Theory Dec. (2012) 73:501–512
DOI 10.1007/s11238-011-9264-5

History as a coordination device

Rossella Argenziano · Itzhak Gilboa

Published online: 18 June 2011

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 article 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 plays a game that exhibits strategic complementarities. We assume a dynamic process that faces different populations with such games for randomly selected values of a parameter. We introduce a belief formation process that takes into account the history of similar games played in the past, not necessarily by the same population. We show that when history serves as a coordination device, the limit behavior depends on the way history unfolds, and cannot be determined from a-priori considerations.

Keywords Belief formation · Similarity · Coordination games · Equilibrium selection

1 Introduction

Games with strategic complementarities typically exhibit multiple equilibria. The game theoretic literature has witnessed many attempts to select equilibria based on the

R. Argenziano (✉)
Department of Economics, University of Essex, Colchester, Essex, UK
e-mail: rargenz@essex.ac.uk

I. Gilboa
HEC, Paris, France
e-mail: tzachigilboa@gmail.com

I. Gilboa
Tel-Aviv University, Tel-Aviv, Israel

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. Other types of considerations attempted to embed the game in a dynamic process (Young 1993; Kandori et al. 1993; Burdzy et al. 2001) or in incomplete information set-up (Carlsson and van Damme 1993).

It is noteworthy that risk dominance has emerged as the preferred selection criterion based on all the above quite different types of considerations. On the other hand, the literature on strategic complementarities arising from 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 that the actual selection of an equilibrium is often left to history, chance, institutional details, or other unmodeled factors.

In this article, we are interested in a dynamic process, according to which large populations are called upon to play a simple coordination game. In the first stage, each player chooses either a low or a high action. In the second stage, nature chooses a low or a high outcome, and nature's move depends on the set of players choosing the high action. Consider the decision of a single player in this game. The optimal action to take depends on his assessment of the probability of a high outcome. We maintain that this assessment would and should be based on the results of past instances of similar games. These games may have been played by the same population or by others. Each past game might differ from the current one by one parameter at most, which is a proxy for the difference in the expected payoff of the two strategies. 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. However, ignoring these past games would hardly seem a rational way of generating beliefs.[1]

Consider, for example, the following "revolution" game played by a large population. In stage 1, each player $i \in [0, 1]$ chooses whether to participate in a revolutionary attempt, or to opt out. In stage 2, 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 who have chosen to participate. After each player determined her choice of participation and nature determined the success of the revolution, the game is over. The payoff of each player depends only on her own choice of participation, and on nature's move. The payoff function $u = u_i$ for every $i \in [0, 1]$ is given by the following matrix:

	S(uccess)	F(ailure)
Participate	1	0
Opt out	$\frac{x+1}{2}$	x

(1)

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

[1] The 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."

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. This worst payoff is normalized 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. This payoff is normalized to 1.

An individual who decides to participate in the revolution therefore decides to bet on its success with the extreme payoffs 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. 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. Their payoff would only equal the arithmetic average between the full benefit, 1, and the status quo, x.

Consider the decision problem of a potential rebel. 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? Suppose that it is common knowledge in the population that revolution games of the type above have been played in the past. We believe that the history of such similar games played in the past should affect the beliefs of the potential rebel.

In this article, we present a simple belief formation process for a class of games that includes the "revolution" example above. The process is such that the probability assigned to a high outcome in the current game is the weighted empirical frequency of high outcomes in past games, where the weights are given by a similarity function that takes into account the differences between past games and the current one. We find that beliefs that are history-dependent may lead to different behavior, depending on the way history unfolds.

The rest of this article is organized as follows. We first discuss related literature. Section 2 describes the stage game. We devote Sect.3 to modeling the way players generate beliefs given history. Section 4 describes the dynamic process and provides the main result of the article. Finally, Sect.5 concludes.

1.1 Related literature

Our article is closely related to the equilibrium selection literature discussed in the introduction. The belief formation process we study can be viewed as a generalization of the process of "fictitious play", first analyzed by Robinson (1951), where players respond optimally to the beliefs defined by the empirical frequency of past play in the game. Research on fictitious play has focused on identifying classes of games with the Fictitious Play Property (FPP), i.e., games for which every fictitious play process converges in belief to equilibrium. Krishna (1992) proves that the FPP holds for games with strategic complementarities, diminishing returns, and a specific tie-breaking rule, Monderer and Shapley (1996) prove it for finite potential games, and Berger (2008) proves it for nondegenerate games with ordinal complementarities and diminishing returns. While these articles prove that every fictitious play process converges to equilibrium in games closely related to the ones we study, they are silent on the issue of equilibrium selection. Similarly, the process of best response dynamics introduced by Gilboa and Matsui (1991) does not imply selection of a specific equilibrium in the long run for games with strategic complementarities.[2] In contrast to this literature, in our set-up different past games can have varying degrees of relevance, and thus empirical frequency is generalized to weighted empirical frequency.

Conceptually, the belief formation process that we assume is closely related to the process studied by LiCalzi (1995), which looks at the case where players give the same similarity weight to the outcome of all the games in a given class, and to the solution concept introduced by Jehiel (2005), in which players form beliefs about their opponents' behavior by grouping nodes in which the opponents play into analogy classes. Our approach assumes a similarity function that is continuous and that does not define a similarity equivalence relation. In particular, for every level of degree of similarity, the relation "as similar as..." in our model is not transitive. Also, in our model the notion of similarity is outcome-dependent, so a past case that resulted in failure may be more or less relevant to a given game than the same past case had it resulted in success.

Steiner and Stewart (2008) study similarity-based learning in games. This article, which is probably the closest to ours, shows that contagion can lead to a unique long-run outcome. Importantly, their model differs from ours in that they consider a continuous state space, in which there is non-negligible probability of the zones in which each pure strategy is dominant.

Our article also relates to the literature on coordination games of regime change. 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 (2008) studied information manipulation in games of regime change, whereas Angeletos et al. (2007) focus on a learning process by which individuals playing such games form beliefs. As in Lohmann (1994) and Angeletos et al. (2007), we study the evolution of beliefs in a game that is played repeatedly. However, as opposed to these articles, 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 See also Matsui and Matsuyama (1995) and Hofbauer and Sigmund (1998).

2 The stage game

We describe a symmetric two-stage extensive form game G_x depending on a parameter $x \in X \equiv \{x^1, \ldots, x^J\}$ for $J > 2$. The cardinal values of the parameter x will be of no import, but their order will. There is a continuum of players [0, 1]. In stage 1, all players move simultaneously. The set of moves for each player i is $S_i = \{0, 1\}$. In stage 2, after each player determined her move in {0, 1}, nature chooses an outcome $\nu \in \{0, 1\}$. 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 $\nu = 1$ with probability $\varphi(\lambda(A))$ where φ is strictly increasing, with $\varphi(0) = 0$ and $\varphi(1) = 1$, and λ stands for Lebesgue's measure. If A is non-measurable, the probability of nature choosing $\nu = 1$ can be defined arbitrarily. At equilibrium, the set A will be measurable. Observe that using Lebesgue's measure here does not involve any significant loss of generality: since the function φ is only assumed to be increasing, one may start with any measure that is absolutely continuous with respect to λ and replace it by λ by an appropriate choice of φ.

After each player determined her choice and nature determined the outcome (by the probability $\varphi(\lambda(A))$), the game is over. The payoff of each player depends only on her own choice and on nature's move. However, ex-ante, the game exhibits strategic complementarities: The change in the expected payoff when own action changes from 0 to 1 is strictly increasing in the measure of players playing 1.

Assume, then, that an individual i attempts to estimate the expected utility of playing 1 versus 0 for a given game G_x with $x \leq x^J$. Suppose that individual i's belief over the measure of other individuals who choose 1 is given by a measure $\mu_{i,x}$ over (the Lebesgue σ-algebra on) [0, 1]. Therefore, for every Lebesgue-measurable set $B \subset [0, 1]$, individual i assigns probability $\mu_{i,x}(B)$ to the event that the measure of individuals who eventually choose 1 (with or without herself) lies in B. Specifically, the subjective probability of individual i that nature will choose 1 in the game x is

$$\widehat{p}_{i,x} = \int_{[0,1]} \varphi(p) \mathrm{d}\mu_{i,x}(p).$$

We assume that for every $x \in X$ there exists a unique $\bar{p}_x \in [0, 1]$ such that playing 1 is optimal if and only if player i believes that nature will choose 1 with probability larger or equal to $\bar{p}_x$. We denote by $\bar{p}$ the vector of thresholds $(\bar{p}_x)_{x \in X}$.

Given beliefs $\mu_{i,x}$, player i's expected payoff from playing 1 in game G_x is greater (smaller) than her expected payoff from playing 0 iff $\widehat{p}_{i,x} > \bar{p}_x$ ($\widehat{p}_{i,x} < \bar{p}_x$). For simplicity we assume that in case of a tie, $\widehat{p}_{i,x} = \bar{p}_x$, player i will play 0.[3]

We assume that $\bar{p}_x$ is strictly increasing in $x \in X$. Therefore, the games are assumed to be ordered according to the difference in the expected payoff of the two strategies. We further assume that $\bar{p}_{x^1} = 0$ and $\bar{p}_{x^J} = 1$. Therefore, in the game G_{x^1}, strategy 1 is dominant, whereas in G_{x^J} – strategy 0 is.

[3] While this assumption will prove immaterial, it simplifies analysis because a random tie-breaking rule requires some additional assumption about the law of large numbers applying to a continuum of i.i.d random variables. See Judd (1985).

At equilibrium, all players will have the same beliefs, hence $\widehat{p}_{i,x} = \widehat{p}_x$, i.e., it is independent of i. Therefore, at equilibrium all players will either play 1 or 0. This implies that a player who has beliefs $\mu_{i,x}$ and who is aware of the entire process, can follow the same reasoning we do and conclude that the probability of nature choosing 1 is, in fact, either 0 or 1, rather than $\widehat{p}_{i,x}$. To accommodate these players, define $\widehat{p}_{i,x}$ as the player's naive beliefs, and the beliefs that result from our analysis—as the player's sophisticated beliefs. Due to strategic complementarities, an act that is optimal with respect to the naive beliefs will also be optimal with respect to the sophisticated beliefs.

It is important to note that if naive beliefs were to be ignored, and players were to have only sophisticated beliefs, then any assignment of 0 or 1 to the games $G_{x^1}, \ldots, G_{x^J}$ could be a consistent set of equilibrium beliefs. However, such a model would not describe the process by which beliefs are formed. The naive belief formation process is the topic of the next section.[4]

3 Belief formation process

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], and for different values of x. The history of similar games played in the past, which is assumed to be common knowledge, determines the beliefs $\widehat{p}_x$ 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 $x_t \in X \equiv \{x^1, \ldots, x^J\}$ in an i.i.d. manner, according to a known discrete distribution. Thus the process is determined by a probability vector $(p_1, \ldots, p_J)$.

Let $H_t = \left((x_\tau, \nu_\tau)_{\tau=1}^{t-1}\right)$ be the history at the beginning of period t, where, for $\tau < t, x_\tau \in X$ denotes the game played at period τ, and $\nu_\tau \in \{0, 1\}$ denotes its outcome. In each period t, all the players are assumed to observe the same history H_t. Before playing, they observe the game G_{x_t} and form an expectation on the probability of a success that is based on the similarity between the current game and previous games that ended, respectively, with a success or a failure.

Let there be two matrices of non-negative numbers $s^+, s^- : X \times X \to \mathbb{R}_+$ with the following interpretation. $s^+(x_\tau, x_t)$ measures the degree of support that a past game x_τ, resulting in $\nu_\tau = 1$, gives to the outcome 1 at the new game x_t. Similarly, $s^-(x_\tau, x_t)$ measures the degree of support that a past game x_τ, resulting in $\nu_\tau = 0$, gives to the outcome 0 at the new game x_t. These degrees of support generate naive beliefs as follows. Denote $S_t = \{\tau < t \mid \nu_\tau = 1\}$ and $F_t = \{\tau < t \mid \nu_\tau = 0\}$, and set

$$\widehat{p}(H_t, x_t) = \frac{\sum_{\tau \in S_t} s^+(x_\tau, x_t)}{\sum_{\tau \in S_t} s^+(x_\tau, x_t) + \sum_{\tau \in F_t} s^-(x_\tau, x_t)} \tag{2}$$

[4] For a discussion of modeling the formation of rational beliefs, see Gilboa et al. (2011).

Note that if the functions s^+, s^- are identically 1, the expression above is simply the relative frequency of 1's in the history H_t. The formula (2) allows different past games to have different weight in the evaluation of probabilities at the current period. Thus, it can be viewed as a generalization of empirical frequencies to weighted empirical frequencies.[5]

Specifically, we assume that games with a lower index x are commonly perceived as a-priori more likely to result in 1 than are games with a higher index x'. (This is in line with the assumption that the difference in the expected payoff of strategies 1 and 0 is strictly decreasing in x.) Thus, a result of 1 in G_x is less surprising than the same result in a game $G_{x'}$. Hence, a result of 1 in G_x lends weaker support to the same result in the current game than would the result 1 in a game $G_{x'}$. Formally, assume that $s^+(x, y)$ is strictly increasing in its first argument and strictly decreasing in its second argument. Similarly, we also assume that $s^-(x, y)$ is strictly decreasing in its first argument and strictly increasing in its second argument. An implication of these assumptions is that $\widehat{p}(H_t, x_t)$ is strictly decreasing in its second argument.

The formula (2) is not well-defined for the first period, $t = 1$. Also, it allows $\widehat{p}(H_t, x)$ to be 0 or 1, if history contains only 0-outcomes or one 1-outcomes, respectively. We find such extreme beliefs unwarranted. Hence, we use Eq. 2 only when history contains both outcomes or one 1-outcomes. Formally, we assume that $t \geq 3$, and that history contains at least one 0-outcome and at least one 1-outcome, so that $\widehat{p}(H_t, x) \in (0, 1)$.

Our model assumes that all players share both history and the similarity functions. The first assumption is roughly equivalent to a complete information assumption: in this model, a case is in the database of one player if and only if it is in the database of all; thus, there is no private information, and this fact may be (implicitly) assumed to be common knowledge among the players. The second assumption, namely, that all players employ the same similarity functions, is a rough analog of the common prior assumption: the model assumes that judgments, which are in principle subjective, happen to coincide. Both assumptions are made for simplicity: our main result is that there exist probability distributions over the state space such that long-run equilibria are not unique. Tautologically, such an existence result still holds in a more general model. Moreover, one can easily see that the result would hold for a range of similarity functions and ranges of empirical frequencies that correspond to different memories.

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 $X = \{x^1, \ldots, x^J\}$ according to probabilities $(p_1, \ldots, p_J)$, beliefs are formed in accordance with Eq. 2, and the players' behavior in G_{x_t} is chosen by the beliefs $\widehat{p}(\cdot)$.

[5] The idea of generating beliefs in a game based on past empirical frequencies is at the heart of "fictitious play", dating back to Robinson (1951). Extending empirical frequencies to similarity-weighted empirical frequencies was suggested and axiomatized in Billot et al. (2005) and Gilboa et al. (2006). Here, we extend the notion of similarity-weighted empirical frequencies to incorporate directional thinking.

A state of the process is fully summarized by a matrix of relative frequencies

$R_t =$

	$x = x^1$	$x = x^2$	...	$x = x^J$
1	$r_{t,11}$	$r_{t,12}$	...	$r_{t,1J}$
0	$r_{t,01}$	$r_{t,02}$	...	$r_{t,0J}$

where $r_{t,ij}$ is the relative frequency, up to time t, of periods in which the game was G_{x^j} and the outcome was i.

Consider the following matrices:

$R^0 =$

	$x = x^1$	$x = x^2$	...	$x = x^J$
1	p_1	0	...	0
0	0	p_2	...	p_J

$R^1 =$

	$x = x^1$	$x = x^2$	...	$x = x^J$
1	p_1	p_2	...	0
0	0	0	...	p_J

We can finally present our main result.

Theorem 1 *For given $\bar{p}$ and s^+, s^-, there exist distributions $(p_1, \ldots, p_J)$ such that there is a positive probability that R_t converges to R^0 and a positive probability that it converges to R^1.*

The main intuition behind this result is that, by choosing a particular draw of games, one can generate histories in which in practically all games everyone expects a particular equilibrium to be played. For example, if we start by a large number of draws of $x = x^J$, when another game (apart from x^1) is played, history contains almost no revolutions, and it suggests that, should a revolution be attempted, it would fail. This results in the population choosing to opt out, and, in turn, in reinforcing the no-revolution outcome in all games but $x = x^1$. Clearly, the history so far described generates relative frequencies similar to the matrix R^0. One needs to verify that such a finite history leads to convergence to R^0 with positive probability. Similarly, one can choose an alternative sequence of games that would converge to R^1.

It will be obvious from the proof of the theorem that there is nothing exceptional about the distributions $(p_1, \ldots, p_J)$ that allow convergence to either of the extreme outcomes. The main condition will be that the probabilities of the extreme games, p_1, p_J be strictly positive but small, relative to the other probabilities (and given X, the similarity functions, and the function $\varphi(\cdot)$). In particular, the set of distributions $(p_1, \ldots, p_J)$ contains open sets.

Observe, however, that if x can take any value in the continuum $[0, 1]$, the probabilities of the extreme values should be zero for our result to hold. Thus, one may prove a similar theorem in case x is continuous, provided that the dominance regions (of the two pure strategies) are only $\{0\}$ and $\{1\}$. This is in contrast to Steiner and Stewart (2008), who employ a continuous model with positive probability strict-dominance regions.

5 Conclusion

Ever since the early days of game theory, there has been a quest for a solution concept that would satisfy existence and uniqueness, with robustness and dynamic stability as additional desiderata. The attempt to narrow down the class of potential predictions was motivated by the desire to make game theory more meaningful and powerful, whether interpreted descriptively or normatively. Clearly, even if uniqueness of equilibria cannot be obtained, tighter theoretical predictions would make the theory more useful, and will thereby reduce the need to resort to extra-theoretical reasoning to select an equilibrium as a likely or a recommended outcome.

The literature on refinements of Nash equilibrium (see van Damme 1983) is generally perceived as falling short of pinpointing unique equilibria in games. However, the more recent literature, viewing a game in the context of similar and related games, have resulted in several results that changed the way we think about equilibrium selection (Carlsson and van Damme 1993; Burdzy et al. 2001; Weinstein and Yildiz 2007). These results may suggest that, in a sufficiently detailed model, a unique equilibrium prediction would exist.

This article is offered as an example, showing that incorporation of additional details into the model may leave the game theoretic prediction ambiguous. We believe that game theoretic analysis is extremely useful, but that, in general, it cannot subsume the need in historical and institutional knowledge. Rather, the formal, mathematical analysis needs to be combined with such knowledge to generate trustworthy predictions.

Acknowledgments We are grateful to Dirk Bergemann, Don Brown, Dino Gerardi, Stephen Morris, Barry O'Neill, Ady Pauzner, Alessandro Pavan, Colin Stewart and Nicolas Vieille for comments and references. Gilboa gratefully acknowledges ISF Grants 975/03 and 396/10 and ERC Grant 269754. This article is part of the Polarization and Conflict Project CIT-2-CT-2004-506084 funded by the European Commission-DG Research Sixth Framework Programme.

Appendix: Proof of theorem 1

First, observe that the relative frequencies of the columns of R_t are governed only by the selection of x, and are independent of the players' behavior.

Under our assumptions, for every history H_t we can predict the outcome of the game played at time t by considering the difference $\widehat{p}(H_t, x_t) - \bar{p}_{x_t}$ (decreasing in x_t). If this difference is strictly positive, all players' expectation $\widehat{p}(\cdot)$ will be above the critical belief $\bar{p}_{x_t}$. They will therefore all play 1, and Nature will select $\nu_t = 1$ with probability 1. Otherwise, all players will play 0 and Nature will select $\nu_t = 0$ with probability 1.

Recall that

$$\widehat{p}\left(H_t, x^k\right) = \frac{\sum_{\tau \in S_t} s^+\left(x_\tau, x^k\right)}{\sum_{\tau \in S_t} s^+\left(x_\tau, x^k\right) + \sum_{\tau \in F_t} s^-\left(x_\tau, x^k\right)}.$$

The assumption that $t \geq 3$, and that history contains at least one 0-outcome and one 1-outcome implies that $\widehat{p}(H_t, x_t) \in (0, 1)$. This in turn implies that the difference $\widehat{p}(H_t, x_t) - \bar{p}_{x_t}$ is strictly positive at $x_t = x^1$ and strictly negative at $x_t = x^J$.

We simplify notation by defining

$$A_{kt} = \sum_{\tau \in S_t} s^+ \left(x_\tau, x^k\right) \tag{3}$$

$$B_{kt} = \sum_{\tau \in F_t} s^- \left(x_\tau, x^k\right) \tag{4}$$

$$\widehat{p}\left(H_t, x^k\right) = \frac{A_{kt}}{A_{kt} + B_{kt}} \tag{5}$$

$$z_{kt} = \left(1 - \bar{p}_{x^k}\right) A_{kt} - \bar{p}_{x^k} B_{kt} \tag{6}$$

so that

$$\widehat{p}(H_t, x^k) - \bar{p}_{x^k} > 0 \Leftrightarrow z_{kt} > 0.$$

Given that A_{kt} is strictly decreasing in x^k, B_{kt} is strictly increasing in x^k, and $\bar{p}_x$ is strictly increasing in x^k, it follows that z_{kt} is strictly decreasing in x^k, hence for history H_t there exists a unique $y_t \in \{x^1, \ldots x^{J-1}\}$ for which $z_{y_t t} > 0 \geq z_{y' t}$ for any y' in X such that $y' > y_t$.

At time t, the expected change in z_{kt} is given by:

$$\begin{aligned}
&E\left[z_{k(t+1)} - z_{kt} | y_t = y\right] \\
&\quad = E\left[\left(1 - \bar{p}_{x^k}\right) A_{k(t+1)} - \bar{p}_{x^k} B_{k(t+1)} | y_t = y\right] - \left[\left(1 - \bar{p}_{x^k}\right) A_{kt} - \bar{p}_{x^k} B_{kt}\right] \\
&\quad = \left(1 - \bar{p}_{x^k}\right) E\left[A_{k(t+1)} - A_{kt} | y_t = y\right] - \bar{p}_{x^k} E[B_{k(t+1)} - B_{kt} | y_t = y] \\
&\quad = \left(1 - \bar{p}_{x^k}\right) \sum_{x^j \leq y} p_j s^+ \left(x^j, x^k\right) - \bar{p}_{x^k} \sum_{x^j > y} p_j s^- \left(x^j, x^k\right).
\end{aligned}$$

This is increasing in y, as $s^+(\cdot)$ and $s^-(\cdot)$ are nonnegative. Also, it is decreasing in x^k because $\bar{p}_x$ and $s^-(\cdot)$ are increasing in x^k and $s^+(\cdot)$ is decreasing in x^k.

Consider histories H_t that contain at least one 1-outcome and a sufficiently long list of 0-outcomes such that $z_{2t} \leq 0$. Such histories have positive probability as long as $p_1, p_J > 0$. Since z_{kt} is strictly decreasing in x^k, $z_{2t} \leq 0$ implies $z_{kt} \leq 0$ for $k \in \{2, \ldots, J\}$.

The expected change in z_{2t} is given by:

$$E\left[z_{2(t+1)} - z_{2t} | z_{2t} \leq 0\right] = \left(1 - \bar{p}_{x^2}\right) p_1 s^+ \left(x^1, x^2\right) - \bar{p}_{x^2} \sum_{j \geq 2} p_j s^- \left(x^j, x^2\right).$$

Let $(p_1, \ldots, p_J)$ be such that $E\left[z_{2(t+1)} - z_{2t} | z_{2t} \leq 0\right] < 0$. (Therefore, assume that $p_1 > 0$ is small enough relative to the other p_k's.)

Since $E\left[z_{k(t+1)} - z_{kt}|y_t = y\right]$ is decreasing in x^k, $E\left[z_{2(t+1)} - z_{2t}|z_{2t} \leq 0\right] < 0$ implies $E\left[z_{k(t+1)} - z_{2t}|z_{kt} \leq 0\right] < 0$ for $k \in \{2, \ldots, J\}$. We argue that, given that $z_{2t} \leq 0$, there is a positive probability that $z_{2\tau} \leq 0$ for all $\tau > t$. To see this, observe that, as long as $z_{2\tau} \leq 0$ for $\tau > t$, $z_{2\tau}$ follows a Markov process. The distribution of $z_{2\tau}$ conditional on $z_{2\tau} > 0$ is not guaranteed to be stationary. However, if we replace it by any stationary distribution, we obtain a new process $\left\{\hat{z}_{2\tau}\right\}_{\tau>t}$ that is Markovian, with $E\left[\hat{z}_{k(t+1)} - \hat{z}_{2t}|\hat{z}_{2t} \leq 0\right] < 0$ and that is identical to $\{z_{2\tau}\}_{\tau>t}$ as long as the latter is non-positive. Since $\left\{\hat{z}_{2\tau}\right\}_{\tau>t}$ has a positive probability of never becoming positive, so does $\{z_{2\tau}\}_{\tau>t}$. This completes the proof that our process has a positive probability of converging to R^0.

A symmetric argument shows that there are probabilities $(p_1, \ldots, p_J)$ for which the process has a positive probability of converging to R^1. Moreover, following the arguments above it is clear that one can find such probabilities for which both events occur with positive probability: basically, one has to guarantee only that $p_1, p_J > 0$ are small enough relative to the other p_k's.

References

Angeletos, G.-M., Hellwig, C., & Pavan, A. (2007). Dynamic global games of regime change: Learning, multiplicity, and the timing of attacks. *Econometrica, 75*, 711–756.

Berger, U. (2008). Learning in games with strategic complementarities reviseted. *Journal of Economic Theory, 143*, 292–301.

Billot, A., Gilboa, I., Samet, D., & Schmeidler, D. (2005). Probabilities as similarity-weighted frequencies. *Econometrica, 73*, 1125–1136.

Burdzy, K., Frankel, D. M., & Pauzner, A. (2001). Fast equilibrium selection by rational players living in a changing world. *Econometrica, 69*, 163–189.

Carlsson, H., & van Damme, E. (1993). Global games and equilibrium selection. *Econometrica, 61*, 989–1018.

Edmond, C. (2008). Information manipulation, coordination and regime change, mimeo

Gilboa, I., & Matsui, A. (1991). Social stability and equilibrium. *Econometrica, 59*, 859–867.

Gilboa, I., Lieberman, O., & Schmeidler, D. (2006). Empirical similarity. *Review of Economics and Statistics, 88*, 433–444.

Gilboa, I., Postlewaite, A., & Schmeidler, D. (2011). Rationality of belief, or: Why Savage's axioms are neither necessary nor sufficient for rationality. *Synthese*, forthcoming.

Hofbauer, J., & Sigmund, K. (1998). *Evolutionary games and replicator dynamics*. Cambridge: Cambridge University Press.

Hume, D. (1748). *An enquiry concerning human understanding*. Oxford: Clarendon Press.

Jehiel, P. (2005). Analogy-based expectation equilibrium. *Journal of Economic Theory, 123*, 81–104.

Judd, K. (1985). The law of large numbers with a continuum of random variables. *Journal of Economic Theory, 35*, 19–25.

Kandori, M., Mailath, G. J., & Rob, R. (1993). Learning, mutation, and long run equilibria in games. *Econometrica, 61*, 29–56.

Katz, M., & Shapiro, C. (1986). Technology adoption in the presence of network externalities. *Journal of Political Economy, 94*, 822–841.

Krishna, V. (1992). Learning in games with strategic complementaties. HBS Working Paper 92-073, Harvard University.

LiCalzi, M. (1995). Fictitious play by cases. *Games and Economic Behavior, 11*, 64–89.

Lohmann, S. (1994). The dynamics of informational cascades: The monday demonstrations in Leipzig, East Germany, 1989–1991. *World Politics, 47*, 42–101.

Matsui, A., & Matsuyama, K. (1995). An approach to equilibrium selection. *Journal of Economic Theory, 65*, 415–434.

Monderer, D., & Shapley, L. S. (1996). Fictitious play property for games with identical interests. *Journal of Economic Theory, 68*, 258–265.

Muller, E. N., & Opp, K.-D. (1986). Rational choice and rebellious collective action. *The American Political Science Review, 80*, 471–488.

Robinson, J. (1951). An iterative method of solving a game. *The Annals of Mathematics, 54*, 296–301.

Schelling, T. (1960). *The strategy of conflict*. Cambridge, MA: Harvard University Press.

Steiner, J., & Stewart, C. (2008). Contagion through Learning. *Theoretical Economics, 3*, 431–458.

van Damme, E. (1983). *Refinements of the Nash equilibrium concept. Lecture Notes in Economics and Mathematical Systems,* vol. 219. Berlin: Springer Verlag.

Weinstein, J., & Yildiz, M. (2007). A structure theorem for rationalizability with application to robust predictions of refinements. *Econometrica, 75*, 365–400.

Young, H. P. (1993). The evolution of conventions. *Econometrica, 61*, 57–84.