

A Note on Equilibrium Selection in Polya–Urn Coordination Games

Giorgio Fagiolo

Department of Economics, University of Verona, Verona (Italy) and Laboratory of Economics and Management, Sant'Anna School of Advanced Studies, Pisa (Italy).

Abstract

We study equilibrium selection in coordination games played by a population of agents whose size increases over time. We assume that, in each time period, a new player enters the economy, observes current strategy shares and irreversibly chooses a strategy on the basis of expected payoffs. We employ a simple Polya–Urn scheme to discuss the efficiency of long–run equilibria under alternative individual decision rules (e.g. best–reply, logit, etc.). We show that the system delivers a predictable outcome only when agents employ either a linear or a logit probability rule. If agents employ deterministic best–reply rules, Pareto–efficient coordination can occur, but the actual outcome depends on initial conditions and chance. In all other cases, coexistence of strategies characterizes equilibrium configurations. Finally, we discuss our findings in the framework of technological adoption models.

Citation: Fagiolo, Giorgio, (2005) "A Note on Equilibrium Selection in Polya–Urn Coordination Games." *Economics Bulletin*, Vol. 3, No. 45 pp. 1–14

Submitted: July 20, 2005. **Accepted:** October 6, 2005.

URL: <http://www.economicbulletin.com/2005/volume3/EB-05C70028A.pdf>

1. Introduction

In the last years, evolutionary-game theory has provided robust solutions to the problem of selection among multiple Nash equilibria in symmetric coordination games (Samuelson, 1997). More specifically, a series of papers has suggested that, whenever the game is played by a population of myopic agents who are repeatedly matched over time and employ best-reply rules to update their strategies, the risk-dominant (RD) Nash equilibrium will be almost always preferred to the Pareto-efficient (PE) one in the long-run¹.

These models have been often employed to address the issue of technological adoption in the presence of several competing standards and network externality effects (cf. e.g., Kirman, 1997). However, in such formalizations, agents typically belong to a fixed-size population and are always able to revise the strategy they currently play. This might seem to be at odds with many “real-world” adoption choices, where agents sequentially enter the market and irreversibly choose one of the existing standards (Katz and Shapiro, 1986; Nelson, Peterhansl, and Sampat, 2004).

On the contrary, irreversible adoption with sequential entry has been studied in Polya-Urn models (Arthur et. al, 1987; Dosi and Kaniovski, 1994). Yet, these models did not explicitly address the issue of selection among equilibria (e.g., adoption shares) that can be ranked according to risk-efficiency criteria.

To fill this gap, this note begins to explore equilibrium selection in symmetric 2×2 coordination games played by myopic agents where: (i) the cardinality of the population increases over time due to a positive net entry rate; (ii) players irreversibly choose a strategy when they enter the economy.

We employ a standard Polya-Urn framework to describe the dynamics of strategy shares. More precisely, we assume that one agent enters the economy in any time period $t = 1, 2, \dots$ and must irreversibly choose among the available strategies with probabilities which depend on expected payoffs. We study the long-run behavior of the system under alternative individual decision rules (e.g. best-replies with or without stochastic mistakes, linear or logit probabilistic rules) and we characterize its efficiency properties.

The rest of this note is organized as follows. Section 2 describes the model. In Section 3 we investigate the long-run behavior of the system and we discuss equilibrium selection issues. Section 4 concludes.

2. A Simple Model

Consider an economy evolving in discrete time-steps $t = 0, 1, 2, \dots$. At time $t = 0$, there are $N_0 > 0$ agents in the economy. Each agent is fully described by a binary pure strategy (*strategy*, henceforth) $s \in S = \{-1, +1\}$. Thus, at any t , the state of the system can be characterized by the current share $x_t \in [0, 1]$ of agents playing $+1$.

¹Cf. e.g. Kandori et al. (1993) and Young (1996). This conclusion holds true also if agents always play the game against their nearest-neighbors (Ellison, 1993) and, to some extent, if they are able to endogenously choose their opponents in the game (cf. Jackson and Watts, 2002; Goyal and Vega-Redondo, 2003; Fagiolo, 2005).

The dynamics runs as follows. Given some initial share x_0 , at any $t > 0$ *one* new agent enters the economy and irreversibly chooses a strategy in S . This implies that the size of the population at time t is $N_t = N_0 + t$. The entrant then ‘plays’ a symmetric 2×2 coordination game against all incumbent agents according to a standard stage-game payoff matrix as in Table , left panel.

	+1	-1		+1	-1
+1	a	b		+1	1 0
-1	c	d		-1	α β

Table 1: Stage-Game Payoff Matrix. Left: Original Payoffs. Right: Normalized Payoffs.

In the left panel of Table , $a > c$ and $b < d$ because the game is a coordination one. We also assume that $a > d$ (i.e. the Nash equilibrium $(+1, +1)$ is Pareto efficient by construction) and that $a \neq b$. In what follows, however, we shall focus on the 2-parameter stage-game payoff matrix in the right panel of Table , obtained from the former – without losing in generality – by letting $\alpha = (c - b)/(a - b)$ and $\beta = (d - b)/(a - b)$. Given the above restrictions, $\alpha < 1$ and $0 < \beta \leq 1$. For the sake of convenience, we also assume that $\alpha \geq 0$. Notice that if $\beta = 1$ and $\alpha = 0$ the game is a pure-coordination one. Moreover, if $\alpha + \beta = 1$ the two Nash equilibria $E^+ = (+1, +1)$ and $E^- = (-1, -1)$ are risk-equivalent. Thus, if $\alpha + \beta < 1$, E^+ is both PE and RD, while, if $\alpha + \beta > 1$, E^+ is PE but E^- is RD.

Since agents are myopic, expected payoffs associated to any given choice $s \in S$ are given by:

$$\pi_t(s) = \pi(s, x_t; \alpha, \beta) = \begin{cases} x_t & \text{if } s = +1 \\ (\alpha - \beta)x_t + \beta & \text{if } s = -1 \end{cases} . \quad (1)$$

Let us call $\Pr\{s^*; \pi_t^{+1}, \pi_t^{-1}\}$ the probability that the agent chooses $s^* \in S$. We assume that entrants choose $s^* \in S$ according to one of the individual decision rules in Table ².

The dynamics can be easily formalized in terms of a Polya-Urn framework (Arthur et al., 1984; Dosi et al., 1994). Since, under any decision rule, $\Pr\{s^*; \pi_t^{+1}, \pi_t^{-1}\}$ only depends – through payoffs – on x_t , α and β , we can define:

$$p(x_t; \alpha, \beta) = \Pr\{s^*; \pi_t^{+1}, \pi_t^{-1}\}. \quad (2)$$

Thus, the dynamics of x_t reads:

$$x_{t+1} = x_t + \frac{p(x_t; \alpha, \beta) - x_t}{N_0 + t + 1} + \frac{\eta(x_t; \alpha, \beta)}{N_0 + t + 1}, \quad (3)$$

where η_t is a zero-mean, dichotomous random variable such that:

$$\Pr\{\eta(x_t; \alpha, \beta) = \tilde{\eta}\} = \begin{cases} p(x_t; \alpha, \beta) & \text{if } \tilde{\eta} = 1 - p(x_t; \alpha, \beta) \\ 1 - p(x_t; \alpha, \beta) & \text{if } \tilde{\eta} = -p(x_t; \alpha, \beta) \end{cases} . \quad (4)$$

²Best-reply rules are studied in Blume (1995). Linear-probability and logit rules are discussed in Anderson, dePalme, and Thisse (1992), Blume (1993), Weisbuch, Kirman, and Herreiner (2000) and Brock and Durlauf (2001).

Acronym	Decision Rule	$\Pr\{s^*; \pi_t^{+1}, \pi_t^{-1}\}$
DBR	Deterministic Best-Reply	$\begin{cases} 1 & \text{if } \pi_t(s^*) > \pi_t(-s^*) \\ 1/2 & \text{if } \pi_t(s^*) = \pi_t(-s^*) \\ 0 & \text{if } \pi_t(s^*) < \pi_t(-s^*) \end{cases}$
PBR	Perturbed Best-Reply	$\begin{cases} 1 - \epsilon & \text{if } \pi_t(s^*) > \pi_t(-s^*) \\ 1/2 & \text{if } \pi_t(s^*) = \pi_t(-s^*) \\ \epsilon & \text{if } \pi_t(s^*) < \pi_t(-s^*) \end{cases}$
LP	Linear Probability	$\frac{\pi_t(s^*)}{\pi_t(+1) + \pi_t(-1)}$
NP	Non-Linear (logit) Probability	$\frac{\exp\{\frac{\gamma}{2} \cdot \pi_t(s^*)\}}{\exp\{\frac{\gamma}{2} \cdot \pi_t(+1)\} + \exp\{\frac{\gamma}{2} \cdot \pi_t(-1)\}}$

Table 2: Individual Decision Rules. PBR: $\epsilon > 0$. NP: $\gamma > 0$. *Note:* (i) PBR converges to a DBR as $\epsilon \downarrow 0$; (ii) NP converges to a DBR as $\gamma \uparrow \infty$.

As t increases, the impact of stochastic entry choices on x_t becomes weaker and weaker. Dosi et al. (1994) show that, under some mild restrictions, x_t converges with probability 1 to a *stable* point belonging to the zeroes of the function $g(x) = p(x) - x$, where $x \in [0, 1]^3$. The properties of p determine the conceivable limit points of the system. However, the dynamics is path-dependent: if there are at least two stable points, which one is actually selected by the process depends on both initial conditions (x_0) and on the stochastic sequence of early entrants' decisions.

3. Long-Run Behavior and Equilibrium Selection

In this Section, we characterize the limit properties of the system under the four decision rules in Table . Let us start with a DBR rule. In this case, it is easy to see that:

Proposition 1 *If agents choose under a DBR rule, then:*

$$p(x_t; \alpha, \beta) = \begin{cases} 1 & \text{if } x_t > x^* \\ \frac{1}{2} & \text{if } x_t = x^* \\ 0 & \text{if } x_t < x^* \end{cases},$$

where $x^* = \beta / (1 - \alpha + \beta)$. Thus: $\Pr\{\lim_{t \rightarrow \infty} x_t \in \{0, 1\}\} = 1$.

Proof. See Appendix A. ■

Thus, full coordination will always emerge in the limit, but we cannot predict if it will occur on E^+ or E^- . Notice, however, that the larger x^* , the smaller the set $\{x \in [0, 1] :$

³See Appendix A. In general, a root $\hat{x} \in (0, 1)$ of $g(x)$ is *stable* if $g(x)$ crosses the line $y = x$ from above at \hat{x} . A similar definition applies for $\hat{x} \in \{0, 1\}$.

$p(x) > x\}$. Since x^* increases with $\beta/(1 - \alpha)$, the more E^- is risk-dominant (i.e. the more $\beta > 1 - \alpha$), the larger is the set of shares for which, on average, x_{t+1} decreases to 0. If the two equilibria are risk-equivalent ($\alpha + \beta = 1$), then $x^* = \frac{1}{2}$. So, if one supposes to draw x_0 from a random variable uniformly distributed over $[0, 1]$, the value of x^* can give us a measure of the likelihood of the RD equilibrium in the deterministic dynamics:

$$E(x_{t+1}|x_t) = x_t + \frac{p(x_t; \alpha, \beta) - x_t}{N_0 + t + 1}. \quad (5)$$

When agents instead employ a PBR, the probability with which the currently best-strategy is chosen becomes $1 - \epsilon$, while the currently inferior option is now selected with probability ϵ . We have that:

Proposition 2 *If agents choose under a PBR rule, then:*

$$p(x_t; \alpha, \beta) = \begin{cases} 1 - \epsilon & \text{if } x_t > x^* \\ \frac{1}{2} & \text{if } x_t = x^* \\ \epsilon & \text{if } x_t < x^* \end{cases},$$

where $x^* = \beta/(1 - \alpha + \beta)$. Thus $\Pr\{\lim_{t \rightarrow \infty} x_t \in \{\epsilon, 1 - \epsilon\}\} = 1$.

Proof. See Appendix A. ■

Coexistence of both +1 and -1 players will always emerge in equilibrium. If the system tends to select the PE (respectively, the RD) equilibrium, then an ϵ -sized ‘niche’ of agents playing the RD, inefficient equilibrium (respectively, the PE one) will always survive. As happens with the DBR rule, the long-run coordination outcome is entirely unpredictable. However, the larger x^* (i.e. the stronger risk-dominance of E^-), the larger the probability that – if one draws initial conditions at random – the deterministic dynamics in (5) will select the RD outcome.

The parameter ϵ is typically interpreted as the probability that an agent makes a ‘mistake’ – with respect to the optimal available decision. This mistake occurs independently of how large (relative) payoffs are. More generally, one might suppose that agents choose only *in probability* the strategy associated to the highest payoffs (LP or NP). If agents employ a LPR, one can prove:

Proposition 3 *If agents choose under a LP rule, then:*

$$p(x_t; \alpha, \beta) = \frac{x_t}{(1 + \alpha - \beta)x_t + \beta},$$

where $\alpha \in [0, 1)$ and $\beta \in (0, 1]$. Hence:

1. If $\alpha = 0$ and $\beta = 1$, then $\forall x \in [0, 1]$, $\Pr\{\lim_{t \rightarrow \infty} x_t = x\} > 0$.
2. If $\alpha \neq 0$, let $x^{**} = \frac{1 - \beta}{1 - \beta + \alpha}$. Then: $\Pr\{\lim_{t \rightarrow \infty} x_t = x^{**}\} = 1$.

Proof. See Appendix A. ■

Under a LP rule, and if the game is not a pure-coordination one, the long-run behavior of the system becomes predictable: the share of agents playing +1 in the limit will converge a.s. to x^{**} . Notice that $x^{**} > 0$ unless $\beta = 1$: full coordination on the RD outcome will never be the case unless E^+ and E^- are Pareto-equivalent. Full coordination on the PE outcome will conversely emerge only if $\alpha = 0$ (that is, if the ‘loss’ from miscoordination is the same for both strategies). But what happens in terms of PE vs. RD equilibria? It is easy to see that $x^{**} > \frac{1}{2} \Leftrightarrow \alpha + \beta < 1$. If the two equilibria are risk-equivalent then $x^{**} = \frac{1}{2}$. Conversely, the larger $\alpha + \beta > 1$, the smaller the fraction of agents playing +1 in equilibrium. However, there will always exist in the limit a small ‘niche’ of players choosing the superior strategy. Similarly, if E^+ is both PE and RD and $\alpha + \beta < 1$ approaches 0, then the system will converge to a $x^{**} > \frac{1}{2}$ which grows with $\alpha + \beta$. Unless $\alpha = 0$, there will always be a persistent niche of -1 players in equilibrium.

Suppose now that entrants choose their strategy according to a non-linear, logit probability rule. Then:

Proposition 4 *If agents choose under a NP rule, then:*

$$p(x_t; \alpha, \beta, \gamma) = \frac{1}{1 + e^{-\gamma\beta[\frac{1-\alpha}{\beta}x_t - 1 + x_t]}} = \frac{1}{1 + e^{-\theta[\sigma x_t - 1 + x_t]}} = \tilde{p}(x_t; \theta, \sigma), \quad (6)$$

where $\theta = \gamma\beta > 0$ and $\sigma = \frac{1-\alpha}{\beta} > 0$. Therefore:

1. If $\alpha + \beta = 1 \Leftrightarrow \sigma = 1$ then:

- If $\theta \leq 2$ then $\Pr\{\lim_{t \rightarrow \infty} x_t = \frac{1}{2}\} = 1$.
- If $\theta > 2$ then $\exists \{\underline{x}(\theta), \bar{x}(\theta)\} : \Pr\{\lim_{t \rightarrow \infty} x_t \in \{\underline{x}(\theta), \bar{x}(\theta)\}\} = 1$, where: $\underline{x}(\theta) < \frac{1}{2}$, $\bar{x}(\theta) > \frac{1}{2}$, and, as $\theta \uparrow \infty$, $\underline{x}(\theta) \downarrow 0$ and $\bar{x}(\theta) \uparrow 1$.

2. If $\alpha + \beta \neq 1 \Leftrightarrow \sigma \neq 1$ then numerical analyses show that:

- If $\theta \leq 2$ then $\exists \hat{x}(\theta, \sigma) : \Pr\{\lim_{t \rightarrow \infty} x_t = \hat{x}(\theta, \sigma)\} = 1$. If $\sigma < 1 \Leftrightarrow \alpha + \beta > 1$ then $\hat{x}(\theta, \sigma) < \frac{1}{2}$ and decreases with θ . If $\sigma > 1 \Leftrightarrow \alpha + \beta < 1$ then $\hat{x}(\theta, \sigma) > \frac{1}{2}$ and $\hat{x}(\theta, \sigma)$ and increases with θ . Moreover, $\hat{x}(\theta, \sigma) \uparrow 1$ as $\sigma \uparrow \infty$ for any $\theta \leq 2$.
- If $\theta > 2$ then (3) admits both 1 stable root and 2 stable roots. However, if $\sigma > 1 \Leftrightarrow \alpha + \beta < 1$ and θ is sufficiently large, then 2 stable roots $\{\underline{\underline{x}}(\theta, \sigma), \bar{\bar{x}}(\theta, \sigma)\}$ always emerge. As before, we have that $\underline{\underline{x}}(\theta, \sigma) < \frac{1}{2}$, $\bar{\bar{x}}(\theta, \sigma) > \frac{1}{2}$ and, as $\theta \uparrow \infty$, $\underline{\underline{x}}(\theta, \sigma) \downarrow 0$, $\bar{\bar{x}}(\theta, \sigma) \uparrow 1$.

Proof. See Appendix A. ■

Under a NP rule, the long-run behavior of the system is predictable if $\theta \leq 2$. This happens either when the payoff of $(-1, -1)$ is small (β small), or when entrants’ choices are not very sensitive to current payoffs (γ small). Thus, risk-dominance of E^+ is not sufficient to guarantee predictability. However, when $\theta \leq 2$ coexistence always occurs and system

efficiency – as measured by $\hat{x}(\theta, \sigma)$ – increases the more E^+ is RD. On the contrary, if $\theta > 2$ and (E^+, E^-) are risk-equivalent, two stable roots always arise and they tend to diverge toward 0 and 1 as θ increases. This regime also occurs when agents become very sensible to current payoffs and E^+ is RD. In such a case, coexistence on efficient ($\bar{x}(\theta, \sigma) > \frac{1}{2}$) or inefficient ($\underline{x}(\theta, \sigma) < \frac{1}{2}$) mixes can always appear.

4. Concluding Remarks

In this note, we studied equilibrium selection in a symmetric 2×2 coordination game played by myopic agents in a population whose size increases over time. We have shown that the system delivers a predictable outcome only if agents employ either a linear probability rule or they use a logit rule, but the latter is not very sensible to current payoffs.

Full coordination on the PE outcome can only occur if agents employ: (i) deterministic best-reply rules (but the actual outcome depends here on initial conditions and chance); or (ii) linear probability rules and the ‘loss’ from miscoordination is the same for both strategies. In all other cases, coexistence of strategies will emerge in the long run and the efficiency level of the equilibrium configuration will greatly vary with initial conditions, early entrants’ choices and stage-game payoffs.

References

- Anderson, S., A. dePalme and J.-F. Thisse (1992), *Discrete Choice Theory of Product Differentiation*. Cambridge, MIT Press.
- Arthur, W., Y. Ermoliev and Y. Kaniovski (1984), “Strong laws for a class of path-dependent stochastic processes with applications”, in Argin, S. and P. Shirayev (eds.), *Proceedings Conference on Stochastic Optimization. Lecture Notes in Control and Information Sciences*. Berlin, Springer Verlag.
- Arthur, W., Y. Ermoliev and Y. Kaniovski (1987), “Path-dependent processes and the emergence of macro-structure”, *European Journal of Operational Research*, 30(1): 294–303.
- Blume, L. (1993), “The statistical mechanics of strategic interaction”, *Games and Economic Behaviour*, 5: 387–424.
- Blume, L. (1995), “The statistical mechanics of best-response strategy revision”, *Games and Economic Behavior*, 11: 111–145.
- Brock, W. and S. Durlauf (2001), “Discrete Choice with Social Interactions”, *Review of Economic Studies*, 68: 235–260.
- Dosi, G., Y. Ermoliev and M. Kaniovski, Yu. (1994), “Generalized urn schemes and technological dynamics”, *Journal of Mathematical Economics*, 23: 1–19.
- Dosi, G. and Y. Kaniovski (1994), “On ‘badly behaved’ dynamics”, *Journal of Evolutionary Economics*, 4: 93–123.

- Ellison, G. (1993), “Learning, local interaction and coordination”, *Econometrica*, 61: 1047–1071.
- Fagiolo, G. (2005), “Endogenous Neighborhood Formation in a Local Coordination Model with Negative Network Externalities”, *Journal of Economic Dynamics and Control*, 29: 297–319.
- Goyal, S. and F. Vega-Redondo (2003), “Network Formation and Social Coordination”, Working Paper No.481, Department of Economics, Queen Mary, University of London.
- Jackson, M. and A. Watts (2002), “On the formation of interaction networks in social coordination games”, *Games and Economic Behavior*, 41: 265–291.
- Kandori, M., G. Mailath and R. Rob (1993), “Learning, mutation and long run equilibria in games”, *Econometrica*, 61: 29–56.
- Katz, M. and C. Shapiro (1986), “Technology Adoption in the Presence of Network Externalities”, *Journal of Political Economy*, 94: 822–841.
- Kirman, A. (1997), “The Economy as an Evolving Network”, *Journal of Evolutionary Economics*, 7: 339–353.
- Nelson, R., A. Peterhansl and B. Sampat (2004), “Why and how innovations get adopted: a tale of four models”, *Industrial and Corporate Change*, 13: 679–699.
- Samuelson, L. (1997), *Evolutionary Games and Equilibrium Selection*. Cambridge, The MIT Press.
- Weisbuch, G., A. Kirman and D. Herreiner (2000), “Market organisation and trading relationships”, *Economic Journal*, 110: 411–436.
- Young, H. (1996), “The Economics of Convention”, *Journal of Economic Perspectives*, 10: 105–122.

Appendix: Proofs

To prove Propositions 1-4 above, we will make use of the following theorems from Dosi, Ermoliev, and Kaniovski (1994). Let x_t the proportion of +1 players at time $t = 0, 1, 2, \dots$ and define the set of zeroes of the function $g(x) = p(x) - x$ as:

$$B = \{x \in [0, 1] : [\underline{\psi}(x), \bar{\psi}(x)] \ni 0\},$$

where $\underline{\psi}(x) = \liminf_{y \rightarrow x} [p(y) - y]$, $\bar{\psi}(x) = \limsup_{y \rightarrow x} [p(y) - y]$ and y belongs to $R(0, 1)$, i.e. the set of rational numbers in $(0, 1)$. Then:

Theorem 1 (Arthur et al., 1984) . *The sequence $\{x_t\}$ converges a.s. to the set B as $t \rightarrow \infty$.*

An isolated point $\xi \in B$ is called *stable* if for every $(\varepsilon_1, \varepsilon_2) \in R_{++}^2$ small enough then $g(x)(x - \xi) < \Delta(\varepsilon_1, \varepsilon_2) < 0$, provided that $\varepsilon_1 \leq |x - \xi| \leq \varepsilon_2$ and x belongs to $R(0, 1)$. Conversely, an isolated point $\xi \in B$ is called *unstable* if for every $\varepsilon > 0$ small enough, then $g(x)(x - \xi) > 0$ for any $x \in R(0, 1) \cap [(\xi - \varepsilon, \xi) \cup (\xi, \xi + \varepsilon)]$. Thus:

Theorem 2 (Dosi et al., 1994) .

1. *If $\xi \in B$ is a stable point in $(0, 1)$ and there exists $(\varepsilon_1, \varepsilon_2) \in R_{++}^2$ such that $p(x) > 0$ for $x \in R(0, 1) \cap (\xi - \varepsilon_1, \xi)$ and $p(x) < 1$ for $x \in R(0, 1) \cap (\xi, \xi + \varepsilon_2)$, then $\Pr\{\lim_{t \rightarrow \infty} x_t = \xi\} > 0$ for every $x_0 \in (\xi - \varepsilon_1, \xi + \varepsilon_2)$.*
2. *If $\xi \in B$ is an unstable point in $(0, 1)$ and one of the following conditions hold: (a) in a neighborhood of ξ the function p is continuous; (b) there exists $\varepsilon > 0$ s.t. for $x \in R(0, 1) \cap (\xi - \varepsilon, \xi + \varepsilon)$ then $g(x)(x - \xi) \geq \lambda|x - \xi|^{1+\eta_1}$ and $p(x)(1 - p(x)) \geq \eta_2 > 0$, where $\eta_1 > \eta_3(2\lambda + 1)^{-1}$, $\eta_1 \in [0, 1]$, $\lambda > 0$ and $\eta_3 = \sup_{x \in R(0, 1) \cap (\xi - \varepsilon, \xi + \varepsilon)} p(x)[1 - p(x)]$; (c) there exists a neighborhood of ξ where for $x \leq \xi$ (resp. $x \geq \xi$) either (a) or (b) holds true and for $x > \xi$ (resp. $x < \xi$) $p(x) = 1$ (resp. $p(x) = 0$). Then: $\Pr\{\lim_{t \rightarrow \infty} x_t = \xi\} = 0$ for every x_0 .*
3. *Let $y_t^0 = N_o x_o / N_t$ and $y_t^1 = 1 + y_t^0 - N_o / N_t$. If $p(y_t^0) < 1$ for $t \geq 0$ and $\sum_{t \geq 0} p(y_t^0) < \infty$ then $\Pr\{\lim_{t \rightarrow \infty} x_t = 0\} > 0$. Also if $p(y_t^1) > 0$ for $t \geq 0$ and $\sum_{t \geq 0} [1 - p(y_t^1)] < \infty$ then $\Pr\{\lim_{t \rightarrow \infty} x_t = 1\} > 0$.*

We are now well-equipped to prove Propositions 1-4 above.

Proof of Propositions 1-2

To obtain $p(x_t; \alpha, \beta)$ in both cases, it suffices to note that $\pi_t(+1) \geq \pi_t(-1)$ if and only $x_t > x^* = \beta / (1 - \alpha + \beta)$. Define for $\epsilon \geq 0$ a generic BR rule as:

$$p(x_t; \alpha, \beta) = \begin{cases} 1 - \epsilon & \text{if } x_t > x^* \\ \frac{1}{2} & \text{if } x_t = x^* \\ \epsilon & \text{if } x_t < x^* \end{cases} .$$

Of course, if $\epsilon = 0$ we have a DBR, while a PBR is recovered for $\epsilon > 0$. It is easy to see that $B = \{\epsilon, x^*, 1 - \epsilon\}$. By applying Theorem 2, points (2-3), above one has that the set of stable points is $B^*(\epsilon) = \{\epsilon, 1 - \epsilon\}$. Therefore, $\Pr\{\lim_{t \rightarrow \infty} x_t \in B^*(\epsilon)\} = 1$. This means that: (i) under a DBR, $\Pr\{\lim_{t \rightarrow \infty} x_t \in \{0, 1\}\} = 1$; and (ii) under a PBR, $\Pr\{\lim_{t \rightarrow \infty} x_t \in \{\epsilon, 1 - \epsilon\}\} = 1$.

Figures 1 and 2 depict two examples of $p(x_t; \alpha, \beta)$ under DBR and PBR rules.

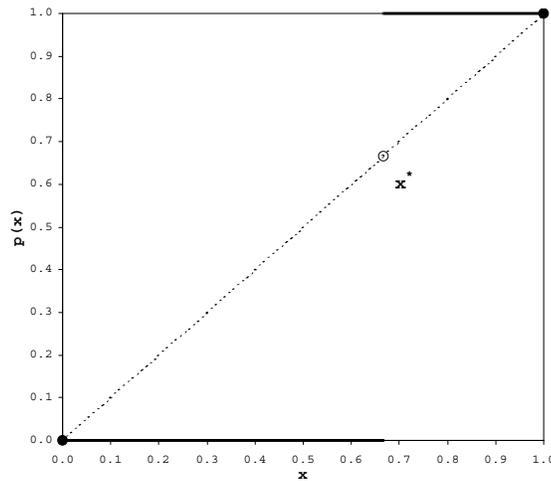


Figure 1: An example of $p(x_t; \alpha, \beta)$ with DBR Rules ($\alpha = 0.3, \beta = 0.5, x^* = 2/3$).

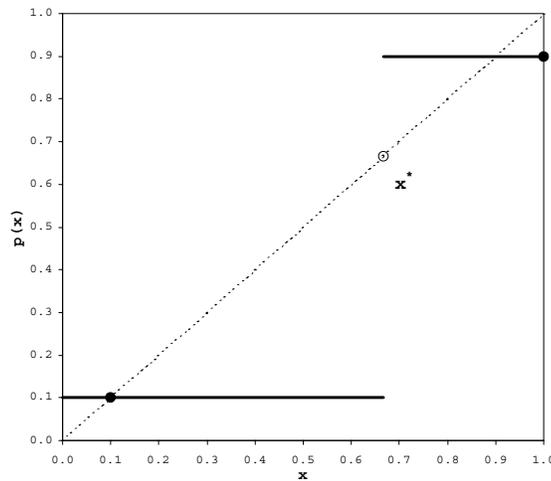


Figure 2: An example of $p(x_t; \alpha, \beta)$ with PBR rules ($\epsilon = 0.1, \alpha = 0.3, \beta = 0.5, x^* = 2/3$).

Proof of Proposition 3

By replacing expected payoffs (1) in:

$$p(x_t; \alpha, \beta) = \frac{\pi_t(+1)}{\pi_t(+1) + \pi_t(-1)},$$

one gets

$$p(x_t; \alpha, \beta) = \frac{x_t}{(1 + \alpha - \beta)x_t + \beta}.$$

Consider first the case $\alpha = 0$ and $\beta = 1$. Here, $p(x_t; \alpha, \beta) = x_t$. Therefore, $B = [0, 1]$ and $\Pr\{\lim_{t \rightarrow \infty} x_t = x\} > 0$ for all $x \in B$. If $(\alpha, \beta) \neq (0, 1)$, it is easy to see that p is C^0 over $[0, 1]$ and $\text{sign}(\partial p / \partial x_t) = \text{sign}(\beta) > 0 \forall x_t \in [0, 1]$. Since $p(0; \alpha, \beta) = 0$ and $p(1; \alpha, \beta) = (1 + \alpha)^{-1} \leq 1$, then $0 \leq p(x_t; \alpha, \beta) \leq 1$ for all $\alpha \in [0, 1)$ and $\beta \in (0, 1]$. Moreover, if the game is not a pure-coordination one ($\alpha = 0, \beta = 1$), p is always (strictly) concave:

$$\partial^2 p / \partial x_t^2 = -\frac{2\beta(1 + \alpha - \beta)}{[(1 + \alpha - \beta)x_t + \beta]^3} < 0 \quad \forall x_t \in [0, 1],$$

and $(\alpha, \beta) \neq (0, 1)$. Finally, notice that solving for the roots of $g(x_t; \alpha, \beta) = p(x_t; \alpha, \beta) - x_t$, one gets $B = \{0, x^{**}\}$, where

$$x^{**}(\alpha, \beta) = \frac{1 - \beta}{1 - \beta + \alpha},$$

and $0 \leq x^{**}(\alpha, \beta) \leq 1$. We have that $x^{**}(\alpha, \beta) = 0 \Leftrightarrow \beta = 1$ and $x^{**}(\alpha, \beta) = 1 \Leftrightarrow \alpha = 0$. By applying Theorem 2, points (2-3), above one has that only $x^{**}(\alpha, \beta)$ is stable, as $p(\varepsilon; \alpha, \beta) > \varepsilon$ for any small $\varepsilon > 0$. Thus $\Pr\{\lim_{t \rightarrow \infty} x_t = x^{**}(\alpha, \beta)\} = 1$.

Figure 3 shows some examples of $p(x_t; \alpha, \beta)$ under a LP rule for different values of α and β .

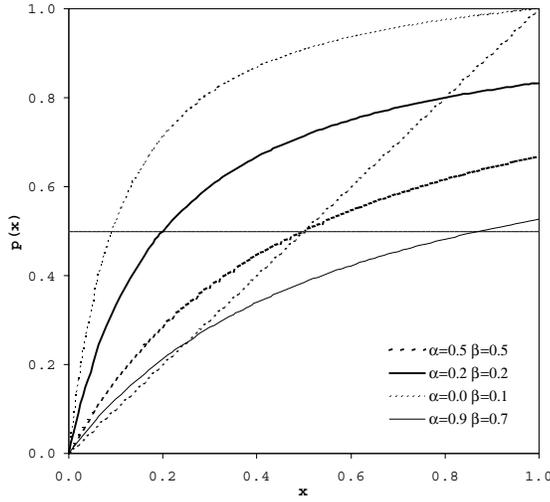


Figure 3: An example of $p(x_t; \alpha, \beta)$ with LP rules for different values of α and β .

Proof of Proposition 4

By replacing expected payoffs (1) in:

$$p(x_t; \alpha, \beta) = \frac{\exp\{\frac{\gamma}{2} \cdot \pi_t(+1)\}}{\exp\{\frac{\gamma}{2} \cdot \pi_t(+1)\} + \exp\{\frac{\gamma}{2} \cdot \pi_t(-1)\}},$$

one gets

$$p(x_t; \alpha, \beta) = \frac{1}{1 + e^{-\gamma\beta[\frac{1-\alpha}{\beta}x_t - 1 + x_t]}}$$

and, defining $\theta = \gamma\beta > 0$ and $\sigma = \frac{1-\alpha}{\beta} > 0$, the more useful two-parameter function:

$$\tilde{p}(x_t; \theta, \sigma) = \frac{1}{1 + e^{-\theta[\sigma x_t - 1 + x_t]}}$$

Being a logistic function, we know that $\tilde{p} \in C^0$ and $\partial\tilde{p}/\partial x_t > 0 \forall x_t \in [0, 1]$. Moreover, $\tilde{p}(0; \theta, \sigma) = (1 + e^\theta)^{-1} < \frac{1}{2}$ and $\tilde{p}(1; \theta, \sigma) = (1 + e^{-\theta})^{-1} > \frac{1}{2}$. Thus, $0 < \tilde{p}(x_t; \theta, \sigma) < 1 \forall x_t \in [0, 1]$ and for all $(\theta, \sigma) \in R_{++}^2$. Notice also that \tilde{p} is convex for all $x_t \leq (1 + \sigma)^{-1} = \beta/(1 - \alpha + \beta) = x^*$ and concave for $x_t \geq x^*$. Hence, at $x_t = x^*$, $\partial^2\tilde{p}/\partial x_t^2 = 0$ for all $(\theta, \sigma) \in R_{++}^2$.

Consider first the case $\alpha + \beta = 1 \Leftrightarrow \sigma = 1$. Here, $\tilde{p}(x^*; \theta, \sigma) = x^* = \frac{1}{2}$, so that $B \ni \frac{1}{2}$. However, $\frac{1}{2}$ does not need to be the unique root for $\tilde{g}(x_t; \theta, 1) = \tilde{p}(x_t; \theta, 1) - x_t = 0$. In fact, it is easy to see that $\frac{1}{2}$ is the unique root for $\tilde{g}(x_t; \theta, 1) = 0$ if and only if $\partial\tilde{p}/\partial x_t|_{x_t=\frac{1}{2}} \leq 1 \Leftrightarrow \theta \leq 2$. In this case, $\frac{1}{2}$ is stable because \tilde{p} is C^0 and crosses $\frac{1}{2}$ from above. Conversely, if $\theta > 2$, $\tilde{g}(x_t; \theta, \sigma) = 0$ will admit two more roots (other than $\frac{1}{2}$), say $\{\underline{x}(\theta), \bar{x}(\theta)\}$, where it must be that $\underline{x}(\theta) < \frac{1}{2}$ and $\bar{x}(\theta) > \frac{1}{2}$. Thus, $\frac{1}{2}$ becomes unstable, while $\underline{x}(\theta)$ and $\bar{x}(\theta)$ are both stable for the same argument. Some examples of $\tilde{p}(0; \theta, 1)$ for different values of θ are shown in Figure 4, while Figure 5 plots the roots of $\tilde{g}(x_t; \theta, 1) = 0$ against θ : notice how, as θ increases, $\underline{x}(\theta) \downarrow 0$ and $\bar{x}(\theta) \uparrow 1$. Hence, Theorem 2 applies and the first point of Proposition 4 is proved.

Consider now the case $\alpha + \beta \neq 1 \Leftrightarrow \sigma \neq 1$. Numerical analyses show the existence of two parameter regions, one where the dynamics of the system admits only one stable root, and another one where two stable roots exist, see Figure 6.

Notice that when $\theta \leq 2$, similarly to the $\sigma = 1$ case, the system admits only one root $\hat{x}(\theta, \sigma)$ for all $\sigma > 0$. Moreover, as Figure 7 shows, if $\sigma < 1$ then $\hat{x}(\theta, \sigma) < \frac{1}{2}$ and decreases as θ increases. Conversely, if $\sigma > 1$, then $\hat{x}(\theta, \sigma) > \frac{1}{2}$ and $\hat{x}(\theta, \sigma) \uparrow 1$ as $\theta \uparrow \infty$. Finally, $\hat{x}(\theta, \sigma) \uparrow 1$ as $\sigma \uparrow \infty$ for any $\theta \leq 2$, cf. Figure 8.

If $\theta > 2$, Figure 6 shows that the process admits both 1 stable root and 2 stable roots. However, if $\sigma > 1$ and θ is sufficiently large, then 2 stable roots $\{\underline{\underline{x}}(\theta, \sigma), \bar{\bar{x}}(\theta, \sigma)\}$ always emerge. As Figure 9 indicates, in this region $\underline{\underline{x}}(\theta, \sigma) < \frac{1}{2}$, $\bar{\bar{x}}(\theta, \sigma) > \frac{1}{2}$. Finally, as $\theta \uparrow \infty$, $\underline{\underline{x}}(\theta, \sigma) \downarrow 0$, $\bar{\bar{x}}(\theta, \sigma) \uparrow 1$. Once again, by applying Theorem 2, the second point of Proposition 4 follows.

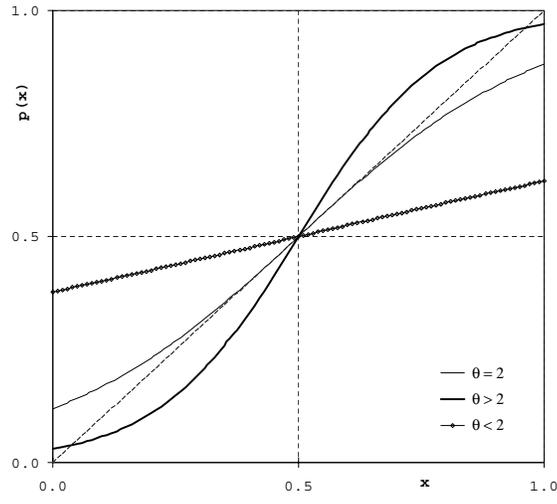


Figure 4: NP rules. The case $\sigma = 1$. The function $\tilde{p}(x_t; \theta, 1)$ for different values of θ .

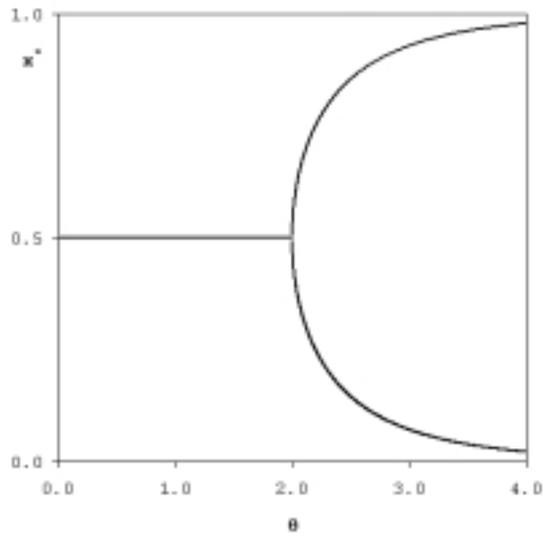


Figure 5: NP rules. The case $\sigma = 1$. Roots of $\tilde{p}(x_t; \theta, 1) - x_t = 0$ as a function of θ .

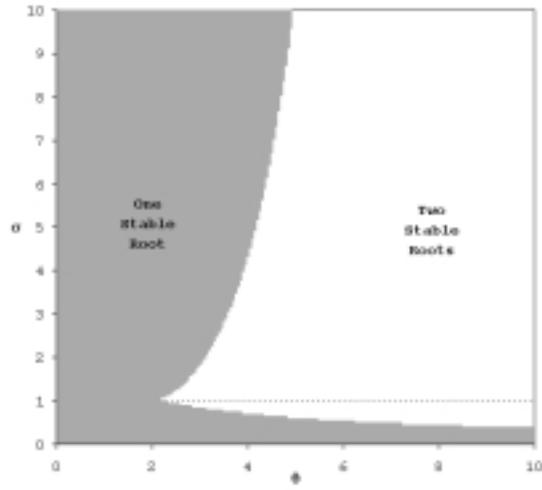


Figure 6: NP rules. The case $\sigma \neq 1$. Stable roots vs. σ and θ .

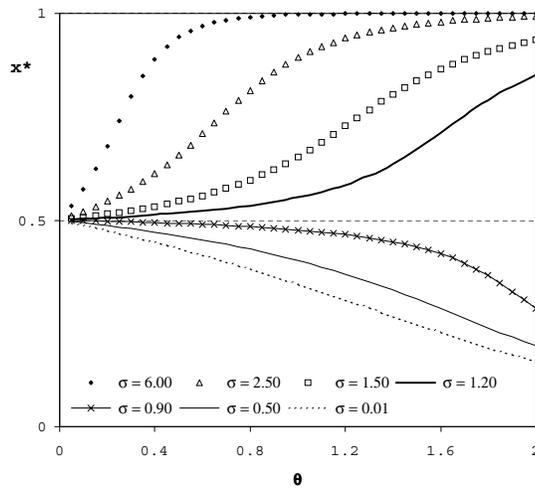


Figure 7: NP rules. The case $\sigma \neq 1$, $\theta \leq 2$. The single stable root as a function of θ for different values of σ .

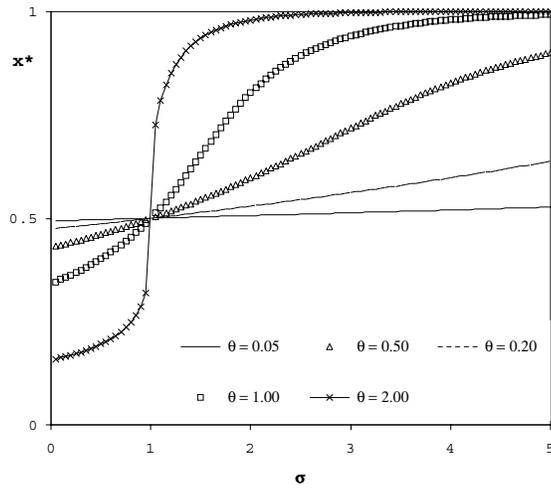


Figure 8: NP rules. The case $\sigma \neq 1$, $\theta \leq 2$. The single stable root as a function of σ for different values of θ .

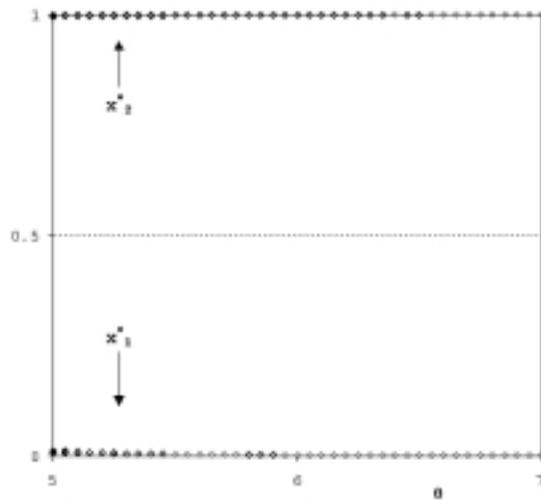


Figure 9: NP rules. The case $\sigma > 1$. The two stable roots vs. $\theta \geq 5$.