Economics Bulletin

Volume 45, Issue 4

Equilibrium selection in generalized Polya-urn coordination games with partial information on population shares

Giorgio Fagiolo
Institute of Economics, Sant'Anna School for Advanced Studies, Pisa, Italy

Abstract

I investigate deterministic best-reply dynamics in generalized Polya-urn coordination games, where a Pareto dominant equilibrium coexists with a risk-dominant one and players have partial information on current population shares. After analytically deriving estimates for the basin of attraction of either equilibrium, I study numerically (and via simulation) equilibrium selection in the long run. Results indicate that risk dominance robustly prevails over Pareto efficiency as a selection principle. Furthermore, I find that the probability of selecting a Pareto-efficient outcome decreases the smaller the sample size drawn by entrants and increases the larger the initial pool of incumbents.

This work was partially supported by the SMaRT COnSTRUCT project (CUP J53C24001460006), in the context of FAIR (PE0000013, CUP B53C22003630006) under the Italian National Recovery and Resilience Plan (Mission 4, Component 2, Line of Investment 1.3) funded by the European Union - NextGenerationEU.
Citation: Giorgio Fagiolo, (2025) "Equilibrium selection in generalized Polya-urn coordination games with partial information on population shares", *Economics Bulletin*, Volume 45, Issue 4, pages 1658-1665
Contact: Giorgio Fagiolo - giorgio.fagiolo@santannapisa.it.
Submitted: November 18, 2025. **Published:** December 30, 2025.

1 Introduction

Equilibrium selection in coordination games has a long-standing tradition in game theory and economics, encompassing both static frameworks (Kim, 1996) and evolutionary, dynamic approaches (Samuelson, 1998). This prominence reflects the wide applicability of such games across various domains (Schelling, 1978; Cooper et al., 1990). A notable area of focus has been the adoption of technological standards in the presence of network externalities (Katz and Shapiro, 1985; Nelson et al., 2004). In dynamic contexts, this problem is typically modeled using evolutionary games (Young, 1996), where a fixed population of agents repeatedly revises strategies in a coordination game, myopically best-responding to global (Kandori et al., 1993) or local (Ellison, 1993; Blume, 1995; Fagiolo, 2005a) adoption shares.

Alternatively, when agents irreversibly select their strategy upon entering a growing population, Polya-urn processes highlight the role of initial conditions, path dependence, and potential inefficiencies (Johnson and Kotz, 1977; Arthur, 1994; Dosi and Kaniovski, 1994; Dosi et al., 1994). In these models, entrants perfectly observe current population shares and respond to global adoption rates using either probabilistic (Arthur et al., 1986) or deterministic/noisy best-reply rules (Fagiolo, 2005b). While this assumption is plausible when adoption rates are freely available, it becomes problematic when information about global shares is costly and agents can only observe a sample of incumbents. In this case, the properties of the model are poorly understood.[1]

Although much of the seminal work on coordination and Polya-urn processes dates back to the 1990s and early 2000s, more recent contributions have been relatively scarce. Nonetheless, the mechanisms underlying equilibrium selection remain of enduring relevance, not least because coordination problems continue to arise in settings with evolving technologies, network externalities, and decentralized information (Agarwal et al., 2016; Raducha and San Miguel, 2022).

Against this background, the present paper extends the Polya-urn coordination framework to environments where agents face partial information. Specifically, I study Polya-urn coordination games where population grows over time due to entrants who cannot observe global adoption rates, but only a sample of incumbents. Upon entry, players sample (without replacement) a share of incumbents and deterministically best-reply to irreversibly choose their strategy. To keep the analysis consistent with existing research (Kandori et al., 1993; Ellison, 1993; Fagiolo, 2005b), I assume that agents face a coordination stage-game where a Pareto-efficient pure-strategy equilibrium coexists with a risk-dominant one (Harsanyi and Selten, 1988). This allows one to study how long-run equilibrium selection depends on sample sizes and risk dominance.

2 The Model

I consider a potentially infinite population of agents playing a simple coordination game with 2 pure strategies $s \in S = \{PE, RD\}$ and a stage-game payoff matrix[2]:

$$\begin{array}{c|cc} & PE & RD \\ \hline PE & 1 & 0 \\ RD & \alpha & \beta \end{array} \tag{1}$$

where $0 < \alpha < 1$ and $0 < \beta < 1$. The game admits two pure-strategy equilibria (PE, PE) and (RD, RD), where (PE, PE) is Pareto efficient. I also assume that (RD, RD) is risk-dominant, i.e. $\alpha + \beta > 1$.

[1] Cf. Chen and Wei (2005) and Chen et al. (2014) for results in the applied probability literature.

[2] This stage-game payoff can be recovered from a standard 4-parameter setup, where $\pi : S \times S \to \mathbb{R}$, $\pi(PE, PE) = a$, $\pi(PE, RD) = b$, $\pi(RD, PE) = c$ and $\pi(RD, RD) = d$, assuming that $a > c$, $c > b$, $d > b$, $a > d$, and $a > b$, subtracting b from all entries, and dividing by $a - b > 0$.

Therefore, the set A of admissible parameters in the game is defined as:

$$A = \{(\alpha, \beta) : \alpha < 1, 0 < \beta < 1, \alpha + \beta > 1\}, \tag{2}$$

implying that $\alpha + \beta < 2$.

Time is discrete. At $t = 0$, there are N_0 players in the game. We assume throughout that the initial strategy configuration is drawn independently at random tossing a fair coin. Agents are irreversibly associated to their chosen strategy and cannot revise their choices across time. In each subsequent time period $t = 1, 2, \ldots$, the dynamics runs as follows:

1. A new agent enters the game;
2. The new agent samples without replacement k_t agents from the pool of $N_t = N_0 + t$ incumbents, where $1 \leq k_t \leq N_t$;
3. The new agent chooses the pure strategy $s^* \in \{PE, RD\}$ that maximizes total payoffs from playing against players in the k_t-agent sample.

Sample sizes $K = \{k_t \leq N_0 + t, t \geq 0\}$ are exogenous and deterministic parameters that can possibly change over time (more on this below in Section 4). Let $1 \leq p_t \leq k_t$ be the number of agents in the sample playing PE at time t. Total payoffs read:

$$\pi(s|p_t; k_t, \alpha, \beta) = \begin{cases} p_t, & \text{if } s = PE \\ \alpha p_t + \beta(k_t - p_t), & \text{if } s = RD \end{cases} \tag{3}$$

where $s \in S$. If players deterministically best-reply to choose s^*, then:

$$s^* = \begin{cases} PE & \text{if } p_t > \theta k_t \\ RD & \text{if } p_t < \theta k_t \end{cases} \tag{4}$$

where $\theta = \beta/(1 - \alpha + \beta) \in (1/2, 1)$ if $(\alpha, \beta) \in A$.[3]. This follows from the assumption that RD is risk-dominant: agents play PE only if the fraction of PE players in the sample is larger than a threshold larger than $1/2$.

3 Analysis

The law of motion for the number P_t of agents playing the PE strategy reads:

$$P_{t+1} = P_t + Z_t, \tag{5}$$

where Z_t is a random variable that, conditionally on P_t, is Bernoulli distributed. The expected value $E[Z_t|P_t]$ can be easily computed noting that, as entrants sample without replacement k_t incumbents from a population of N_t units, P_t of which play PE, the number p_t of PE players in the sample is distributed as a hypergeometric distribution, i.e.:

$$\text{Prob}\{p_t = p|N_t, P_t; k_t\} = \frac{\binom{P_t}{p}\binom{N_t - P_t}{k_t - p}}{\binom{N_t}{k_t}}, \tag{6}$$

[3]Note that one can safely disregard here payoff ties. This is because we can either assume throughout that players randomize if a payoff tie occurs as in Fagiolo (2005b); or choose θ in such a way that the condition $p_t = \theta k_t$ never occurs if $k_t \in K$.

with $\max\{0, k_t - (N_t - P_t)\} \le p \le \min\{P_t, k_t\}$. Consequently:

$$E[Z_t|P_t] \equiv \eta(N_t, P_t; k_t, \theta) = 1 - F(\lfloor \theta k_t \rfloor; N_t, P_t, k_t) \equiv H(\lfloor \theta k_t \rfloor; N_t, P_t, k_t), \tag{7}$$

where $F(\cdot; N_t, P_t, k_t)$ is the cdf of the hypergeometric distribution with parameters (N_t, P_t, k_t). Eqs. (5–7) fully characterize the evolution of the process once initial values and parameters (θ, K) are given.

Let $X_t = P_t/N_t$ be the share of PE players in the current pool of incumbents and $\tilde{x}_t = p_t/k_t$ the correspondent share in the sample drawn by the entrant. Hence, (5) becomes:

$$X_{t+1} = \frac{N_t}{N_{t+1}} X_t + \frac{Z_t}{N_{t+1}} = X_t + \frac{f(\tilde{x}_t) - X_t}{N_{t+1}} + \lambda(\tilde{x}_t; X_t), \tag{8}$$

where:

$$f(\tilde{x}_t) = \begin{cases} 1 & \text{if } \tilde{x}_t > \theta \\ 0 & \text{if } \tilde{x}_t < \theta \end{cases} \tag{9}$$

and:

$$\lambda(\tilde{x}_t; X_t) = \begin{cases} 1 - f(\tilde{x}_t), & \text{prob} = \eta(N_t, P_t; k_t, \theta) \\ -f(\tilde{x}_t), & \text{prob} = 1 - \eta(N_t, P_t; k_t, \theta) \end{cases} \tag{10}$$

Since $E[\lambda(\tilde{x}_t; X_t)|X_t] = 0$, we get that:

$$E[X_{t+1}|X_t] = X_t + \frac{f(\tilde{x}_t) - X_t}{N_{t+1}}. \tag{11}$$

I now investigate convergence of the stochastic process defined in Eq. (8). If $k_t = N_t$, for all $t > 0$ —i.e., entrants can observe the current global share X_t— then in Eq. (6) one obtains that $\text{Prob}\{p_t = p|N_t, P_t\} = 0$ if $p \neq P_t$ and $\text{Prob}\{p_t = P_t|N_t, P_t\} = 1$. Hence, $\tilde{x}_t = X_t$ in Eqs. 5–11, with $\eta(N_t, P_t; k_t, \theta) = f(X_t)$, and the model boils down to a standard generalized Polya-urn coordination game with best-response dynamics (Fagiolo, 2005b). In this case, the sequence X_t converges a.s. to the set $B = \{0, 1\}$ as $t \to \infty$ (cf. Dosi et al., 1994, , Theorems 1-4). This is because, given Eq. 9, the only fixed points of the deterministic dynamics described by Eq. 11 lie at the extremes of the unit interval. Thus, full coordination will always emerge in the limit, but whether a Pareto efficient or risk-dominant equilibrium emerges is unpredictable. Of course, the larger θ, the larger the basin of attraction (henceforth, BoA) of the RD equilibrium. When $(\alpha + \beta) \downarrow 1$ then $\theta \downarrow 1/2$ (i.e., the two pure-strategy equilibria tend to be risk-equivalent), we expect the process to converge with equal probability to either full-coordination configurations.

If $1 \le k_t < N_t$, and $(X_t, \tilde{x}_t)$, Eq. (11) still describes a deterministic law of motion of the process and we can apply again results in Arthur et al. (1986); Dosi et al. (1994) to conclude that X_t converges a.s. to $\{0, 1\}$. In order to evaluate the average BoA of the full-coordination RD equilibrium, it is useful to take expected values of both sides of (8) conditional on X_t only. Since $E[f(\tilde{x}_t)|X_t] = \eta(N_t, P_t; k_t, \theta)$ and $E[\lambda(\tilde{x}_t; X_t)|X_t] = 0$, one gets:

$$\begin{aligned} E[X_{t+1}|X_t] =& X_t + \frac{E[f(\tilde{x}_t)|X_t] - X_t}{N_{t+1}} + E[\lambda(\tilde{x}_t; X_t)|X_t] = \\ &= X_t + \frac{1}{N_{t+1}} \left[\eta(N_t, P_t; k_t, \theta) - \frac{P_t}{N_t} \right] = \\ & X_t + \frac{1}{N_{t+1}} \left[H(\lfloor \theta k_t \rfloor; N_t, P_t, k_t) - \frac{P_t}{N_t} \right] = \\ &= X_t + \frac{W(P_t; N_0 + t, k_t, \theta)}{N_0 + t + 1}. \end{aligned} \tag{12}$$

Eq. 12 can be employed to evaluate the expected BoA of long-run equilibria of the process. Define

$P^*(N_0, t, k_t, \theta)$ as the value of P_t such that:

$$W(P^*(N_0, t, k_t, \theta); N_0 + t, k_t, \theta) = 0 \tag{13}$$

i.e. a zero of the function W wrt P_t. The larger P^*, the smaller the BoA of the PE equilibrium. Unless $k_t = N_t, \forall t \geq 0$, the BoA depends not only on parameters (K, θ), but also on initial population size (N_0) and time. Therefore, the quantity $Q^*(N_0, t, k_t, \theta)$, defined as:

$$Q^*(N_0, t, k_t, \theta) = 1 - \frac{P^*(N_0, t, k_t, \theta)}{N_0 + t} \tag{14}$$

provides the expected share of PE players at time t as parameters change.[4]

4 Design of the Experiments

I study convergence to full coordination equilibria of the stochastic process defined by Eqs. (5–7), when stage-game parameters (α, β) vary and assumptions as to K are made. I mostly do so numerically, exploring the behavior of $Q^*(N_0, t, k_t, \theta)$ as parameters change. Furthermore, I validate numerical experiments via Monte Carlo exercises, i.e., simulating Eqs. (5–7) to check whether Monte Carlo frequencies match with those retrieved by numerically computing Q^*.

To simplify the setup, let further transform stage-game parameters as follows:

$$\psi = \frac{\alpha + \beta - 1}{\beta} = 2 - \frac{1}{\theta} \in (0, 1), \tag{15}$$

with $\theta = \beta/(1 - \alpha + \beta) \in (1/2, 1)$. Here, the parameter ψ controls for how much the strategy RD is risk dominant —i.e., how much $\alpha + \beta > 1$— and goes to zero if PE and RD tend to be risk equivalent. Conversely, ψ approaches one as both α and β tend to one —i.e., maximum risk dominance.

As to the sequence of exogenous sample sizes K, I consider a very simple baseline scenario where k_t is constant over time and reads:

$$k_t = \lfloor \phi N_0 \rfloor, \tag{16}$$

where $\phi \in (0, 1]$ controls for the share of initial population size (N_0). This assumption can be justified if one assumes that, as population size $N_t = N_0 + t$ increases with t, entrants might face sampling costs preventing them to enlarge their observation sample.

To check for robustness, I also explore a second, alternative, scenario where sampling costs are negligible and sample sizes k_t are allowed linearly to increase with t in such a way that the ratio k_t/N_t remains constant over time:

$$k_t = \lfloor \phi N_t \rfloor, \tag{17}$$

and ϕ controls for the share of current incumbents (N_t) sampled by the entrants.[5]

5 Results

Suppose first that sample sizes are constant, i.e. K as in Eq. (16). Figure 1, top panels, reports expected shares of PE players in the long run, when $t \to \infty$ and $N_0 \in 100, 1000$. PE shares are obtained computing $Q^*(N_0, t, \lfloor \phi N_0 \rfloor, \theta)$ and setting t = 1M. Notice, first, that PE shares decrease, everything else being

[4] A fully closed-form solution for long-run dynamics in generalized Polya-urn settings is typically intractable. This explains why much of the literature relies on numerical or simulation-based approaches (Arthur et al., 1987; Dosi and Kaniovski, 1994; Fagiolo, 2005b). The analytical formulation of Q^* presented here thus serves as a tractable benchmark.

[5] In this setup, if $\phi = 1$, all entrants are able to observe the current global share X_t and the process follows a standard generalized Polya-urn coordination game with deterministic best-response dynamics (Fagiolo, 2005b).

constant, as ψ increases. This is expected, as a larger ψ implies a stronger risk-dominance of the RD equilibrium.

More interestingly, the smaller the sample drawn by entrants (i.e., the smaller ϕ), the smaller the probability that the system converges to the PE equilibrium. This happens because the RD strategy is more rewarding when entrants face mixed populations, where the PE strategy is not dominant. Since Polya-urn processes are strongly path dependent (Arthur et al., 1987; Arthur, 1994) and initial conditions are perfectly mixed, the RD strategy is more rewarding insofar the sample drawn by the entrant is sufficiently mixed. However, the probability of drawing mixed samples is governed by Eq. (6) and it is well known that (for a given population size and number of successes in the population), the probability of drawing the same sufficiently mixed sample (e.g., fifty-fifty) is larger the smaller sample size. Therefore, the smaller sample size, the larger the probability that first entrants choose RD, thus setting the stage to a path-dependent history converging in the limit to coordination on the RD equilibrium.

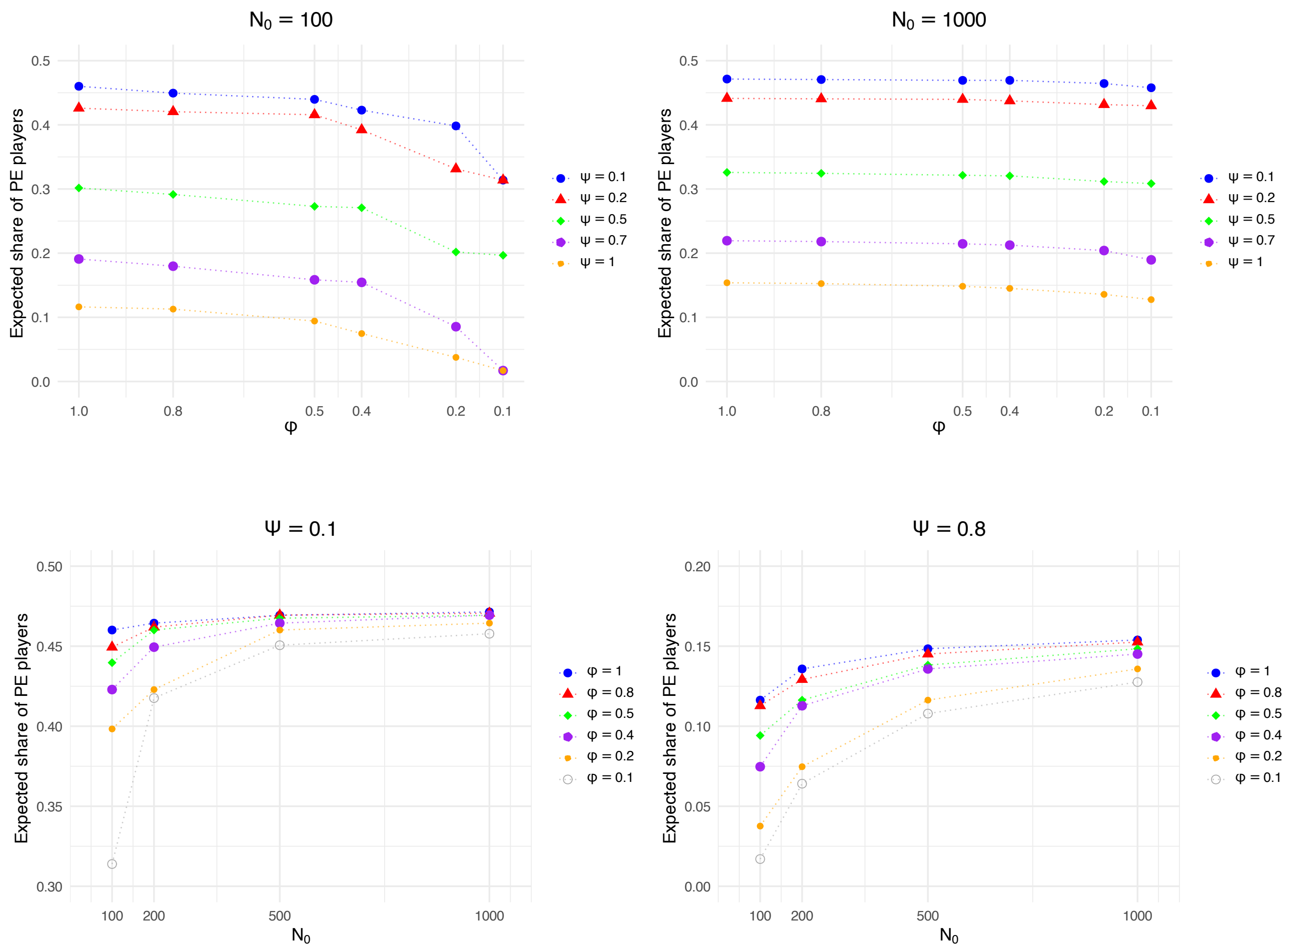

Figure 1: Expected share of PE players when $k_t = \lfloor \phi N_0 \rfloor$, $t \to \infty$. Top panels: $Q^*(N_0, t, k_t, \theta)$ as a function of (ϕ, ψ) with $N_0 \in 100, 1000$. Bottom panels: $Q^*(N_0, t, k_t, \theta)$ as a function of N_0 for different values $(\phi, \psi.)$

Figure 1, top-right panel, also suggests that expected PE shares depend on initial population size as, given the same (ϕ, ψ), the BoA of the RD equilibrium seems to shrink. This behavior is exemplified in the two bottom panels of Figure 1, where expected PE shares are plotted against N_0 for different values of sample sizes and two levels of ψ (small vs. large RD dominance). Results confirm that, everything else being constant, a larger initial population sizes implies larger expected PE shares. Using again the intuition above, this finding can be rationalized noting that (given the same sample ratio) a similar level of mixing in the sample is more probable —as per Eq. (6)— when the population size is smaller. As a result, in larger initial populations, the PE strategy is more rewarding than in smaller ones, leading to larger long-run expected PE shares.

Next, I explore what happens when entrants are allowed to draw a sample of incumbents that grows linearly with population size, as in Eq. 17. To begin with, recall that long-run shares are highly dependent on the choices made by first entrants, who are more uncertain, but whose decision has a larger impact on the future history of the process. Therefore, an increasing sample size should not have a detectable impact on long-run expected shares, as key decisions are made with smaller samples. Those entering later in the process, who can already count on almost-settled shares, enjoy a further boost of information, being allowed to sample a larger fraction of incumbents. This is indeed what Figure 2 shows, where for similar values of (ϕ, ψ) I plot expected shares for the case $N_0 = 100$.

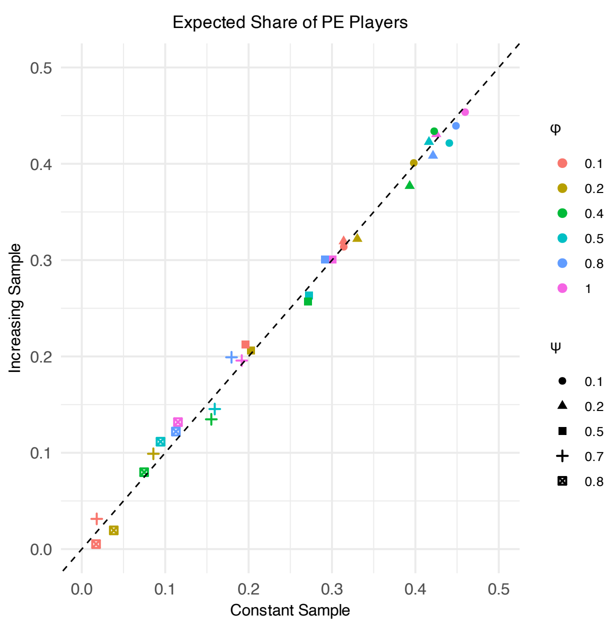

Figure 2: Constant ($k_t = \lfloor \phi N_0 \rfloor$) vs Increasing ($k_t = \lfloor \phi N_t \rfloor$) Sample Size. Expected shares $Q^*(N_0, t, k_t, \theta)$ with $t \to \infty$.

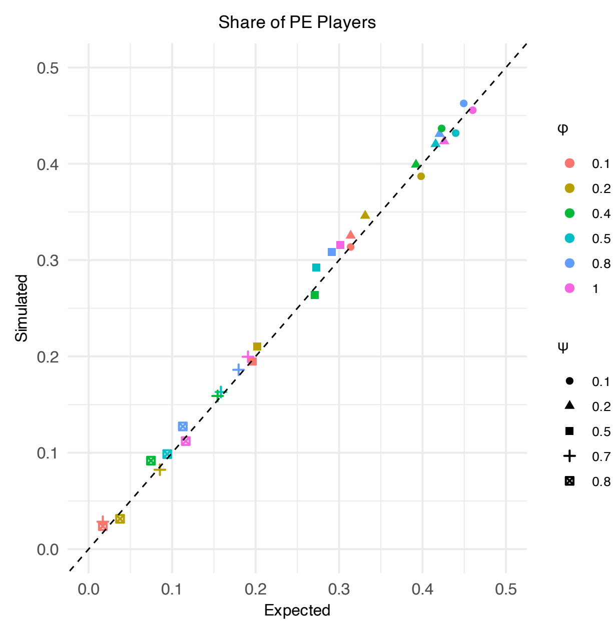

Figure 3: Expected vs Simulated Shares $N_0 = 100$. Simulated shares computed via Monte Carlo simulation of Eqs. 5–7. Sample Size: $M = 100,000$. Constant sample setup.

Finally, I validate via Monte Carlo simulations the results obtained with expected shares $Q^*(N_0, t, \lfloor \phi N_0 \rfloor, \theta)$. The experiment is done repeating for a sample of $M = 100,000$ independent replications (where parameters and initial conditions are kept fixed) the simulation of the process according to Eqs. 5–7. Figure 3 shows (for the case where sample sizes are constant and $N_0 = 100$) a very good match between expected and simulated PE shares. This result is confirmed for larger values of N_0 as well.

6 Conclusions

The foregoing results bear some interesting implications in terms of equilibrium selection in Polya-urn coordination games. To begin with, although long-run shares remain unpredictable and dynamics is governed by path dependence, risk dominance prevails over Pareto efficiency across the whole parameter range. This is in line with findings obtained in fixed-population with reversible choices by the evolutionary-game literature (Kandori et al., 1993; Ellison, 1993).

Furthermore I find that, when players cannot observe global shares, the probability of selecting a Pareto-efficient outcome decreases the smaller the sample size drawn by entrants. This implies that risk-dominance is reinforced as a selection criterion when entrants choose on the basis of a noisier, partial signal. However, Pareto-efficiency can partially offset this tendency the larger the initial pool of incumbents. This finding is somewhat at variance with Lane and Vescovini (1996). The contrast stems from the distinct informational mechanism at work in their contagion model: there, Bayesian optimization with larger samples may perversely entrench inferior options, whereas simple heuristics can ensure convergence on the superior alternative. By contrast, in my setting smaller samples exacerbate noise and favor risk-dominance over Pareto efficiency. The apparent divergence thus reflects the opposite role that sample size plays across the two frameworks.

While the foregoing analysis provides a consistent characterization of equilibrium selection, it is partly constrained by the difficulty of deriving fully closed-form dynamics. For this reason, the paper combines some analytical results —notably the formulation of Q^* as the zero of W— with numerical exploration and simulation-based validation, in line with established practice in the evolutionary-game and Polya-urn literatures.

My results can be checked for robustness against a number of different assumptions about sample-size dynamics (e.g., decreasing sample size) and player choice criteria (e.g., linear vs. non-linear probabilistic decision rules). Furthermore, an interesting avenue for future research would be to relax the assumption that entrants project sample shares directly onto the population and instead allow for Bayesian belief updating based on observed samples. Finally, since the baseline coordination game analyzed here already corresponds to a restricted stag-hunt structure, further extensions could instead explore genuinely different payoff setups, such as prisoner's dilemma matrices or coordination games with asymmetric payoffs.

Acknowledgments

This work was partially supported by the SMaRT COnSTRUCT project (CUP J53C24001460006), in the context of FAIR (PE0000013, CUP B53C22003630006) under the Italian National Recovery and Resilience Plan (Mission 4, Component 2, Line of Investment 1.3) funded by the European Union - NextGenerationEU.

References

Agarwal, S., Ghosh, D. and Chakrabarti, A. S. (2016) "Self-organization in a distributed coordination game through heuristic rules" *The European Physical Journal B* **89**, 266.

Arthur, W. B., Ermoliev, Y., and Kaniovski, Y. (1986) "Strong laws for a class of path-dependent stochastic processes with applications" in *Stochastic Optimization* by Arkin, V. I., Shiraev, A., and Wets, R., Eds., Springer Berlin: Berlin, Heidelberg, 287–300.

Arthur, W. B. (1994) *Increasing returns and path-dependency in economics*, University of Michigan Press: Ann Arbor.

Arthur, W. B., Ermoliev, Y., and Kaniovski, Y. (1987) "Path-dependent processes and the emergence of macro-structure" *European Journal of Operational Research* **30**, 294–303.

Blume, L. (1995) "The statistical mechanics of best-response strategy revision" *Games and Economic Behavior* **11**, 111–145.

Chen, M. and Wei, C. (2005) "A new urn model" *Journal of Applied Probability* **42**, 964–976.

Chen, M.-R., Hsiau, S.-R., and Yang, T.-H. (2014) "A new two-urn model" *Journal of Applied Probability* **51**, 590–597.

Cooper, R. W., DeJong, D. V., Forsythe, R., and Ross, T. W. (1990) "Selection Criteria in Coordination Games: Some Experimental Results" *American Economic Review* **80**, 218–233.

Dosi, G., Ermoliev, Y., and Kaniovski, Y. (1994) "Generalized urn schemes and technological dynamics" *Journal of Mathematical Economics* **23**, 1–19.

Dosi, G. and Kaniovski, Y. (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. (2005a) "Endogenous neighborhood formation in a local coordination model with negative network externalities" *Journal of Economic Dynamics and Control* **29**, 297–319.

Fagiolo, G. (2005b) "A note on equilibrium selection in polya-urn coordination games" *Economics Bulletin* **3**, 1–14.

Harsanyi, J. C. and Selten, R. (1988) *A general theory of equilibrium selection in games*, The MIT Press: Cambridge, MA.

Johnson, N. L. and Kotz, S. (1977) *Urn Models and Their Application*, John Wiley & Sons: New York, NY.

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

Katz, M. L. and Shapiro, C. (1985) "Network Externalities, Competition, and Compatibility" *American Economic Review* **75**, 424–440.

Kim, Y. (1996) "Equilibrium selection in N-Person coordination games" *Games and Economic Behavior* **15**, 203–227.

Lane, D. and Vescovini, R. (1996) "Decision rules and market share: Aggregation in an information contagion model" *Industrial and Corporate Change* **5**, 127–146.

Nelson, R., Peterhansl, A., and Sampat, B. (2004) "Why and how innovations get adopted: a tale of four models" *Industrial and Corporate Change* **13**, 679–699.

Raducha, T. and San Miguel, M. (2022) "Coordination and equilibrium selection in games: the role of local effects" *Scientific Reports* **12**, 3373.

Samuelson, L. (1998) *Evolutionary Games and Equilibrium Selection*, The MIT Press: Cambridge, MA.

Schelling, T. C. (1978) *Micromotives and Macrobehavior*, W. W. Norton & Company: New York, NY.

Young, H.P. (1996) "The economics of convention" *Journal of Economic Perspectives* **10**, 105–122.