in which the game of period k were identical to the game in period t . This assumption is supposed to reflect the following reasoning, "The revolution in period t is unlikely to succeed. Even in period k , when people were hungrier, the revolution failed. Why would it succeed now, when the status quo is better?"

Finally, if $x_t < x_k$ the logic is reversed: a success in period k lends

support greater than 1 to a success in period t , because in period t , a-priori, there is a stronger motivation to rebel, and a failure in period k lends support lower than 1 for a failure in period t for the same reason.

Observe that if the same game is repeatedly played over and over again, that is, if $x_k = x_t$ for every k , then the expected probability of a success is simply the observed relative frequency of success in the past. Hence, our belief formation process can be viewed as generalizing empirical frequencies (as in fictitious games, Robinson (1951)) to the case in which the game that is played is not identical to past games, but only similar to them.

The following notation may prove useful. Let N_t^S be the cardinality of S_t , and N_t^F – the cardinality of F_t . Equation (2) can be written as

$$\hat{p}_t(x_1, x_2, \dots, x_t) = \frac{N^S - N^S x_t + \sum_{k \in S_t} x_k}{N^S + N^F + (N^F - N^S) x_t + \sum_{k \in S_t} x_k - \sum_{k \in F_t} x_k}$$

It seems natural that the expected probability of a successful revolution, $\hat{p}_t$, be a monotonically decreasing function of the well-being at the status quo, x_t . Indeed,

$$\frac{\partial \hat{p}_t(\cdot)}{\partial x_t} = \frac{+N^S \left(\sum_{k \in F_t} x_k - 2N^F \right) - N^F \sum_{k \in S_t} x_k}{\left[N^S + N^F + (N^F - N^S) x_t + \sum_{k \in S_t} x_k - \sum_{k \in F_t} x_k \right]^2} < 0 \quad (5)$$

(The last inequality holds because the numerator is the sum of two negative terms).

4 The Dynamic Process

We now wish to study the dynamic process in which at every stage $t \geq 1$ x_t is drawn from $\{0, \frac{1}{2}, 1\}$ according to probabilities $(p_0, p_{\frac{1}{2}}, p_1)$, beliefs are

formed in accordance with equation (2), and an equilibrium in G_{x_t} is chosen by the beliefs $\hat{p}_t(\cdot)$.

Under our assumptions, for almost every history we can predict the outcome of the game by looking at a graph representing the two curves $p^*(x_t)$ (increasing in x_t) and $\hat{p}_t(\cdot)$ (decreasing in x_t). Note that the two curves intersect exactly once in $[0, 1]$ because their difference $\hat{p}_t(\cdot) - p^*(x_t)$ is continuous, strictly decreasing, takes a nonnegative value at zero and a non-positive value at one. Let α_t denote the value such that $p^*(\alpha_t) = \hat{p}_t(\alpha_t)$.

If the current value x_t is to the left of the intersection point, that is, $x_t < \alpha_t$, all players' expectation $\hat{p}_t(\cdot)$ will be above the critical belief $p^*(x_t)$, and they will therefore all play 1, resulting in a successful revolution. If, however, the current x_t is to the right of the intersection point, that is, $x_t > \alpha_t$, then all players will play 0 and there will be a failure with probability 1. If $\hat{p}_t(\cdot) = p^*(x_t)$, namely, $x_t = \alpha_t$, then all individuals are indifferent between playing 0 and playing 1. For simplicity we assume that they break ties in favor of the status quo and choose 0. Observe that this is the only case in which we *assume* that the equilibrium is symmetric. For other values of x_t symmetry is a result of optimization.³

Next, we observe that a state of the process is fully summarized by a matrix of relative frequencies

$$R_t = \begin{array}{c|ccc} & x = 0 & x = \frac{1}{2} & x = 1 \\ \hline 1 & r_{t,10} & r_{t,1\frac{1}{2}} & r_{t,11} \\ \hline 0 & r_{t,00} & r_{t,0\frac{1}{2}} & r_{t,01} \\ \hline \end{array}$$

where $r_{t,ij}$ is the relative frequency, up to time t , of periods in which the game was G_j and all players played i . We are interested in the limit of R_t as $t \rightarrow \infty$.

³Other assumptions are possible, allowing players to select different strategies, and/or to play mixed strategies. Note that in the latter case one has to assume that the law of large numbers holds (see Judd (1985)). The main point of this paper does not depend on these assumptions.

Given our assumption that $t \geq 3$, and that history contains at least one success and at least one failure, $\hat{p}_t \in (0, 1)$, which in turn implies that $r_{t,00} = r_{t,11} = 0$ for all t . Observe that the relative frequencies of the columns are governed only by the selection of x , and are independent of the players' behavior. Hence the only candidates for limit frequencies can be the following matrices:

$$\begin{aligned}
 L_1 &= \begin{array}{c|ccc} & x=0 & x=\frac{1}{2} & x=1 \\ \hline 1 & p_0 & p_{\frac{1}{2}} & 0 \\ \hline 0 & 0 & 0 & p_1 \end{array} \\
 L_2 &= \begin{array}{c|ccc} & x=0 & x=\frac{1}{2} & x=1 \\ \hline 1 & p_0 & 0 & 0 \\ \hline 0 & 0 & p_{\frac{1}{2}} & p_1 \end{array} \\
 L_3 &= \begin{array}{c|ccc} & x=0 & x=\frac{1}{2} & x=1 \\ \hline 1 & p_0 & wp_{\frac{1}{2}} & 0 \\ \hline 0 & 0 & (1-w)p_{\frac{1}{2}} & p_1 \end{array}
 \end{aligned}$$

for $w \in (0, 1)$.

We can finally present our main result.

Theorem 1 *For every $(p_0, p_{\frac{1}{2}}, p_1)$, R_t converges to L_1 or to L_2 with probability 1. For an open and convex set of vectors $(p_0, p_{\frac{1}{2}}, p_1)$, containing $(\frac{1}{3}, \frac{1}{3}, \frac{1}{3})$, there is a positive probability that R_t converges to L_1 , and also a positive probability that R_t converges to L_2 .*

The proof is relegated to an appendix.

The theorem states that games $G_{\frac{1}{2}}$, which are partly similar to games G_0 but also to games G_1 , will eventually be played either like G_0 or like G_1 with probability 1. That is, history will determine the outcome of the non-trivial games: either they all result, in the limit, in successful revolutions (if they end up being played like G_0), or they all result in failed ones (if they end up being played like G_1). The main point of the theorem is its second part, which states that the limit distribution of R_t cannot be computed a-priori.

That is, the fundamentals of the game and of the dynamic process do not suffice for a unique determination of the limit behavior in the non-trivial games. The random process generating the sequence x_t will determine the results of the games $G_{\frac{1}{2}}$. Specifically, if the process starts with a large proportion of $x_t = 0$, then for $x_t = \frac{1}{2}$ most similar games will be found to have resulted in a successful revolution, and therefore the game at stage t will also end in success. This will establish an equilibrium at which revolutions succeed in games $G_{\frac{1}{2}}$. Conversely, if the process starts with many draws of $x_t = 1$, every new game $G_{\frac{1}{2}}$ will be considered similar to games in which revolutions failed, and will therefore also result in a failed revolution.

5 Appendix: Proof of the Theorem

We have observed that the behavior of the process depends on whether $\hat{p}_t(x_1, \dots, \frac{1}{2})$ is above or below $p_t^*(\frac{1}{2}) = \frac{2}{3}$.

Recall that

$$\hat{p}_t(x_1, x_2, \dots, x_t) = \frac{\sum_{k \in S_t} s^+(x_k, x_t)}{\sum_{k \in S_t} s^+(x_k, x_t) + \sum_{k \in F_t} s^-(x_k, x_t)}$$

We simplify notation by defining

$$A_t = \sum_{k \in S_t} s^+\left(x_k, \frac{1}{2}\right) \quad (6)$$

$$B_t = \sum_{k \in F_t} s^-\left(x_k, \frac{1}{2}\right) \quad (7)$$

$$c = p_t^*\left(\frac{1}{2}\right) = \frac{2}{3} \quad (8)$$

$$\hat{p}_t(x_1, \dots, \frac{1}{2}) = \frac{A_t}{A_t + B_t} \quad (9)$$

$$z_t = A_t - cB_t - cA_t \quad (10)$$

so that

$$\widehat{p}_t(x_1, \dots, \frac{1}{2}) > c \Leftrightarrow \frac{A_t}{A_t + B_t} > c \Leftrightarrow \quad (11)$$

$$A_t - cB_t - cA_t > 0 \Leftrightarrow z_t > 0 \quad (12)$$

Suppose we start from $z_t > 0$.

Then,

- with probability p_0 the next draw will be 0 , there will be a success and z_t will change (increase) by $(1 - c)(A_{t+1} - A_t) = (1 - c)(1 + (0 - \frac{1}{2})) = (1 - c)\frac{1}{2} = +\frac{1}{6}$
- with probability $p_{\frac{1}{2}}$ the next draw will be $\frac{1}{2}$, there will be a success and z_t will change (increase) by $(1 - c)(A_{t+1} - A_t) = (1 - c)1 = +\frac{1}{3}$
- with probability p_1 the next draw will be 1 , there will be a failure and z_t will change (decrease) by $-c(B_{t+1} - B_t) = -c(1 - (1 - \frac{1}{2})) = -\frac{c}{2} = -\frac{1}{3}$.

For $z_t > 0$

Prob.	$z_{t+1} - z_t$
p_0	$\frac{1}{6}$
$p_{\frac{1}{2}}$	$\frac{1}{3}$
p_1	$-\frac{1}{3}$

(*)

Similarly, if we start from $z_t \leq 0$ what will happen is that:

- with probability p_0 the next draw will be 0 , there will be a success and z_t will change (increase) by $(1 - c)(A_{t+1} - A_t) = (1 - c)(1 + (0 - \frac{1}{2})) = (1 - c)\frac{1}{2} = +\frac{1}{6}$
- with probability $p_{\frac{1}{2}}$ the next draw will be $\frac{1}{2}$, there will be a failure and z_t will change (increase) by $-c(B_{t+1} - B_t) = -c(1) = -\frac{2}{3}$
- with probability p_1 the next draw will be 1 , there will be a failure and z_t will change (decrease) by $-c(B_{t+1} - B_t) = -c(1 - (1 - \frac{1}{2})) = -\frac{c}{2} = -\frac{1}{3}$.

For $z_t \leq 0$	Prob.	$z_{t+1} - z_t$	(**)
	p_0	$\frac{1}{6}$	
	$p_{\frac{1}{2}}$	$-\frac{2}{3}$	
	p_1	$-\frac{1}{3}$	

We conclude that z_t is a Markov process, which, for non-negative values is governed by (*), and for negative ones – by (**).

To establish convergence, we wish to show that, with probability 1, z_t ends up being always positive or always negative from some point on. To see that this is the case, we note that for every $(p_0, p_{\frac{1}{2}}, p_1)$, (i) the process (*) has a positive drift, or (ii) the process (**) has a negative drift, or both. If (i) holds but not (ii), then z_t will be always positive from some point on with probability 1. Conversely, it will be always negative (from some point on with probability 1) if (ii) holds but not (i). We are left with the case in which both (i) and (ii) hold. In this case, starting with any positive value of z_t , there is a positive probability (independent of t) that z_t will never be non-positive, and vice versa for negative values of z_t . Hence the probability of switching infinitely many times between positive and negative values is zero.

To establish convergence, it remains to note that if, for some T , for all $t \geq T$ we have $z_t \geq 0$, then R_t converges to L_1 , whereas if, for some T , for all $t \geq T$ we have $z_t < 0$, then R_t converges to L_2 .

Finally, we observe that both (i) and (ii) hold when $p_0 + 2p_{\frac{1}{2}} > 2p_1$ and $p_0 - 4p_{\frac{1}{2}} < 2p_1$. This defines a convex and non-empty set of vectors $(p_0, p_{\frac{1}{2}}, p_1)$, of which $(\frac{1}{3}, \frac{1}{3}, \frac{1}{3})$ is a member. $\square$

References

- [1] Angeletos, G.-M., C. Hellwig, and A. Pavan (2004): “On the Dynamics of Information, Coordination, and Regime Change”, miemo.

- [2] Burdzy, K., D.M. Frankel and A. Pauzner (2001): "Fast Equilibrium Selection by Rational Players Living in a Changing World", *Econometrica*, **69**: 163-189.
- [3] Carlsson, H., and E. van Damme (1993): "Global Games and Equilibrium Selection", *Econometrica*, **61**: 989-1018.
- [4] Edmond, C. (2003): "Information and the Limits of Autocracy", mimeo.
- [5] Harsanyi, J.C., and R. Selten (1988): *A general theory of equilibrium selection in games*. Cambridge, Mass. and London: MIT Press, pp. xiii, 378.
- [6] Hume, D., (1748) *An Enquiry concerning Human Understanding*. Oxford: Clarendon Press.
- [7] Judd, K. (1985), "The Law of Large Numbers with a Continuum of Random Variables", *Journal of Economic Theory*, **35**: 19-25..
- [8] Kandori, M., G.J. Mailath and R. Rob (1993): "Learning, Mutation, and Long Run Equilibria in Games", *Econometrica*, **61**: 29-56.
- [9] Katz, M. and C. Shapiro (1986), "Technology Adoption in the Presence of Network Externalities", *Journal of Political Economy*, **94**: 822-841.
- [10] Lohmann, S. (1994), "The Dynamics of Informational Cascades: The Monday Demonstrations in Leipzig, East Germany, 1989-91", *World Politics*, **47**: 42-101.
- [11] Muller, E. N. and K.-D. Opp, (1986), "Rational Choice and Rebellious Collective Action", *The American Political Science Review*, **80**: 471-488.
- [12] Robinson, J. (1951), "An Iterative Method of Solving a Game", *The Annals of Mathematics*, **54**: 296-301.

- [13] Schelling, T. (1960), *The Strategy of Conflict*. Cambridge, MA: Harvard University Press.
- [14] van Damme E. (1983): *Refinements of the Nash equilibrium concept*. Lecture Notes in Economics and Mathematical Systems, vol. 219. Berlin: Springer Verlag.
- [15] Young, H.P. (1993): “The Evolution of Conventions”, *Econometrica*, **61**: 57-84

*THE FOERDER INSTITUTE FOR ECONOMIC RESEARCH and
THE SACKLER INSTITUTE FOR ECONOMIC STUDIES
The Eitan Berglas School of Economics
Tel-Aviv University*

Recent List of Working Papers

1-2006	Philippe Jehiel Ady Pauzner	Partnership dissolution with interdependent values
2-2006	Yoram Weiss Pierre-Andre Chiappori	Divorce, Remarriage and Child Support
3-2006	Eddie Dekel Barton L. Lipman Aldo Rustichini	Temptation-Driven Preferences
4-2006	Elhanan Helpman Gene M. Grossman	Separation of Powers and the Budget Process
5-2006	Elhanan Helpman	Trade, FDI, and the Organization of Firms
6-2006	Manuel Trajtenberg	Innovation Policy for Development: an Overview
7-2006	Leonardo Leiderman Rodolfo Maino Eric Parrado	Inflation Targeting in Dollarized Economies
8-2006	Itzhak Gilboa Rossella Argenziano	History as a Coordination Device

Foerder Institute for Economic Research, Tel-Aviv University, Tel-Aviv, 69978 Israel, Tel: 972-3-640-9255;
fax: 972-3-640-5815; e-mail: foerder@post.tau.ac.il
Papers from 2000 onwards (and some from 1998 and 1999) can be downloaded from our website as follows:
http://econ.tau.ac.il/research/search_workingPapers.asp