Heterogeneous Noise and Stable Miscoordination

Srinivas Arigapudi^{*} Yuval Heller[†] Amnon Schreiber[§]

August 15, 2023

Abstract

Coordination games admit two types of equilibria: pure equilibria, where players coordinate their actions, and mixed equilibria, where players frequently miscoordinate. The existing literature shows that under many evolutionary dynamics, populations always converge to a pure equilibrium. By contrast, we show that under plausible learning dynamics, where agents observe the actions of a random sample of their opponents, stable miscoordination can arise when there is heterogeneity in the sample sizes. This occurs when some agents make decisions based on small samples, while others rely on large samples. Finally, we demonstrate the empirical relevance of our results in the housing market.

Keywords: sampling best response dynamics, action-sampling dynamics, coordination games, hawk-dove games, evolutionary stability, logit dynamics. JELClassification: C72, C73.

1 Introduction

Many real-life situations can be modeled as coordination games, in which the best response against an opponent's action is to play the same action (possibly after relabeling

^{*}Department of Economic Sciences, IIT Kanpur. arigapudi@iitk.ac.in.

[†]Department of Economics, Bar-Ilan University and U. of California San Diego. yuval.heller@biu.ac.il. [‡]Department of Economics, Bar-Ilan University. amnon.schreiber@biu.ac.il.

[§]This paper replaces an earlier working paper titled "Sampling Dynamics and Stable Mixing in Hawk-Dove Games."We are grateful to Vince Crawford, Josef Hofbauer, Bill Zame, and to seminar and confer-

ence participants at Bar-Ilan University, Haifa University, Learning, Evolution and Games, SWET 2023 (UC Irvine), UC Los Angeles, UC Riverside, UC San Diego, and University of South California, for many helpful comments. YH gratefully acknowledges the financial support of the European Research Council (#677057), the Israeli Science Foundation (#448/22), and the US-Israel Binational Science Foundation (#2020022). SA is supported in part at the Technion by a Fine Fellowship.

the actions of one of the players). Two-player, two-action coordination games admit three equilibria of two distinct types: two strict pure equilibria, where the players successfully coordinate their actions, and a mixed equilibrium, where players frequently experience miscoordination. A fundamental result in evolutionary game theory is that under a broad set of learning dynamics, the mixed equilibrium is unstable and populations in which agents are randomly matched to play coordination games must converge to everyone playing one of the pure equilibria (as surveyed in Section 2). By contrast, in this paper we show that the mixed equilibrium with miscoordination can be stable if the populations are heterogeneous in the sense that some (but not all) of the agents rely on anecdotal evidence induced by small samples.

Highlights of the Model Consider a setup in which pairs of agents from two infinite populations are repeatedly randomly matched to play a (one-shot) coordination game.¹ Agents occasionally die and are replaced by new agents (or, alternatively, agents occasionally receive opportunities to revise their actions). The new agents do not have precise information about the aggregate behavior in the opponent's population, and estimate this from sampling the opponent's population. Specifically, each population *i* is characterized by a distribution of sample sizes θ_i , such that $\theta_i(k)$ is the frequency of agents with sample size *k*. Each such agent observes the behavior of *k* random opponents, and then adopts the action that is a best response to her sample (with an arbitrary tie-breaking rule).²

These learning dynamics, which seem plausible in various setups, are called *sampling best-response dynamics* (Sandholm, 2001; Osborne and Rubinstein, 2003; Oyama, Sandholm, and Tercieux, 2015, henceforth abbreviated as *sampling dynamics*).

As explained in Appendix A.1, any two-action coordination game (such as the battle of the sexes or the stag hunt) can be represented w.l.o.g. by the payoff matrix presented in Table 1 that has two positive parameters (u_1, u_2) that represent the players' payoffs when both play the first action: the payoffs when coordinating on the second action are normalized to 1, and the payoffs for miscoordinating are normalized to 0.

The following definition will be helpful for presenting our results. An action is q-dominant (Morris, Rob, and Shin, 1995) to player i if it is the player's best response

¹Appendix A.2 extends our results to symmetric coordination games played within a single population.

²This behavior can be interpreted by each new agent utilizing her own sample to calculate a maximum likelihood estimation for the overall behavior of the opposing population.

	a_2	b_2	
a_1	$u_1^{u_2}$	0 0	
b_1	0	1	

Table 1: Standard Representation of Two-Action Coordination Game

against any opponent's mixed action that assigns a mass of at least q to the opponent playing the same action. In particular, action a_i (resp., b_i) is $\frac{1}{2}$ -dominant for player i if it is the best response against an opponent's uniform play, which is the case iff $u_i > 1$ (resp., $u_i < 1$).

Global Convergence to Miscoordination Theorem 1 presents a full characterization for environments in which the populations converge to states with miscoordination from almost all initial states. This happens if (and essentially only if): (1) each population has a different $\frac{1}{2}$ -dominant action (i.e., $u_2 < 1 < u_1$), (2) the product of the mass of agents with sample size 1 times the expected sample size of the other population is larger than one (i.e., $\forall i, \theta_i(1) \cdot \mathbb{E}(\theta_j) > 1$), and (3) each population has a q-dominant action for a sufficiently low q (which is satisfied iff each u_i is sufficiently far from 1).

The proof idea is as follows. Global convergence to miscoordination occurs iff both pure equilibria are unstable. Assume that $u_2 < 1 < u_1$. Consider a slightly perturbed state near the equilibrium (a_1, a_2) , in which $\epsilon_i << 1$ of the agents in population *i* plays b_i . Events in which a new agent observes multiple occurrences of the opponent's rare action b_j in her sample are negligible $(O(\epsilon_j^2))$. Neglecting these very rare events implies that new agents of population 1 will adopt the risk-dominated action b_1 only when they have sample size 1, and they have observed the rare action b_2 (the probability of this is $\theta_1(1) \cdot \epsilon_2$). By contrast, if $u_2 < 1$ is sufficiently small, then a single occurrence of the rare action b_1 in a sample of size k (which occurs with probability of $k \cdot \epsilon_1$) is sufficient to induce a new agent of population 2 to play her $\frac{1}{2}$ -dominant action b_2 . This implies that the total share of new agents of population 2 who play action b_2 is $\mathbb{E}(\theta_2) \cdot \epsilon_1$. This, in turn, implies that the product of the number of agents playing the rare action in each population increases iff $\theta_1(1) \cdot \mathbb{E}(\theta_2) > 1$.

Heterogeneity and Stable Miscoordination Our second main result (Theorem 2) shows that heterogeneity in the sample sizes is necessary for stable miscoordination.

Specifically, we show that if all agents in each population i have the same sample size $k_i \ge 2$, then all states with miscoordination are unstable (the case in which everyone has sample size 1 is discussed in Remark 1). Our final result (Theorem 3) shows that many heterogeneous distributions of sample sizes in which some agents have relatively small samples and others have sufficiently large samples induce locally stable states with miscoordination if u_1 and u_2 are not too close to 1.

The intuition for why heterogeneity in the sample sizes is important for the stability of interior states with miscoordination is as follows. Consider homogeneous populations with a fixed sample size k in a stationary interior state. In such a state the random sample of size k frequently yields both outcomes for which action a_i is a best response and outcomes for which b_i is a best response. We show that in such situations the probability of each action being a best response is sensitive to small perturbations in the opponent's distribution of actions. That is, if ϵ more of the opponent's population play a_j , it increases the probability that a_i is the best response to the random sample by more than ϵ .

Next consider a heterogeneous population in which some agents have relatively small samples, while other agent have large samples. The stationary interior state with miscoordination typically does not coincide with the mixed Nash equilibrium, which implies that almost all agents with large samples play the same action (the unique best response to the true distribution of the opponents' actions), and that their play is insensitive to small perturbations of the opponents' behavior. This allows the overall sensitivity of the entire population to small perturbations to be sufficiently small to allow stable miscoordination.

Theorem 2's proof relies on deriving a property of binomial distributions (Proposition 3), which may be of independent interest. This property states that a composition of two binomial cumulative distributions has a unique interior fixed point.

Numerical Analysis and Insights In Section 7 we demonstrate our main insight that heterogeneity in the sample sizes induces stable miscoordination does not depend on the specific details of the sampling dynamics. Specifically, we numerically study the commonly used logit dynamics, in which agents play a noisy best response to their opponents' aggregate behavior, with η_i describing the noise level in population *i*. We first demonstrate that if the noise level in each population is homogeneous, then one can induce stable miscoordination only with implausibly high levels of noise. By contrast, when we introduce an extension of logit dynamics that allows heterogeneity in the level of noise in each population, we show that stable miscoordination can be supported by moderate heterogeneous levels of noise.

Taken together, our results show that the conventional wisdom that miscoordination is unstable is not accurate. Miscoordination can be stable in heterogeneous populations in which some agents rely on anecdotal evidence or noisy data, while other agents have access to more accurate data. The experimentally testable implications of our results are discussed in Section 2.

Bargaining Application and Empirical Relevance Demonstrating the empirical relevance of our theoretical findings requires real-life scenarios in which it is plausible to have: (1) heterogeneous sample sizes, (2) persistent miscoordination within an interior state, and (3) a substantial proportion of agents not maximizing their payoffs with respect to the aggregate behavior of the other population. We argue that these conditions might exist in bargaining situations in both housing and used-car markets. Coordination games can model simple bargaining in these markets, where each agent has two strategies: a high-price strategy or low-price strategy. For example, a seller can either demand a high price (with the risk of bargaining failure) or agree to a low price, and a buyer can either agree to a high price or insist on a low price (as modeled by hawk-dove coordination games; see Example 3 in Appendix A.1). Heterogeneity in the sample sizes is likely to exist in these markets, where both professional players (real-estate investors and car dealers) and inexperienced players (those who have only bought or sold houses or cars a couple of times) participate. Professional players are likely to have access to large samples, while inexperienced players may rely on anecdotal evidence.

Koster and Rouwendal (2021) conducted an empirical study on the Dutch housing market. They found that a significant number of sellers set low list prices, which can be consistent with profit maximization only if the sellers have high annual discount factors of up to 50%. This behavior is consistent with an interior stable state in our model in which some sellers rely on small, noisy samples to choose non-payoff-maximizing low prices. Koster and Rouwendal's data does not allow separating between different types of sellers. By contrast, Genesove and Mayer (2001) does allow separating the sellers to owner-occupants and investors, and show that the two groups have systematic differences in their list prices.

Larsen (2021) empirically studied the efficiency of bargaining following wholesale usedcar auctions in which the highest bid is lower than the reserve price, and found that bargaining fails in about 35% of these cases (which is consistent with an interior stable state with miscoordination). The data suggests that this inefficiency results in substantial losses, representing 12%-23% of ex-post gains from trade. Furthermore, Larsen's analysis suggests that only a small part of this loss is due to incomplete information constraints.

Structure Section 2 presents the related literature. Our model is described in Section 3. Section 4 presents a "complete" characterization of global convergence to states with miscoordination. In Section 5, we show that homogeneous populations always converge to one of the pure equilibria. Section 6 shows that many heterogeneous populations can induce locally stable states with miscoordination. In Section 7, we numerically analyze logit dynamics, and demonstrate that our main insights hold in this setup as well. We conclude in Section 8. Formal proofs and additional results are presented in the appendix. Appendix A.2 extends our results to one-population dynamics. Appendix A.3 states and proves an interesting general result on binomial distributions.

2 Related Literature

Instability of Miscoordination It is well-known that strict equilibria satisfy strong stability refinements, while mixed equilibria with miscoordination do not satisfy even weak stability refinements. In particular, strict equilibria are evolutionarily stable (Maynard Smith and Price, 1973), while the mixed equilibrium does not satisfy neutral stability (Maynard Smith, 1982) or even the mild refinement of weak stability (Heller, 2017). Moreover, it is well known that interior stationary states in all multiple-population games cannot be asymptotically stable under the widely studied replicator dynamics (see, e.g., Sandholm, 2010, Theorem 9.1.6). Instability of interior states is further shown for various classes of learning dynamics in Crawford (1989).

Moreover, various papers in the literature have proven that populations playing coordination games converge to one of the pure equilibria from almost all initial states under various dynamics. Kaniovski and Young (1995) show that there is global convergence to one of the pure equilibria for sampling dynamics with sufficiently large samples. Oprea, Henwood, and Friedman (2011) show that there is global convergence to one of the pure equilibria under monotone dynamics (i.e., under any dynamics in which an action becomes more frequent iff it yields a higher payoff than the alternative action).³ Recently, Oyama, Sandholm, and Tercieux (2015) proved the global stability of the $\frac{1}{2}$ -dominant pure equilibrium in symmetric coordination games for distributions of sample sizes in which sufficiently many agents have sufficiently small samples. The stochastic evolutionary dynamics literature (pioneered by Kandori, Mailath, and Rob, 1993; Young, 1993; see also the recent application to hawk-dove games in Bilancini, Boncinelli, Ille, and Vicario, 2022) shows that only pure equilibria can be stochastically stable in large finite populations in which agents most of the time best respond to a large sample, but occasionally mistakenly play the other action.

Thus, taken together, the various existing studies suggest that states with miscoordination are unstable. Our contribution is in showing that this is not the case in plausible learning dynamics in which agents base their behavior on sampling the actions of the other population, and there is substantial heterogeneity in the sample sizes in each population.

Sampling Dynamics Sampling best-response dynamics (henceforth, sampling dynamics) were pioneered by Sandholm (2001) and Osborne and Rubinstein (2003). As argued by Oyama, Sandholm, and Tercieux (2015), the deterministic nature of sampling dynamics implies that when there is convergence to a stable state (which is always the case in our setup; see Proposition 2) the convergence is fast.⁴ Recently, Heller and Mohlin (2018) studied the conditions on the expected sample size that implies global convergence for all payoff functions and all sampling dynamics.⁵

Salant and Cherry (2020) (see also Sawa and Wu, 2023) generalized sampling dynamics by allowing new agents to use various procedures to infer from their samples the aggregate behavior of the opponents (in addition to allowing for payoff heterogeneity in the population). Salant and Cherry paid special attention to *unbiased* inference proce-

³Oprea, Henwood, and Friedman (2011) proved this for hawk-dove games, but the proof can be extended to all coordination games.

⁴Conditions in which stochastic dynamics induce fast convergence are studied in Kreindler and Young (2013) and Arieli, Babichenko, Peretz, and Young (2020).

⁵Hauert and Miekisz (2018) used the term "sampling dynamics" to refer to a variant of the replicator dynamics, in which when an agent samples another agent and mimics the other agent's behavior, it is more likely that these two agents will be matched with each other. This is less related to our use of the notion of "sampling dynamics," which is in line with the literature cited above.

dures in which the agent's expected belief about the share of opponents who play an action coincides with the sample mean. Examples of unbiased procedures are maximum likelihood estimation, beta estimation with a prior representing complete ignorance, and a truncated normal posterior around the sample mean. In our setup, the payoffs are linear in the share of agents who play action a_j , which implies that the agent's perceived best response depends only on the expectation of her posterior belief. This implies that the behavior of agents in our model can be interpreted as each new agent utilizes her own sample to calculate an unbiased estimation procedure (such as, maximum likelihood estimation) for the overall behavior of the opposing population. remains the same with any unbiased inference procedure.

Experimental Literature and Testable Predictions Our model yields a novel testable prediction: miscoordination can persist when there is heterogeneity in the amount of information that the agents have about the opponents' behavior, and for each population one of the actions is q-dominant for a sufficiently low q. Brunner, Camerer, and Goeree (2011) compared the predictive power of various learning models when applied to the experimental data of McKelvey, Palfrey, and Weber (2000), Goeree, Holt, and Palfrey (2003), and Selten and Chmura (2008), in which (1) agents are randomly and repeatedly matched to play a one-shot two-action two-player game against an anonymous opponent, and (2) agents get feedback only about their most recent opponent's play (and can rely on their memory of her previous feedback about past opponents). Brunner, Camerer, and Goeree's findings show that sampling dynamics explain well the aggregate experimental behavior, and that sampling dynamics' predictions are: (1) as good as the quantal response equilibrium in all games and significantly better in some games, (2) significantly better than the Nash equilibrium in all games, and (3) significantly improve when focusing on the less noisy second halves of the treatments (see, Brunner, Camerer, and Goeree, 2011, Figure 4).

Brunner, Camerer, and Goeree's results suggest that sample sizes of 3–12 fit the data best (which is roughly in line with the typical estimates of people's short-run memory capacity). The prediction of global convergence to miscoordination (Theorem 1) requires that a substantial part of the population to have a small sample size of 1. This may be induced by having some subjects playing in different rounds different underlying games. For example, in each round, one of several underlying games will be played. Each subject will be informed about the payoff matrix of the current game, and will be reminded about her most recent past opponent's behavior in the previous period in which the same game was played. It seems plausible that many subjects will rely only on this feedback when deciding how to play (i.e., essentially have a sample of size 1), while only a few subjects will exert effort to remember relevant past feedback from less recent periods in which the same game game was played.

Recently, Lyu, Li, and Xu (2022) applied a more elaborate experimental setup that aims to implement sampling dynamics. In their implementation the subjects are endowed with a specific "default" action to play in the first round. This novel design component allows us to test dynamic predictions with respect to specific initial states. At the end of each round each agent observes the k most recent actions (k = 2 and k = 7 were applied in their treatments) of random opponents among the 14 players in their matching group. Their underlying games were symmetric coordination games with Pareto-ranked pure equilibria. Lyu, Li, and Xu showed that action-sampling dynamics (with respect to the most recent k observed actions) fits about 80% of the subjects' behavior.

A central reason why some subjects deviated from the predictions of sampling dynamics (especially in early rounds) was that they played the action that is part of the Pareto-dominant equilibrium in order to "teach" the other subjects in the matching group to move from the Pareto-dominated equilibrium to the Pareto-dominant equilibrium. These "teaching" incentives were justified, as they often helped groups starting in the Pareto-dominated equilibrium to shift to the Pareto-dominant equilibrium. We think that "teaching" incentives would be substantially reduced, and the fit of sampling dynamics would be substantially improved, if either (1) the strict equilibria of the underlying coordination games were not Pareto-ranked (e.g., battle of the sexes), or (2) the matching groups were substantially larger.

The predictions of Theorem 3 (local stability of miscoordination) require both an appropriate initial state (which can be implemented à la Lyu, Li, and Xu, 2022), and that some agents will have accurate information about the opponents' aggregate behavior. This can be implemented by providing some agents with full feedback about the behavior of all opponents in the previous round (while providing the remaining agents with feedback about the behavior of only their own matched opponents).

3 Model

3.1 Coordination Games

Let $G = \{A_1 \times A_2, u\}$ denote a normal-form two-action two-player coordination game. Let $i \in \{1, 2\}$ be an index denoting one of the players ("she"), and let j denote her opponent ("he"). For each $i \in \{1, 2\}$, let $A_i = \{a_i, b_i\}$ denote the actions of player i.

The standard two-parameter payoff matrix of a coordination game is given in Table 2: the players get a low payoff (normalized to 0) if they miscoordinate (i.e., one player plays a_i and the opponent plays b_j), they get a high payoff (normalized to 1) if they coordinate on both playing (b_1, b_2) , and they get a payoff of $u_i > 0$ if they coordinate on both playing (a_1, a_2) . By relabeling the actions we can assume w.l.o.g. that action profile **a** is the weakly preferred outcome of Player 1; i.e., we assume that $u_1 \ge 1$. We say that the coordination game is (1) symmetric if $u_1 = u_2$, and (2) antisymmetric if $u_1 = \frac{1}{u_2}$. Antisymmetric coordination games can be interpreted as battle of the sexes games, in which Player 1's preference for the first coordinated outcome **a** over the second coordinated outcome **b** is as strong as Player 2's preference for **b** over **a**.

In Appendix A.1 we formally show that this standard two-parameter representation captures w.l.o.g. two-action coordination games (including battle of the sexes, stag hunt, and hawk-dove games). This is so because our dynamics (as defined in Eq. (3.1) and Eq. (A.1))) depend only on the differences between the payoffs of a player in action profiles in which the opponent plays the same action. This implies that the dynamics are invariant to (1) adding a constant to the two payoffs of Player 1 (resp., Player 2) in the same row (resp., column), and (2) dividing all the payoffs of a player by a positive constant.

We extend the game to mixed actions in the standard linear way. We identify each mixed action with the probability it assigns to the first action (a_i) , and we denote it by $p_i \in [0,1]$. We identify the degenerate mixed action 1 (resp., 0) with the pure action a_i (resp., b_i). Observe that the coordination game admits three Nash equilibria: two pure equilibria, (a_1, a_2) and (b_1, b_2) , in which the players coordinate their actions, and a mixed equilibrium $\mathbf{p}^{NE} = \left(\frac{1}{1+u_2}, \frac{1}{1+u_1}\right)$ in which the players frequently experience miscoordination.

	a_2	b_2		
a_1	$u_1^{u_2}$	0 0		
b_1	0	1 1		
$u_1 \ge 1, u_2 > 0$				

Table 2: Standard Representation of Two-Action Coordination Game

¹/₂-Dominance Fix $q \in [0, 1]$. We say that action a_i (resp., b_i) is *q*-dominant (Morris, Rob, and Shin, 1995; Oyama, Sandholm, and Tercieux, 2015) for player *i* if it is a strict best response to any opponent's mixed action that assigns a mass of at least *q* to the counterpart action a_j (resp., b_j).

Notice that both actions are 1-dominant (which is equivalent to being part of a strict equilibrium). Additionally, it can be noted that as q decreases, the q-dominance condition becomes more stringent. In other words, if an action is q-dominant, it also satisfies r-dominance for any q between r and 1. Lastly, it can be observed that action a_i is $\frac{1}{1+u_i}$ -dominant, while action b_i is $\frac{u_i}{1+u_i}$ -dominant.

We say that the pure equilibrium **a** (resp., **b**) is $\frac{1}{2}$ -dominant if each action a_i (resp., b_i) is $\frac{1}{2}$ -dominant; in this case we say that the remaining equilibrium is risk dominated. Observe that the game admits a $\frac{1}{2}$ -dominant equilibrium iff $u_1, u_2 > 1$. By contrast, if $u_1 > 1 > u_2$, then a_1 is $\frac{1}{2}$ -dominant for player 1 and b_2 is $\frac{1}{2}$ -dominant for player 2.

Observe that q-dominance depends only on the differences between the payoffs a player can get by playing the different actions. This implies that q-dominance is invariant to both of the payoff transformations mentioned above (and described in detail in Appendix A.1). This implies that an action is q-dominant in the standard representation of Table 2 iff it is q-dominant in the original representation as a general coordination game (as described in Appendix A.1). By contrast, payoff dominance is not invariant to the first of these two transformations. Specifically, adding a constant to the two payoffs of Player 1 (Player 2) in the same row (column) might change a Pareto-dominated equilibrium into a Pareto-dominant equilibrium.

3.2 Evolutionary Dynamics

We assume that there are two unit-mass continuums of agents and that agents in population 1 are randomly matched with agents in population 2. Aggregate behavior at time $t \in \mathbb{R}^+$ is described by a state $\mathbf{p}(t) = (p_1(t), p_2(t)) \in [0, 1]^2$, which is equivalent to a mixed action profile (i.e., $p_i(t)$ represents the share of agents playing action a_i at time t in population i). A state $\mathbf{p} = (p_1, p_2)$ is interior (i.e., mixed) if $p_1, p_2 \in (0, 1)$.

Agents die at a constant rate of 1, and are replaced by new agents (or, equivalently, agents get opportunities to revise their actions). The evolutionary process is represented by a function $\mathbf{w} : [0,1]^2 \to [0,1]^2$, which describes the frequency of new agents in each population who play action a_i as a function of the current state. Thus, the instantaneous change in the share of agents in population *i* who play a_i is given by the following dynamics (where w_i is as defined in Eq. (A.1) below):

$$\dot{p}_i = w_i \left(p_j \right) - p_i. \tag{3.1}$$

Sample Sizes We allow heterogeneity in the sample sizes used by new agents. Let $\theta_i \in \Delta(\mathbb{Z}_+)$ denote the distribution of sample sizes of new agents in population *i*. We assume that θ_i has a finite support. A share of $\theta_i(k)$ of the new agents have a sample of size *k*. Let supp (θ_i) denote the support of θ_i . If there exists some *k*, for which $\theta_i(k) = 1$, then we use *k* to denote the degenerate (homogeneous) distribution $\theta_i \equiv k$.

Remark 1. We assume that some agents have sample sizes larger than 1 (i.e., $\forall i \max(\sup p(\theta_i)) > 1$). This rules out the trivial case in which all agents have sample size 1. In this case, the entire diagonal $\{(p, p) | p \in [0, 1]\}$ is Lyapunov stable, and no state is asymptotically stable (as shown in Sethi, 2000, Example 7 in a related setup).⁶

Definition 1. An *environment* is a tuple $(\boldsymbol{u}, \boldsymbol{\theta}) = ((u_1, u_2), (\theta_1, \theta_2))$, where (u_1, u_2) describe the payoffs of the underlying coordination game, and (θ_1, θ_2) describes the distributions of sample sizes in each population.

Sampling Best-Response Dynamics Sampling best-response dynamics (henceforth *sampling dynamics*; Sandholm, 2001; Oyama, Sandholm, and Tercieux, 2015) fit situations in which agents do not know the exact distribution of actions in the other pop-

⁶See Appendix A.4 for the formal definitions of Lyapunov and asymptotic stability. Lyapunov stability requires that populations starting near the state remain close to the state, and asymptotic stability further requires that the populations eventually converge to the state. The intuition for the stability of the entire main diagonal is that new agents with sample size 1 always adopt the action that they observe. If all new agents have sample size 1, this implies that the share of new agents who play a_i coincides with the share of incumbents of the other population who play a_j . This, in turn, implies that the populations move from any initial state toward the closest point in the main diagonal, until reaching a stationary state in which both populations have the same aggregate behavior.

ulation. New agents estimate this unknown distribution by sampling the opponent's actions. Specifically, each new agent with sample size k (henceforth, k-agent) samples k randomly drawn agents from the other population and then plays the action that is the best response against the sample. One possible interpretation of this behavior is that each new agent utilizes her own sample to calculate a maximum likelihood estimation for the overall behavior of the opposing population (or by any other unbiased estimation procedure, à la Salant and Cherry, 2020).

To simplify notation, we assume that in case of the tie, the new agent plays a_i . Our results are essentially the same for any tie-breaking rule.

Let $X(k, p_j) \sim Bin(k, p_j)$ denote a random variable with binomial distribution with parameters k (number of trials) and p_j (probability of success in each trial), which is interpreted as the number of a_j -s in the sample. Observe that the sum of payoffs of playing action a_i against the sample is $u_i \cdot X(k, p_j)$ and the sum of payoffs of playing action b_i against the sample is $k - X(k, p_j)$.

This implies that action a_i is a best response to a sample of size k iff $u_i \cdot X(k, p_j) \ge k - X(k, p_j) \Leftrightarrow X(k, p_j) \ge \frac{k}{u_i+1}$. This, in turn, implies that the sampling dynamics for environment $(\boldsymbol{u}, \boldsymbol{\theta})$ is given by

$$w_i(p_j) = \sum_{k \in \text{supp}(\theta_i)} \theta_i(k) \cdot \Pr\left(X(k, p_j) \ge \frac{k}{u_i + 1}\right).$$
(3.2)

Observe that $\Pr\left(X\left(k,p_{j}\right) \geq \frac{k}{u_{i}+1}\right) = \sum_{l=\left\lceil \frac{k}{u_{i}+1} \right\rceil}^{k_{i}} \binom{k}{l} p_{j}^{l} \left(1-p_{j}\right)^{k-l}$, which is a polynomial of p_{j} of degree k. This implies that $w_{i}\left(p_{j}\right)$ is a polynomial with a finite degree of max supp $(\theta_{i}) > 1$.

Remark 2. Our model deals with coordination games played between two different populations (or, equivalently, with games played within a single population where an agent can condition her play on the role she was assigned to in the game: Player 1 or Player 2). In Appendix A.2 we show that our results remain the same under the alternative *one-population dynamics* that can be applied to symmetric coordination games ($u_1 = u_2$) in which a player cannot condition her play on her assigned role of Player 1 or Player 2.

4 Global Convergence to Miscoordination

As discussed in Section 2, various existing papers have shown that populations playing coordination games will always converge to pure equilibria under various dynamics. In this section, we fully characterize the conditions for which the opposite result holds under sampling dynamics; i.e., the populations converge from almost any initial state to interior states with miscoordination.

4.1 Analysis of $w_i(\mathbf{p})$

The characteristics of sampling dynamics are closely related to the properties of the polynomials $w_1(p_2)$ and $w_2(p_1)$ and their intersection points, which are analyzed in this subsection.

Figure 4.1 illustrates the phase plots of sampling dynamics and the properties of the curves $w_2(p_1)$ and $w_1(p_2)$. We refer to the latter curve also as $w_1^{-1}(p_1)$. The left panel illustrates a symmetric coordination game with $u_1 = u_2 = 3$, and the right panel illustrates an antisymmetric coordination game with $u_1 = 3$ and $u_2 = \frac{1}{3}$. In both panels all agents have sample size of 3 ($\theta_i \equiv 3$). The dashed orange curve is the polynomial $p_2 = w_2(p_1)$, which describes the states in which $\dot{p}_2 = 0$. In all states above (resp., below) this curve $\dot{p}_2 < 0$ (resp., $\dot{p}_2 > 0$). The solid purple curve is the polynomial $p_1 = w_1(p_2)$, which describes the states in which $\dot{p}_1 = 0$. In all states to the right (resp., left) of this curve $\dot{p}_1 < 0$ (resp., $\dot{p}_1 > 0$). Observe that in both panels, all nonstationary initial states converge to a pure state. In the left panel there is global convergence to (1, 1), while in the right panel some states converge to (1, 1) and others to (0, 0).

The following fact is immediate given the basic properties of binomial random variables.

Fact 1. $w_i(p_j)$ is a strictly increasing polynomial function that satisfies $w_i(0) = 0$ and $w_i(1) = 1$. This implies that the inverse function $w_i^{-1} : [0,1] \to [0,1]$ exists, is continuously differentiable, and that $w_i^{-1}(0) = 0$ and $w_i^{-1}(1) = 1$.

Fact 1 implies that the two curves intersect at (0,0) and (1,1).

Appendix A.4 presents the standard definitions of stationary states, asymptotically stable states, and unstable states. Observe that a state is stationary (i.e., it is a fixed point of the dynamics) iff it is an intersection point of the two curves w_1 and w_2 .



The figure illustrates the phase plots of sampling dynamics for two environments: (1) $\forall i \ \theta_i \equiv 3 = u_i$ (left panel, a symmetric coordination game), and (2) $\forall i \ \theta_i \equiv 3 = u_1 = \frac{1}{u_2}$ (right panel, an antisymmetric coordination game). The solid purple (resp., dashed orange) curve of $w_1 (p_2)$ (resp., $w_2 (p_1)$) shows the states for which $\dot{p}_1 = 0$ (resp., $\dot{p}_2 = 0$). The intersection points of these curves are the stationary states. A solid (resp., hollow) dot represents an asymptotically stable (resp., unstable) stationary state.

Fact 2. State **p** is stationary $\iff p_2 = w_2(p_1)$ and $p_1 = w_1(p_2) \implies p_1 = w_1(w_2(p_1))$ and $p_2 = w_2(w_1(p_2))$.

4.2 Asymptotic Stability of Pure Equilibria

To state our next results, it will be helpful to consider the condition in which a single appearance of a rare action can change the behavior of a new agent. Specifically, consider a new agent in population i with a sample size of k. Observe that:

- 1. Action a_i induces a weakly higher payoff against a sample with an opponent's action a_j (and k-1 opponents' actions b_j) iff $u_i \ge k-1 \Leftrightarrow k \le u_i+1$.
- 2. Action b_i induces a strictly⁷ higher payoff against a sample with an opponent's action b_j (and k-1 opponents' actions a_j) iff $1 > (k-1)u_i \Leftrightarrow k < \frac{1}{u_i} + 1$.

Next we define m-bounded expectation as the expected value of a probability distribution when taking into account only values smaller than m. Formally:

⁷We require strictly higher payoffs for action b_i and weakly higher payoffs for action a_i due to our tie-breaking rule in favor of action a_i .

Definition 2. The *m*-bounded expectation $\mathbb{E}_{\leq m}$ (resp., $\mathbb{E}_{< m}$) of distribution θ_i with support on integers is⁸ $\mathbb{E}_{\leq m}(\theta_i) = \sum_{1 \leq k \leq m} \theta_i(k) \cdot k$ (resp., $\mathbb{E}_{< m}(\theta_i) = \sum_{1 \leq k < m} \theta_i(k) \cdot k$).

Proposition 1 characterizes the asymptotic stability of the pure equilibria. It shows that asymptotic stability depends only on whether the product of the bounded expectations of the distributions of sample sizes is larger or smaller than one, where the bound of each distribution is the maximal sample size for which a single appearance of a rare action can change the behavior of a new agent. Formally:⁹

Proposition 1.

- 1. $\mathbb{E}_{<\frac{1}{u_1}+1}(\theta_1) \cdot \mathbb{E}_{<\frac{1}{u_2}+1}(\theta_2) = \theta_1(1) \cdot \mathbb{E}_{<\frac{1}{u_2}+1}(\theta_2) > 1 \Rightarrow \mathbf{a} = (a_1, a_2) \text{ is unstable;}$ 2. $\mathbb{E}_{<\frac{1}{u_1}+1}(\theta_1) \cdot \mathbb{E}_{<\frac{1}{u_2}+1}(\theta_2) = \theta_1(1) \cdot \mathbb{E}_{<\frac{1}{u_2}+1}(\theta_2) < 1 \Rightarrow \mathbf{a} \text{ is asymptotically stable;}$
- 3. $\mathbb{E}_{\leq u_1+1}(\theta_1) \cdot \mathbb{E}_{\leq u_2+1}(\theta_2) > 1 \Rightarrow \boldsymbol{b} = (b_1, b_2)$ is unstable; and
- 4. $\mathbb{E}_{\leq u_1+1}(\theta_1) \cdot \mathbb{E}_{\leq u_2+1}(\theta_2) < 1 \Rightarrow \boldsymbol{b} = (b_1, b_2)$ is asymptotically stable.

Sketch of Proof. Consider a slightly perturbed state $(1 - \epsilon_1, 1 - \epsilon_2)$ near $\boldsymbol{a} = (1, 1)$ (the argument for \boldsymbol{b} is analogous) in which almost all agents play action a_i . The event of two rare actions $(b_i$ -s) appearing in a sample of a new agent has a negligible probability of $O(\epsilon_i^2)$. If a new agent has a sample size of k, then the probability of a rare action appearing in the sample is approximately $k \cdot \epsilon_j$. This rare appearance changes the perceived best response of a new agent in population i iff k is smaller than $\frac{1}{u_i} + 1$. Thus, the probability that a new agent in population i adopts the rare action b_i is equal to $\mathbb{E}_{<\frac{1}{u_i}+1}(\theta_i)$. This implies that the product of the share of new agents adopting a rare action in each population is $\epsilon_1 \cdot \mathbb{E}_{<\frac{1}{u_1}+1}(\theta_1) \cdot \epsilon_2 \cdot \mathbb{E}_{<\frac{1}{u_1}+1}(\theta_2)$. This shows that the share of agents playing rare actions gradually increases (resp., decreases) if $\mathbb{E}_{<\frac{1}{u_1}+1}(\theta_1) \cdot \mathbb{E}_{<\frac{1}{u_1}+1}(\theta_2) < 1$), which implies instability (resp., asymptotic stability). Finally, observe that our assumption that $u_1 \ge 1$ implies that $\frac{1}{u_1} + 1 \le 2$, which, in turn, implies that $\theta_1(1) = \mathbb{E}_{<\frac{1}{u_1}+1}(\theta_1)$. See Appendix A.5 for a formal proof.

An interesting implication of Proposition 1 is the substantial difference in the stability of a $\frac{1}{2}$ -dominant equilibrium and a risk-dominated equilibrium (related results for

⁸Observe that in our notation the parameter k takes only (positive) integer values (although we allow the upper bound m to be a noninteger).

⁹Replacing the a_i -favorable tie-breaking rule with a b_i -favorable one would replace the "<"-s and the " \leq "-s in the bounded expectations in the statement of Proposition 1.

symmetric coordination games are derived in Oyama, Sandholm, and Tercieux, 2015). A risk-dominated equilibrium (say, **b**) is unstable as long as sufficiently many agents have samples that are not too large (i.e., samples below $u_i + 1$). The process inducing instability is as follows. A small perturbation of a few agents who play a_i induces a slightly larger number of new agents to observe a_i at least once in their samples, which induces them to play a_i as well, which allows the small perturbations to gradually increase, until, at the end of the process, everyone plays the $\frac{1}{2}$ -dominant action a_i .

By contrast, any $\frac{1}{2}$ -dominant equilibrium is asymptotically stable. To see this, observe that if \boldsymbol{a} is a weakly $\frac{1}{2}$ -dominant equilibrium (i.e., if $u_2 \geq 1$, where the weak inequality is due to the tie-breaking rule in favor of a_i) then $\frac{1}{u_2} + 1 \leq 2$, which implies that $\mathbb{E}_{<\frac{1}{u_2}+1}(\theta_2) = \theta_2(1) < 1$. This implies that $\mathbb{E}_{<\frac{1}{u_1}+1}(\theta_1) \cdot \mathbb{E}_{<\frac{1}{u_2}+1}(\theta_2) = \theta_1(1) \cdot \theta_2(1) < 1$, and thus \boldsymbol{a} is asymptotically stable. This implies that:

Corollary 1. Any $\frac{1}{2}$ -dominant equilibrium is asymptotically stable.

Thus, only in games in which neither equilibrium is $\frac{1}{2}$ -dominant (i.e., those in which $u_2 < 1 \le u_1$), might it be possible for both pure equilibria to be unstable.

4.3 Global Convergence Results

The following definition of neighboring stationary states will be helpful for stating our results.

Definition 3. Two stationary states $\underline{p}, \overline{p}$ are said to be neighbors if there does not exist any stationary state \hat{p} such that $\underline{p}_i < \hat{p}_i < \overline{p}_i$.

We first show that the populations always converge to a stationary state.¹⁰

Proposition 2. $\lim_{t\to\infty} \mathbf{p}(t)$ converges to a stationary state for any $\mathbf{p}(0)$.

Proof. Since the function $w_i(\cdot)$ is a polynomial of degree $\max(\text{supp}(\theta_i)) > 1$, it follows that $w_i(w_j(\cdot))$ is a polynomial of a finite degree strictly larger than 1. As the stationary states are characterized by the intersection of the curves $p_1 = w_1(p_2)$ and $p_2 = w_2(p_1)$, we can conclude that there are a finite number of stationary states. The result now follows from the following claims whose proofs are presented in Appendix A.6.

¹⁰One could present an alternative proof to Proposition 2 that relies, in part, on the Bendixson–Dulac Theorem (see Theorem 9.A.6 of Sandholm, 2010) to show that there are no closed orbits. However, we prefer presenting a direct proof due to its simplicity, and because some of its arguments are used later in the results of Section 5.

Claim 1 (Any trajectory reaches the area between the curves). Either $\lim_{t\to\infty} \mathbf{p}(t)$ converges to a stationary state, or there exists $t < \infty$ such that either $w_1^{-1}(p_1(t)) \leq p_2(t) \leq w_2(p_1(t))$ or $w_2(p_1(t)) \leq p_2(t) \leq w_1^{-1}(p_1(t))$.

The intuition for Claim 1 (see Figure 4.1) is that populations starting at an initial state below (resp., above) the two curves move upward and to the left (resp., downward and to the right) until reaching the area between the two curves.

Claim 2 (Convergence from the area between the curves to a neighboring stationary state). Let $\mathbf{p}(t)$ be an interior state for some $t \ge 0$ and let $\underline{\mathbf{p}}, \overline{\mathbf{p}}$ be neighboring stationary states with $\underline{p}_i \le p_i(t) \le \overline{p}_i$. We have,

1. If
$$p_{2}(t) \in \left[w_{1}^{-1}(p_{1}(t)), w_{2}(p_{1}(t))\right]$$
, then $\lim_{t\to\infty} \mathbf{p}(t) = \overline{\mathbf{p}}$, and
2. If $p_{2}(t) \in \left[w_{2}(p_{1}(t)), w_{1}^{-1}(p_{1}(t))\right]$, then $\lim_{t\to\infty} \mathbf{p}(t) = \underline{\mathbf{p}}$.

The intuition for Claim 2 (see Figure 4.1) is that in the area between the two curves in the interval between \underline{p} and \overline{p} , the populations either always move upward and to the right (in which case they must converge to \overline{p}), or they always move downward and to the left (in which case they must converge to p).

Next we show that if any initial interior state converges to one of the pure equilibria, then this equilibrium is asymptotically stable (as defined in Appendix A.4).

Lemma 1. Assume that $\mathbf{p}(0) \neq (0,0)$ and $\lim_{t\to\infty} \mathbf{p}(t) = (0,0)$; then (0,0) is asymptotically stable. The same result holds when (1,1) replaces (0,0).

Proof. See Appendix A.7.

We now show that the populations converge from almost any initial state to an interior stationary state if (and essentially only if) the product of the bounded expectation of the distribution of sample sizes in each population and the share of agents with sample size 1 in the other population is larger than 1.

Theorem 1. Assume that $u_2 < 1 \le u_1$ (no $\frac{1}{2}$ -dominant equilibrium). Then

1. Global convergence to miscoordination: Assume that

 $\theta_1(1) \cdot \mathbb{E}_{<\frac{1}{u_2}+1}(\theta_2) > 1 \quad and \quad \theta_2(1) \cdot \mathbb{E}_{\le u_1+1}(\theta_2) > 1.$

If $p(0) \notin \{(0,0), (1,1)\}$, then $\lim_{t\to\infty} \mathbf{p}(t) \notin \{(0,0), (1,1)\}$.

2. Local convergence to coordination: Assume that

$$\theta_1(1) \cdot \mathbb{E}_{<\frac{1}{u_2}+1}(\theta_2) < 1 \text{ or } \theta_2(1) \cdot \mathbb{E}_{\le u_1+1}(\theta_2) < 1.$$

Then at least one of the pure equilibria is asymptotically stable.

Proof.

- Proposition 1 implies that both pure stationary states are unstable. Proposition 2 and Lemma 1 imply that from almost any initial state, the population converges to an interior stationary state.
- Proposition 1 implies that at least one of the pure equilibria is asymptotically stable, which implies that some interior initial states converge to a pure equilibrium.

Theorem 1 shows that global convergence to miscoordination requires heterogeneity in the sample sizes in each population that includes both agents with a small sample size of one, and agents with larger samples (but not too large, as they must be below the bound for which a single observation of a rare action can influence behavior). Specifically, in each population it is required that the product of (1) the share of agents with a sample size of 1 and (2) the bounded expected sample size should be sufficiently large.

Observe that the farther the u_i -s are from 1, the higher (i.e., less restrictive) the bounded expected value is. That is, games in which for each population one of the actions is much riskier than the other action (i.e., for population 2, action a_2 is much riskier than action a_1 , while for population 1, action a_1 is much riskier than action a_2) are more likely have stable miscoordination.

The following example demonstrates global convergence to miscoordination, and the fact that the stability of pure states is nonmonotone in the sample sizes.

Example 1 (Nonmonotone Impact of Sample Size). Consider a battle of the sexes game in which $u_1 = 5$ and $u_2 = 0.2$ and both populations have the same distribution of sample sizes. Consider 3 distributions of sample sizes, in each of which half of the players in the population have sample size 1. In the first distribution (left panel of Figure 4.2) the remaining half of the players have sample size 3, in the second distribution (middle panel) they have sample size 5, and in the third distribution (right panel of Figure 4.2) they have sample size 7. Observe that the second distribution satisfies the condition for global



The figure illustrates the phase plots for three environments. In each environment, the underlying coordination game is antisymmetric with $u_1 = 5$ and $u_2 = 0.2$, and 50% of the agents in each population have sample size 1. In the environment illustrated in the left (resp., middle, right) panel the remaining half of the players in the population have sample size 2 (resp., 5, 7). The middle panel shows global convergence to the interior state with miscoordination (0.37, .0.63), while the other two panels show global convergence to one of the pure equilibria.

convergence to miscoordination, i.e.,

$$\theta_1(1) \cdot \mathbb{E}_{<\frac{1}{u_2}+1}(\theta_2) = \theta_1(1) \cdot \mathbb{E}_{<6}(\theta_2) = \theta_1(1) \cdot \mathbb{E}(\theta_2) = 0.5 \cdot 3 > 1$$

(and the same holds for population 2). Indeed, the middle phase plot shows that the populations converge from any interior state to the state (0.37, 0.63) with substantial miscoordination (specifically, the agents miscoordinate and get a payoff of zero in 46% $\approx 2 \cdot 0.37 \cdot 0.63$ of the interactions). By contrast, either decreasing the larger sample size from 5 to 2, or increasing it to 7, yields a product $\theta_1(1) \cdot \mathbb{E}_{<\frac{1}{u_2}+1}(\theta_2)$ that is strictly smaller than 1 ($\theta_1(1) \cdot \mathbb{E}(\theta_2) = 0.5 \cdot 1.5 < 1$ in the first case, and $\theta_1(1) \cdot \mathbb{E}_{<6}(\theta_2) = \theta_1(1) \cdot \theta_1(1) = 0.5 \cdot 0.5 < 1$ in the second case). The left and right panels of Figure 4.2 illustrate that in both cases almost all initial states converge to a pure state. Thus, changing sample sizes of agents has a nonmonotone impact on the stability of miscoordination.

5 Homogeneity and Unstable Miscoordination

In this section, we show that heterogeneity is necessary for stable miscoordination. Specifically, we show that any environment in which all agents in each population have the same sample size admits at most one interior stationary state, and that this state is unstable. Auxiliary Results We begin by showing that a stationary state $\hat{\boldsymbol{p}}$ is asymptotically stable iff the curve w_2 is above the curve w_1^{-1} in a left neighborhood of $\hat{\boldsymbol{p}}$, and it is below the curve w_1^{-1} in a right neighborhood of $\hat{\boldsymbol{p}}$.

Lemma 2. Let \hat{p} be a stationary state. \hat{p} is asymptotically stable if both of the following conditions hold:

- 1. Left neighborhood: If $\hat{p}_1 > 0$ then there exists $\underline{p}_1 \in (0, \hat{p}_1)$ such that $w_2(p_1) > w_1^{-1}(p_1)$ for any $p_1 \in (\underline{p}_1, \hat{p}_1)$, and
- 2. **Right neighborhood**: If $\hat{p}_1 < 1$, then there exists $\overline{p}_1 \in (\hat{p}_1, 1)$ such that $w_2(p_1) < w_1^{-1}(p_1)$ for any $p_1 \in (\hat{p}_1, \overline{p}_1)$.

Moreover, if either of the above two conditions is not satisfied, then \hat{p} is unstable.

Proof. See Appendix A.8.

Lemma 2 implies that the neighbor of an asymptotically stable state must be unstable.

Corollary 2. Let $\underline{p} \neq \overline{p}$ be two neighboring stationary states. If \underline{p} is asymptotically stable, then \overline{p} is unstable.

Proof. Without loss of generality assume that $\underline{p}_1 < \overline{p}_1$. Due to Claim 2, the fact that \underline{p} is asymptotically stable implies that $w_2(p_1) > w_1^{-1}(p_1)$ for any $p_1 \in (\underline{p}_1, \overline{p}_1)$ which, in turn implies that \overline{p} is unstable.

A Property of Binomial Distributions The main result is implied by a property of binomial distributions (which may be of independent interest). Recall our notation of $X(k,p) \sim Bin(k,p)$ denoting a random variable with a binomial distribution. Define $F_m^k(p) \equiv \Pr(X(k,p) \ge m)$ as the probability of having at least m successes in k trials when the probability of success in each trial is p. Observe that $F_m^k(0) = 0$, $F_m^k(1) = 1$, and $(F_m^k)' > 0$. It is known that F_m^k has at most one interior fixed point, i.e.,

Fact 3 (Green, 1983, Theorem 1). Fix arbitrary integers $0 < m \le k$. Then there is at most one $p \in (0,1)$ such that $F_m^k(p) = p$.

We extend it by showing that the same is true also for a composition of any two cumulative binomial distributions, i.e., that $F_{m_1}^{k_1} \circ F_{m_2}^{k_2}$ has at most one interior fixed point.

Proposition 3. Fix arbitrary integers satisfying $0 < m_1 \le k_1$ and $0 < m_2 \le k_2$. There exists at most one $p \in (0, 1)$ such that $\left(F_{m_1}^{k_1} \circ F_{m_2}^{k_2}\right)(p) = p$.

The proof, which is detailed in Appendix A.3, shows that $F_{m_1}^{k_1} \circ F_{m_2}^{k_2}$ has at most one inflection point, which implies that there exists at most one interior fixed point.

Main Result We now present the main result of this section: any environment with homogeneous sample sizes admits at most one interior stationary state that is unstable. This implies that almost all initial states converge to one of the pure equilibria.

Theorem 2. Assume that $\theta_i \equiv k_i > 1$ for each $i \in \{1, 2\}$. There exists at most one interior stationary state, and this state (if it exists) is unstable.

Proof. The fact that $\theta_i \equiv k_i > 1$ implies that $w_i(p) = F_{m_i}^{k_i}(p)$ for some $1 \leq m_i \leq k_i$. This implies that any stationary state \boldsymbol{p}^* must satisfy $\left(F_{m_1}^{k_1} \circ F_{m_2}^{k_2}\right)(p_1^*) = p_1^*$. Proposition 3 implies that this holds for at most one interior state $\hat{\boldsymbol{p}}$. Therefore, the stationary state $\hat{\boldsymbol{p}}$ (if it exists) is a neighbor of both pure stationary states (0,0) and (1,1). Proposition 1 and the fact that no agents have sample size of 1 implies that at least one of these pure states is asymptotically stable. Finally, Corollary 2 implies that $\hat{\boldsymbol{p}}$ is unstable.

6 Heterogeneity and Stable Miscoordination

The conditions presented for global convergence to miscoordination in Section 4 are somewhat narrow in the sense of requiring sufficiently many agents with sample size one. In this section we show that a much broader set of heterogeneous distributions of sample sizes can induce asymptotically stable states with miscoordination.

Specifically, we show that essentially any distribution of sample sizes that combines agents with small samples and agents with large samples induces locally stable miscoordination if the u_i -s are not too close to one. As demonstrated below, this type of heterogeneity in the sample sizes is plausible in various setups.

Example 2. In housing markets (in which the bargaining situations can be modeled as hawk-dove coordination games), it is often the case that each population includes two types of agents: (1) professional real-estate investors, and (2) people who buy/sell houses only a couple of times during their life. It seems plausible that the real-estate investors

have reliable information on the aggregate behavior (captured in our model by having large samples), while the remaining agents are likely to have limited, anecdotal evidence about the aggregate behavior (modeled by having small samples).

Result Recall (Theorem 2) that all interior equilibria are unstable if all agents in each population have the same size. Our final result shows that, perhaps surprisingly, one can always add agents with large samples, and obtain an asymptotically stable interior equilibrium with miscoordination, provided that the u_i -s are not too close to one.

We formally define adding agents with large samples as follows. Given two sample sizes k, \bar{k} and $\alpha \in (0, 1)$, let $k\alpha \bar{k}$ be the distribution of sample sizes that assigns mass α to sample size k and mass $1 - \alpha$ to sample size \bar{k} .

Theorem 3. Fix any pair of sample sizes $k_1, k_2 > 1$. Then there exist $r, \alpha_1, \alpha_2 \in (0, 1)$, such that for any $u_1, u_2 \notin (r, \frac{1}{r})$ the environment $\left(\boldsymbol{u}, \left(k_1\alpha_1\bar{k}, k_2\alpha_2\bar{k}\right)\right)$ admits an asymptotically stable interior state for any sufficiently large \bar{k} .

Proof idea (see Appendix A.9 for a formal proof). We present one argument for games with a $\frac{1}{2}$ -dominant equilibrium, and another for those without. The arguments are illustrated in the phase plots of Figure 6.1.

- 1. Games with a $\frac{1}{2}$ -dominant equilibrium: Assume that $u_i > k_i$ for each *i*. Proposition 1 implies that the risk-dominated equilibrium **b** is unstable when all agents have sample size k_i . Observe that adding agents with large samples decreases the average probability of new agents playing a_i in state (ϵ_1, ϵ_2) near **b**. We choose α_i , such that the share of new agents playing action a_i (1) is above ϵ_j in state (ϵ_1, ϵ_2) , and (2) is below p_j in an interior state $\hat{\boldsymbol{p}}$ (this is possible because $\frac{w_i(p_j)}{p_j}$ is decreasing in small p_j -s). This, in turn, implies that there is a stable interior state between **b** and $\hat{\boldsymbol{p}}$.
- 2. Games without a $\frac{1}{2}$ -dominant equilibrium: Observe that $w_i^{k_i}(p_j)$ remains the same for all values of u_i that are sufficiently far from one $(u_1 > k_1, u_2 < \frac{1}{k_2})$, while the interior Nash equilibrium p_1^{NE} (resp., p_2^{NE}) converges to 0 (resp., 1) as u_1 converges to infinity (resp., as u_2 converges to zero). Thus, for u_i -s sufficiently far from 1, $p_2^{NE} < \frac{w_2^{\theta_2}(\frac{1}{2})}{2}$. We show that this implies that there must be an unstable stationary state with $p_1 < \frac{1}{2}$ and $p_2 \simeq p_2^{NE} < \frac{1}{2}$. By an analogous argument, there is an unstable stationary state with $p_1 \simeq p_1^{NE} > \frac{1}{2}$ and $p_2 > \frac{1}{2}$. This implies that

there must be an asymptotically stable interior state between these two unstable states. $\hfill \square$



Figure 6.1: Illustrative Phase Plots for Theorem 3 (Locally Stable Miscoordination)

The left panel illustrates the symmetric game $u_1 = u_2 = 2.5$, where 40% (resp., 60%) of the agents in each population have sample size 3 (resp., 1,000). The pure risk-dominated equilibrium (0,0) is unstable, and its neighbor the interior stationary state with miscoordination (0.18, 0.18) is asymptotically stable. The right panel illustrates the antisymmetric game $u_1 = \frac{1}{u_2} = 5$, where 50% (resp., 50%) of the agents have sample size 2 (resp., 1,000). The mixed Nash equilibrium of this game is $(\frac{5}{6}, \frac{1}{6})$. The environment admits three interior stationary states: two unstable interior stationary states where one of the coordinates is close to the Nash equilibrium: (0.6, 0.18) and (0.82, 0.3), and a stable interior stationary state at (0.73, 0.27).

Remark 3. For simplicity, Theorem 3 deals with a heterogeneous population that combines two sample sizes, k_i and \bar{k} . The proof's argument can be adapted in a straightforward way to populations combining some agents who have arbitrary sample sizes (i.e., an arbitrary distribution θ_i replaces the arbitrary k_i), and some agents who have sufficiently large sample sizes (i.e., not necessarily all of the agents with large samples have the same sample size of \bar{k}).

Non-Payoff-Maximizing Behavior of Agents with Small Samples Observe the interior stable state is (generically) not a mixed Nash equilibrium. For example, in the environment illustrated in the left panel of Figure 3 in which $u_i = 2.5$, $\theta_i(3) = 40\%$ and $\theta_i(1,000) = 60\%$ for each population *i*, the interior stable state is (0.17, 0.17), while the mixed Nash equilibrium is (0.28, 0.28). This implies that one of the actions is the unique true best response to the opponent's behavior (action b_i in the example). This unique best response is played essentially by all agents with a large sample size of 1,000. By contrast, about 45% of the agents with a small sample size of 3 observe action a_{-i} at least once in their sample, and play action a_i due to an erroneous belief (based on their small sample) that the frequency of a_{-i} is at least $\frac{1}{3}$. As a result, the average payoff to agents with a small sample is only 0.75, which is 8% smaller then the payoff of 0.82 to agents with a large sample.

This raises the question of why the agents with small samples do not learn to increase their samples, or why they are not driven out of the population. We suggest two explanations for this. The first is setups like our motivating example of a housing market that combine both inexperienced players (who have little experience in buying/selling houses) and experienced players (real-estate investors). In such setups, it seems likely that the inexperienced agents might have very limited (or costly) access to large samples (and a continuous inflow of inexperienced traders prevents them from being driven out of the market). Another reason why some players do not exert effort to increase their sample sizes is the "the law of small numbers," which is the commonly believed bias that small samples well represent large populations (Tversky and Kahneman, 1971).

7 Comparison with Logit Dynamics

The main candidate in the existing literature to induce stable miscoordination in coordination games is logit dynamics. In this section we numerically demonstrate that (1) standard logit dynamics with a homogeneous level of noise in each population can induce stable miscoordination only with high levels of noise that seem implausible, and (2) a variant of logit dynamics with heterogeneity in the noise level can induce stable miscoordination with substantially lower levels of noise. This suggests that our key insight that stable miscoordination is induced by heterogeneous noise might hold in various classes of dynamics, and not only in sampling dynamics.¹¹

Standard (Homogeneous) Logit Dynamics Logit dynamics (introduced in Fudenberg and Levine, 1995; see Sandholm, 2010, Section 6.2.3 for a textbook exposition) are characterized by a single parameter η_i that describes the *noise level* for each population

¹¹Preliminary calculations suggest that this insight holds also under best experienced payoff dynamics (see, e.g., Sethi, 2000) and under dynamics induced by representative sampling à la Danenberg and Spiegler (2022).

i. If player *i* plays action a_i , she will get a payoff of $p_j \cdot u_j$. If she plays action b_i she will get a payoff of $(1 - p_j) \cdot 1$. Logit dynamics assume that the probability of revising agents playing action a_i is proportional to $e^{\frac{\text{Payoff of } a_i}{\eta}}$. Specifically, logit dynamics are given by

$$w_i(p_j) \equiv w_i(\mathbf{p}) = \frac{e^{\frac{p_j \cdot u_j}{\eta_i}}}{e^{\frac{1-p_j}{\eta_i}} + e^{\frac{p_j \cdot u_j}{\eta_i}}}.$$
(7.1)

Trivially, logit dynamics can induce substantial miscoordination by having high values of noise. The interesting question is whether stable miscoordination can be supported by a low level of noise. Our numerical analysis suggests that the answer is negative. In what follows we demonstrate that this is indeed the case. For example, when we revisit the two examples of Figure 6.1 $(u_1 = u_2 = 2.5 \text{ and } u_1 = \frac{1}{u_2} = 5)$, then the minimal level of noise that is required to sustain an asymptotically stable interior state in which each action is played with a probability of at least 10% is $\eta = 1$ (see the left panel of Figure 7.1 for an illustration of the case of $u_1 = u_2 = 2.5$). Such a high level of noise implies that 27% of the revising agents make the obvious mistake of playing a_i when facing a population in which almost everyone plays b_i ; by contrast, this obvious mistake is never made under action-sampling dynamics. Moreover, the average expected payoff obtained by revising agents who follow logit dynamics against an opponent population in which the share of agents playing action a_i is distributed uniformly is 85% (resp., 71%) of the maximal payoff that can be obtained by payoff-maximizing agents in the first (resp., second) environment with $u_1 = u_2 = 2.5$ (resp., $u_1 = \frac{1}{u_2} = 5$). By contrast, this average expected payoff is 98% (resp., 95%) of the maximal payoff under the sampling dynamics. Thus, stable cooperation can be supported by standard (homogeneous) logit dynamics only when the agents have high levels of noise.

Heterogeneous Logit Dynamics Next, consider a variant of logit dynamics in which there is heterogeneity in the level of noise for agents in each population. Specifically, in a population in which there are n groups, the size of the l-th group is μ_i^l and its members have a noise level of η_i^l , the heterogeneous logit dynamics are given by

$$w_{i}(p_{j}) \equiv w_{i}(\mathbf{p}) = \sum_{l} \mu_{i}^{l} \cdot \frac{e^{\frac{p_{j} \cdot u_{j}}{\eta_{i}^{l}}}}{e^{\frac{1-p_{j}}{\eta_{i}^{l}}} + e^{\frac{p_{j} \cdot u_{j}}{\eta_{i}^{l}}}}.$$
(7.2)



Figure 7.1: Supporting Stable Coordination with Heterogeneous Logit Dynamics

The figure revisits the symmetric game presented in the left panel of Figure 6.1 in which $u_1 = u_2 = 2.5$. The left panel shows the phase plot of the minimal homogeneous level of noise, $\eta_i = 1$, that sustains an asymptotically stable state in which each action is played with a probability of at least 10%. The right panel shows the phase plot of a heterogeneous variant of logit dynamics in which 55% of the the agents in each population have a moderate level of noise $\eta_i = 0.55$ and 45% have a small level of noise $\eta_i = 0.01$.

The numerical calculations demonstrate that heterogeneous noise levels can induce asymptotically stable miscoordination with relatively low levels of noise. Specifically, in both of the above examples ($u_1 = u_2 = 2.5$, which is illustrated in the right panel of Figure 7.1, and $u_1 = \frac{1}{u_2} = 5$), populations in which 55% of the agents have a moderate level of noise (i.e., $\eta = 0.55$) and 45% have a small level of noise (i.e., $\eta = 0.01$) induce asymptotically stable states with miscoordination ((0.21, 0.21) in the right panel of Figure 7.1). Given these heterogeneous levels of noise, only 8% of the agents make the mistake of playing action a_i when facing a population in which everyone plays a_j , and the average expected payoff obtained by playing against opponent populations in which the share of agents playing action a_i is distributed uniformly is 96% (resp., 89%) of the maximal payoff that can be obtained by payoff-maximizing revising agents in the environment with $u_1 = u_2 = 2.5$ (resp., $u_1 = \frac{1}{u_2} = 5$).

8 Conclusion

The conventional wisdom, which is supported by key results in evolutionary game theory, is that only pure (coordinated) outcomes are reasonable long-run predictions of behavior in coordination games. By contrast, we show that plausible learning dynamics, in which new agents rely on samples to estimate and best respond to the behavior of the opponents' population, can induce stable miscoordination. This happens if there is heterogeneity in the sample sizes: some agents have accurate information based on large samples of the opponents' aggregate behavior, while other agents rely on anecdotal evidence induced by small samples. We further show that stable miscoordination holds under a broad set of heterogeneous distributions of sample sizes, if for each player one of the actions in the underlying game is q-dominant for a sufficiently small q.

Although our analytical results focus on the specific family of sampling dynamics, the numerical results for logit dynamics suggest that qualitatively similar results are likely to hold under other learning dynamics, i.e., heterogeneous noise levels can induce stable miscoordination. Such heterogeneity is plausible in many applications, such as bargaining in housing markets, where some participants are professional investors, while other participants are inexperienced. The predictions of our model can be experimentally tested, as outlined Section 2.

A Appendix

A.1 Standard Representation of General Coordination Games

In this subsection we explain why the two-parameter standard representation (Table 2) captures w.l.o.g. any two-action coordination game (in line with Harsanyi and Selten's (Harsanyi and Selten) Axiom 2 of best-reply invariance).

The most general definition of a two-action coordination game is a game that admits two strict Nash equilibria. By relabeling the actions of player 1, we can assume w.l.o.g. that these two pure equilibria are (a_1, a_2) and (b_1, b_2) (i.e., if the two pure equilibria are (a_1, b_2) and (b_1, a_2) , then we switch the labels of player 1's actions: $a_1 \leftrightarrow b_1$). This implies that the left panel of Table 3 shows a parametric representation of all coordination games.

Sampling dynamics (as defined in 3.1 and A.1) depend only on the differences between the payoffs a player can get by playing different actions (the same property holds for best-response dynamics and logit dynamics, which implies that the sets of Nash equilibria, quantal response equilibria, and evolutionary stable strategies depend only on these

Original Representation		Standard Representation				
	a_2	b_2			a_2	b_2
a_1	u_{11} v_{11}	$u_{12}^{\ \ v_{12}}$	\Rightarrow	a_1	$\frac{u_{11} - u_{21}}{u_{22} - u_{12}} \stackrel{\frac{v_{11} - v_{12}}{v_{22} - v_{21}}}{=}$	0 0
b_1	$u_{21}^{v_{21}}$	$u_{22}^{\ \ v_{22}}$		b_1	0 0	1
$u_{11}>u_{21}, u_{22}>u_{12}, v_{11}>v_{12}, v_{22}>v_{21}$						

Table 3: Normalization of General Two-Action Coordination GamesOriginal RepresentationStandard Representation

differences).¹² These differences are invariant to subtracting a constant from all the payoffs of a player while fixing the opponent's action (e.g., subtracting u_{21} from all of Player 1's first-column payoffs). Moreover, sampling dynamics (as well as all the other dynamics and solution concepts mentioned above) are invariant to dividing all of a player's payoff by a positive constant (which preserves the vN–M utility). The left matrix in Table 3 is reduced to the right matrix by the following steps (none of which affect the sampling dynamics):

- Three changes to Player 1's payoffs: (I) subtracting u₂₁ from Player 1's payoffs in her first column, (II) subtracting u₁₂ from Player 1's payoffs in her second column, and (III) dividing Player 1's payoff by u₂₂-u₁₂; and
- 2. Three changes to Player 2's payoffs: (I) subtracting v_{12} from Player 1's payoffs in her first row, (II) subtracting v_{21} from Player 1's payoffs in her first row, and (III) dividing Player 1's payoff by $v_{22} v_{21}$.

Observe that the assumption that $u_1 = \frac{u_{11}-u_{21}}{u_{22}-u_{12}} \ge 1$ in the standard representation of Table 2 is w.l.o.g.. If $\frac{u_{11}-u_{21}}{u_{22}-u_{12}} < 1$, then we can multiply all of Player 1's payoffs by $\frac{u_{22}-u_{12}}{u_{11}-u_{21}}$ and all of Player 2's payoffs by $\frac{v_{22}-v_{21}}{v_{11}-v_{12}}$, relabel the actions $a_i \leftrightarrow b_i$ for both players, and obtain a standard representation as in Table 2 in which $u_1 \ge 1$.

Example 3 (Hawk-Dove Games). Consider a hawk-dove (aka, chicken) game, which can be interpreted as a game of bargaining over the price of an asset (e.g., a house) between a buyer and a seller. Each player can either insist on a more favorable price ("hawk") or agree to a less favorable price in order to close the deal ("dove"). The left panel of Table

¹²A notable exception is best experienced payoff dynamics (see, e.g., Osborne and Rubinstein, 1998; Cárdenas, Mantilla, and Sethi, 2015; Mantilla, Sethi, and Cárdenas, 2018; Sandholm, Izquierdo, and Izquierdo, 2019, 2020; Sethi, 2021; Arigapudi, Heller, and Milchtaich, 2021), where the dynamics depend directly on the payoffs, and not only on payoff differences.

Original Representation			Standard Representation			tion
	h_2	d_2			$a_2 = h_2$	$b_2 = d_2$
h_1	0	1+g $1-l$	\Rightarrow	$a_1 = d_1$	$\frac{1-l}{g}$ $\frac{g}{1-l}$	0
d_1	1- <i>l</i> 1+ <i>g</i>	1		$b_1 = h_1$	0 0	1

Table 4: Normalization of Hawk-Dove Games $(g, l \in (0, 1))$

4 shows the payoffs of a hawk-dove game. Two doves agree on an equally favorable price (which gives both players a relatively high payoff normalized to one). A hawk obtains a favorable price when matched with a dove (which increases the payoff to the hawk by $g \in (0, 1)$, while reducing the dove's payoff by $l \in (0, 1)$), but faces a high probability of bargaining failure when matched with another hawk (which would yield a low payoff of zero to both hawks).

Observe that a hawk-dove game can be transformed to our standard representation of a coordination game (the right panel of Table 4) as follows: (1) relabel the actions of player 1 such that $a_1 = d_1$ and $b_1 = h_1$ (while keeping the actions of Player 2 as $a_2 = h_2$ and $b_2 = d_2$), (2) subtract a payoff of 1 from Player 1's payoffs in her second column and from Player 2's payoffs in her first column, and (3) divide all the payoffs of Player 1 by g, and all payoffs of Player 2 by 1 - l. Observe that the induced standard representation is antisymmetric, i.e., $u_1 = \frac{1-l}{g} = \frac{1}{u_2}$.

A.2 One-Population Dynamics

In the main text, we derived results for general coordination games under two-population dynamics. In this appendix, we show that analogous results hold for symmetric coordination games under one-population dynamics.

Adapted Definitions An environment $((u_1, u_2), (\theta_1, \theta_2))$ is symmetric if $u_1 = u_2$ and $\theta_1 = \theta_2$. As we focus on symmetric environments here, we omit the subscript *i* referring to the population from the parameters u_i, θ_i , the actions a_i, b_i , and the sampling dynamics w_i . Thus, a symmetric environment is denoted by (u, θ) , where *u* is the payoff to each player when coordinating on *a*, and θ is the distribution of sample size of each player.

Under one-population dynamics, the state is denoted by a single number $p \in (0, 1)$ that describes the share of players who play action a in the population. At each time t, a constant flow (normalized to one) of agents die, and are replaced by new agents. The sample sizes of the new agents are distributed according to θ . Each k-agent (a new agent with sample size k) randomly and independently samples k players and observes the actions that they played. Thus, the one-dimensional sampling dynamics is given by $\dot{p} = w(p) - p$, where

$$w(p) = \sum_{k \in \text{supp}(\theta)} \theta(k) \cdot \Pr\left(X(k, p) \ge \frac{k}{u+1}\right).$$
(A.1)

In the above equation, X(k, p) denotes a binomial random variable with parameters k and p. Analogous arguments to those presented in Appendix A.1 imply that any symmetric coordination game can be normalized w.l.o.g. to the standard form with a single parameter u.¹³

Adaptation of Results We next show that in symmetric coordination games, the set of stationary states and the set of asymptotically stable stationary states coincide under the one-population dynamics and the two-population dynamics. This implies that all our main-text results hold also for one-population dynamics in symmetric coordination games.

We first show that the sets of stationary states coincide in both sets of dynamics.

Proposition 4. Let (u, θ) be a symmetric environment. State (p_1, p_2) is stationary under two-population dynamics iff the state is symmetric (i.e., $p \equiv p_1 = p_2$), and p is stationary under one-population dynamics.

Proof. Assume that (p_1, p_2) is a stationary state under two-population dynamics. Assume to the contrary that $p_1 \neq p_2$. w.l.o.g., assume that $p_1 < p_2$. The fact that w is strictly increasing implies that $p_2 = w(p_1) < w(p_2) = p_1$, and we get a contradiction. Thus, we know that all stationary states under two-population dynamics are symmetric (i.e., of the form (p, p) in symmetric environments). Observe that a symmetric state (p, p)is stationary under two-population dynamics iff p is stationary under one-population

¹³Under one-population dynamics it is important that the strict equilibria are obtained by both players playing the same action (i.e., being on the main diagonal of the payoff matrix), as one cannot relabel the actions of only one of the players in the game without breaking the symmetry. Thus, under one-population dynamics, hawk-dove games (in which the pure equilibria are off the main diagonal) are no longer coordination games (by contrast, they are asymmetric coordination games under two-population dynamics). As demonstrated in Herold and Kuzmics (2020), introducing pre-play cheap-talk can induce some hawk-dove games to essentially be coordination games also in one-population dynamics.



Figure A.1: Illustration of Equivalence Results for One- and Two-Population Dynamics

The figure illustrates the phase plots of three symmetric environments for two-population dynamics (upper panels) and the corresponding one-population dynamics (lower panels). Notice that the phase plots of one-population dynamics are similar to the part of phase plots of the corresponding two-population dynamics along the main diagonal (marked with red arrows). In all environments u = 2.4. In the left panels all agents have sample size k = 3, in the middle panels all agents have sample size 7, and in the right panels 40% of the agents have sample size 3, and 60% have sample size 1000. Observe that in all cases, the dynamic predictions under both one-population dynamics and two-population dynamics are the same.

dynamics, because stationarity under both dynamics is characterized by the equation p = w(p).

The fact that one-population dynamics are one-dimensional implies the following simple and well-known fact:

Fact 4. Stationary state *p* is asymptotically stable under one-population dynamics iff both of the following conditions hold:

- 1. If p > 0, then w(p) > p in a left neighborhood of p.
- 2. If p < 1, then w(p) < p in a right neighborhood of p.

We conclude this section by showing that the sets of asymptotically stable states coincide under both one-population dynamics and two-population dynamics. Fact 4 and Propositions 4–5 are illustrated in Figure A.1.

Proposition 5. Let (u, θ) be a symmetric environment. State (p_1, p_2) is asymptotically stable under two-population dynamics iff the state is symmetric (i.e., $p \equiv p_1 = p_2$), and p is asymptotically stable under one-population dynamics.

Proof. From Proposition 4, it follows that all asymptotically stable states under twopopulation dynamics are symmetric. The strict monotonicity of w implies that state psatisfies conditions (1) and (2) in Fact 4 iff the symmetric state (p, p) satisfies conditions (1) and (2) in Lemma 2. This, in turn, implies that state p is asymptotically stable under one-population dynamics iff the symmetric state (p, p) is asymptotically stable under two-population dynamics.

A.3 General Result for Binomial Distributions

Recall our notation of $X(k,p) \sim Bin(k,p)$ denoting a random variable with binomial distribution, and of the function $F_m^k(p) \equiv \Pr(X(k,p) \ge m)$.

Proposition. 3 Fix arbitrary integers satisfying $0 < m_1 \le k_1$ and $0 < m_2 \le k_2$. There exists at most one $p \in (0, 1)$ such that $\left(F_{m_1}^{k_1} \circ F_{m_2}^{k_2}\right)(p) = p$.

Proof. Let $w_i(p) \equiv F_{m_i}^{k_i}(p)$ for each $i \in \{1, 2\}$, $F(p) \equiv \left(F_{m_1}^{k_1} \circ F_{m_2}^{k_2}\right)(p) \equiv (w_1 \circ w_2)(p)$, and G(p) = F(p) - p. In what follows we show that the function $G(\cdot)$ has at most one $p \in (0, 1)$ such that G(p) = 0, which proves the result.

We have G(0) = G(1) = 0. Assume to the contrary that there exist two different interior points $0 < \underline{p} < \overline{p} < 1$ such that $G(\underline{p}) = G(\overline{p})$. Then G equals zero at four points in the interval [0, 1]. By Rolle's theorem, this implies that G' is equal to zero in at least three points in the interval (0, 1), which further implies that G'' is equal to zero in at least two interior points in the interval (0, 1). Observe that $G'' \equiv F''$. Thus, in order to obtain a contradiction we have to show that F''(p) = 0 in at most one interior point. Recall that (see, e.g., Green, 1983, Eq. (5)):

$$w_{i}'(p) = m_{i} \begin{pmatrix} k_{i} \\ m_{i} \end{pmatrix} p^{m_{i}-1} (1-p)^{k_{i}-m_{i}} \Rightarrow \frac{w_{i}''(p)}{w_{i}'(p)} = \frac{m_{i}-1}{p} - \frac{k_{i}-m_{i}}{1-p}.$$
 (A.2)

For $p \in (0, 1)$, using Eq. (A.2), we compute as follows:

$$F'(p) = w'_1(w_2(p)) w'_2(p)$$

$$F''(p) = w''_1(w_2(p)) (w'_2(p))^2 + w'_1(w_2(p)) w''_2(p)$$

$$= w'_1(w_2(p)) w'_2(p) \left[\left(\frac{m_1 - 1}{w_2(p)} - \frac{k_1 - m_1}{1 - w_2(p)} \right) w'_2(p) + \frac{m_2 - 1}{p} - \frac{k_2 - m_2}{1 - p} \right]$$

The fact that each $w_i(p)$ is strictly increasing implies that F''(p) = 0 iff

$$\begin{pmatrix} \frac{m_1 - 1}{w_2(p)} - \frac{k_1 - m_1}{1 - w_2(p)} \end{pmatrix} w_2'(p) = \frac{k_2 - m_2}{1 - p} - \frac{m_2 - 1}{p} \Leftrightarrow$$

$$m_2 \begin{pmatrix} k_2 \\ m_2 \end{pmatrix} \left(\frac{m_1 - 1}{w_2(p)} - \frac{k_1 - m_1}{1 - w_2(p)} \right) p^{m_2 - 1} (1 - p)^{k_2 - m_2} = \frac{k_2 - m_2}{1 - p} - \frac{m_2 - 1}{p} \Leftrightarrow$$

$$m_2 \begin{pmatrix} k_2 \\ m_2 \end{pmatrix} \left(\frac{(m_1 - 1)p^{m_2}(1 - p)^{k_2 - m_2}}{\sum_{l=m_2}^{k_2} \binom{k_2}{l} p^l (1 - p)^{k_2 - l}} - \frac{(k_1 - m_1)p^{m_2}(1 - p)^{k_2 - m_2}}{\sum_{l=0}^{m_2 - 1} \binom{k_2}{l} p^l (1 - p)^{k_2 - l}} \right) \frac{1}{p} = \frac{k_2 - m_2}{1 - p} - \frac{m_2 - 1}{p} \Leftrightarrow$$

$$m_2 \begin{pmatrix} k_2 \\ m_2 \end{pmatrix} \left(\frac{(m_1 - 1)}{\sum_{l=m_2}^{k_2} \binom{k_2}{l} \binom{p}{1 - p}^{l - m_2}} - \frac{k_1 - m_1}{\sum_{l=0}^{m_2 - 1} \binom{k_2}{l} \binom{1 - p}{p}^{m_2 - l}} \right) \frac{1}{p} = \frac{k_2 - m_2}{1 - p} - \frac{m_2 - 1}{p} \Leftrightarrow$$

One could verify that the left-hand side of the above equation is strictly decreasing and the right-hand side is strictly increasing in p. Therefore, there can be at most one point $p^* \in (0, 1)$ where $F''(p^*) = 0$. This completes the proof.

A.4 Standard Definitions of Dynamic Stability

For completeness, we present in this appendix the standard definitions of dynamic stability that are used in the paper (see, e.g., Weibull, 1997, Chapter 5).

A state is said to be stationary if it is a rest point of the dynamics.

Definition 4. State $\mathbf{p}^* \in [0,1]^2$ is a stationary state if $w_i(\mathbf{p}^*) = p_i^*$ for each $i \in \{1,2\}$.

Let $\mathcal{E}(w)$ denote the set of stationary states of w, i.e., $\mathcal{E}(w) = \{\mathbf{p}^* | w_i(\mathbf{p}^*) = p_i^*\}$. Under monotone dynamics, an interior (mixed) state $\mathbf{p}^* \in (0, 1)^2$ is a stationary state iff it is a Nash equilibrium (Weibull, 1997, Prop. 4.7). By contrast, under nonmonotone dynamics (such as the sampling dynamics analyzed below) the two notions differ.

A state \mathbf{p}^* is Lyapunov stable if a population beginning near \mathbf{p}^* stays close to it. \mathbf{p}^* is asymptotically stable if it is Lyapunov stable and, in addition, nearby populations eventually converge to it. A state is unstable if it is not Lyapunov stable. It is well known (see, e.g., Weibull, 1997, Section 6.4) that every Lyapunov stable state must be a stationary state. Formally:

Definition 5. A stationary state $\mathbf{p}^* \in [0, 1]^2$ is *Lyapunov stable* if for every neighborhood U of \mathbf{p}^* there is a neighborhood $V \subseteq U$ of \mathbf{p}^* such that if the initial state $p(0) \in V$, then $\mathbf{p}(t) \in U$ for all t > 0. A state is *unstable* if it is not Lyapunov stable.

Definition 6. A stationary state $\mathbf{p}^* \in [0, 1]^2$ is asymptotically stable (or locally stable) if it is Lyapunov stable and there is some neighborhood U of \mathbf{p}^* such that all trajectories initially in U converge to \mathbf{p}^* , i.e., $\mathbf{p}(0) \in U$ implies $\lim_{t\to\infty} \mathbf{p}(t) = \mathbf{p}^*$.

A.5 Proof of Proposition 1 (Stability of Pure States)

We are interested in deriving conditions for the stability of pure stationary states. In what follows, we compute the Jacobian of sampling dynamics in the pure state $\boldsymbol{a} = (0,0)$ (resp., $\boldsymbol{b} = (1,1)$). For this, we consider a slightly perturbed state with a "very small" ϵ_i share of agents playing b_i (resp., a_i) in population *i*. By "very small," we mean that higher-order terms of ϵ_i and ϵ_j are neglected.

Consider a new agent in population i who has a sample size of k_i . Action b_i (resp., a_i) has a weakly (resp., strictly) higher mean payoff against a sample size of k_i iff (neglecting rare events of having multiple b_j -s (resp., a_j -s) in the sample): (1) the sample includes the opponent's action b_j (resp., a_j), and (2) $k_i < \frac{1}{u_i} + 1$ (resp., $k_i \le u_i + 1$). The probability of (1) is $k_i \cdot \epsilon_j + o(\epsilon_j)$, where $o(\epsilon_j)$ denotes terms that are sublinear in ϵ_j , and, thus, it will not affect the Jacobian as $\epsilon_j \to 0$. This implies that the probability that a new agent in population i (with a random sample size distributed according to θ_i) who has a higher mean payoff for action b_i (resp., a_i) against her sample is $w_i(1-\epsilon_j) = \epsilon_j \cdot \mathbb{E}_{<\frac{1}{u_i}+1}(\theta_i) + o(\epsilon_j)$ (resp., $w_i(\epsilon_j) = \epsilon_j \cdot \mathbb{E}_{\le u_i+1}(\theta_i) + o(\epsilon_j)$). Therefore, the sampling dynamics at (ϵ_1, ϵ_2) (resp., $(1 - \epsilon_1, 1 - \epsilon_2)$) can be written as follows (ignoring the higher-order terms of ϵ_1 and ϵ_2):

$$\dot{\epsilon}_{i} = \epsilon_{j} \cdot \mathbb{E}_{\leq \frac{1}{u_{i}}+1}(\theta_{i}) - \epsilon_{i} \quad (\text{resp.}, \dot{\epsilon}_{i} = \epsilon_{i} - \epsilon_{j} \cdot \mathbb{E}_{\leq u_{i}+1}(\theta_{i})). \tag{A.3}$$

Define: $a_{\theta_i} = \mathbb{E}_{<\frac{1}{u_i}+1}(\theta_i)$ (resp., $b_{\theta_i} = \mathbb{E}_{\le u_i+1}(\theta_i)$). The Jacobian of the above system of Equations (A.3) is then given by $J_a = \begin{pmatrix} -1 & a_{\theta_1} \\ a_{\theta_2} & -1 \end{pmatrix}$ (resp., $J_a = \begin{pmatrix} 1 & -b_{\theta_1} \\ -b_{\theta_2} & 1 \end{pmatrix}$). The eigenvalues of J_a (resp., J_b) are $-1 - \sqrt{a_{\theta_1}a_{\theta_2}}$ and $-1 + \sqrt{a_{\theta_1}a_{\theta_2}}$ (resp., $-1 - \sqrt{b_{\theta_1}b_{\theta_2}}$ and

 $-1 + \sqrt{b_{\theta_1}b_{\theta_2}}$). Observe that: (1) if $a_{\theta_1}a_{\theta_2} < 1$ (resp., $b_{\theta_1}b_{\theta_2}>1$) then both eigenvalues are negative, which implies that the pure state \boldsymbol{a} (resp., \boldsymbol{b}) is asymptotically stable, and (2) if $a_{\theta_1}a_{\theta_2} > 1$ (resp., $b_{\theta_1}b_{\theta_2}>1$) then one of the eigenvalues is positive, which implies that this state is unstable (see, e.g., Perko, 2013, Theorems 1 and 2 in Section 2.9).

A.6 Proof of Proposition 2

A.6.1 Proof of Claim 1 (Reaching the Area between the Curves)

We say that state **p** is above (resp., below) curve $w_2(p_1)$ if $p_2 > w_2(p_1)$ (resp., $(p_2 < w_2(p_1))$). Similarly, we say that state **p** is to the *right* (resp., *left*) of the curve $w_1(p_2)$ if $p_1 > w_1(p_2)$ (resp., $p_1 < w_1(p_2)$). We say that the state **p** is on the curve $w_i(p_j)$ if $p_i = w_i(p_j)$. Due to the fact that the two curves are strictly increasing, we identify the notion of being above a curve and being to the left of the curve, and similarly we identify the notion of being below a curve and being to the right of the curve. The states on the curve $w_2(p_1)$ (resp., $w_1(p_2)$) are characterized by having $\dot{p}_2 = 0$ (resp., $\dot{p}_1 = 0$). Observe that $\dot{p}_2 > 0$ (resp., $\dot{p}_2 < 0$) in any state **p** above and to the left (resp., below and to the right) of the curve $w_2(p_1)$. Similarly, $\dot{p}_1 > 0$ (resp., $\dot{p}_1 < 0$) in any state **p** above and to the left (resp., below and to the right) of the curve $w_2(p_1)$.

Any state $\mathbf{p} \in [0, 1]$ can be classified in one of $9 = 3 \cdot 3$ classes, depending on its relative location with respect to the two curves, i.e., whether \mathbf{p} is below, above, or on each of the two curves $w_i(p_j)$. If state \mathbf{p} is on (resp., above, below) the curve $w_2(p_1)$, then \dot{p}_2 is zero (resp., negative, positive). Similarly, if state \mathbf{p} is on (resp., above, below) the curve $w_1(p_2)$, then \dot{p}_1 is zero (resp., positive, negative). In particular, any state \mathbf{p} that is above (resp., below) both curves must satisfy $\dot{p}_1 > 0 > \dot{p}_2$. This implies that any trajectory that begins above (resp., below) both curves must always move downward and to the right. This (together with the fact that $\dot{\mathbf{p}}(t) = (0,0)$ only in the intersections of the two curves) implies that the trajectory must either converge to a stationary state (i.e., an intersection of the two curves), or cross one of the curves, and reach a state $\mathbf{p}(t)$ that satisfies either $w_1^{-1}(p_1(t)) \leq p_2(t) \leq w_2(p_1(t))$ or $w_2(p_1(t)) \leq p_1(t) \leq w_1^{-1}(p_1(t))$.

A.6.2 Proof of Claim 2 (Convergence to Stationary States)

We first show why we can assume w.l.o.g. that p(t) is strictly between the two curves. If p(t) crosses one of the curves and is strictly above (resp., below) the remaining curve, then the dynamics must take the populations to a state that is strictly below one of the curves and strictly above the remaining curve. This is so because on the crossing point one of the \dot{p}_i -s is zero and the remaining derivative \dot{p}_j is negative (resp., positive), which implies that the dynamics take the trajectory below and to the right (resp., above and to the left) of the curve that was crossed.

Next assume that $p_i(t) \in (w_1^{-1}(p_1(t)), w_2(p_1(t)))$ (resp., $p_i(t) \in (w_2(p_1(t)), w_1^{-1}(p_1(t)))$). By the classification presented in the proof of Claim 1, the trajectory must move upward and to the right, i.e., $\dot{p_1}, \dot{p_2} > 0$ (resp., downward and to the left, i.e., $\dot{p_1}, \dot{p_2} < 0$). This implies that the trajectory must cross one of the curves. The crossing point cannot be only on the curve of $w_2(p_1)$ (resp., $w_1(p_2)$), because at such a point the trajectory moves horizontally to the left (vertically upward), which implies that it must cross the lower (resp., higher) curve $w_2(p_1)$ (resp., $w_1(p_2)$) from the left side (resp., from below), and we get a contradiction. This implies that the crossing point must be the closest intersection points of the two curves to the right (resp., left) of $p_i(t)$, namely, \overline{p} (resp., \underline{p}).

A.7 Proof of Lemma 1 (Convergence to Pure States)

If **p** is below (resp., above) both curves, then by the classification presented in the proof of Claim 1 it must be that $\dot{p}_2 > 0$ (resp., $\dot{p}_1 > 0$), which implies that convergence to (0,0) is possible only if the trajectory passes through a state that is strictly between the two curves, and that the closest intersection point of the two curves to the left of this state is (0,0). By the classification presented in the proof of Claim 2 it must be that the curve of $w_2(p_1)$ is strictly below the curve of $w_1(p_2)$ in a right neighborhood of (0,0), which implies that (0,0) is asymptotically stable because any sufficiently close initial state would converge to (0,0).

A.8 Proof of Lemma 2

Assume that Conditions (1) and (2) hold. Let $p \neq \hat{p}$ be any sufficiently close state. By Claim 1 any trajectory beginning at p will enter one of the two areas between the two curves on either side of \hat{p} . By Claim 2, Condition (1) (resp., (2)) implies convergence to \hat{p} if the trajectory has entered the area between the curves to the left (resp., right) of \hat{p} . This implies that any trajectory that starts sufficiently close to \hat{p} must converge to \hat{p} , and, thus, \hat{p} is asymptotically stable.

Next assume that $\hat{p}_1 > 0$ and Condition (1) (resp., $\hat{p}_1 < 1$ and Condition (2)) is not satisfied. This implies that $w_2(p_1) < w_1^{-1}(p_1)$ (resp., $w_2(p_1) > w_1^{-1}(p_1)$) for any p_1 that is sufficiently close to \hat{p}_1 from the left (resp., right). By Claim 2, this implies that a trajectory starting at (p_1, p_2) with p_1 sufficiently close to \hat{p}_1 from the left (resp., right) and with $p_2 \in (w_2(p_1), w_1^{-1}(p_1))$ (resp., $p_2 \in (w_1^{-1}(p_1), w_2(p_1))$) converges to the neighboring stationary point on the left (resp., right) side of \hat{p} . Thus, \hat{p} is unstable.

A.9 Proof of Theorem 3 (Stability of Miscoordination)

Let $w_i^{k_i}(p_j)$ (resp., $w_i^{\theta_i}(p_j)$) denote the sampling dynamics induced by an agent with sample size k_i (resp., agents with the distribution θ_i).

1. Fix any $r < \min\left(\frac{1}{k_1}, \frac{1}{k_2}\right)$. Assume that $u_i > \frac{1}{r}$ for each i (which implies that $u_i > k_i$). Observe that

$$w_i^{k_i}(p_j) = \Pr\left(X\left(k_i, p_j\right) \ge 1\right) = 1 - (1 - p_j)^{k_i} = k_i p_j - \binom{k_i}{2}(p_j)^2 + O\left((p_j)^3\right).$$

Fix a sufficiently small $\epsilon > 0$. Let $\hat{p}_1 \in (0, p_1^{NE} - \epsilon)$ be sufficiently small such that the term $O((p_j)^3) < \epsilon$ is negligible for any $p_j < \hat{p}_1$. For each $i \in \{1, 2\}$, let $\alpha_i \in (0, 1)$ be such that: (1) $\alpha_i k_i > 1$ and (2) $\alpha_i \left(k_i - {k_i \choose 2} p_j\right) < 1 - 2\epsilon$. This implies that $w_i^{k_i}(p_j) > \frac{p_j}{\alpha_i}$ in a right neighborhood of zero, and $w_i^{k_i}(p_j) < \frac{p_j}{\alpha_i}$ in a left neighborhood of \hat{p}_1 . Observe that $\lim_{k \to \infty} w_2^k(\hat{p}_1) = 0$ and $\lim_{k \to \infty} (w_1^k)^{-1}(\hat{p}_1) = p_2^{NE}$. This implies that there exists \bar{k} sufficiently large such that $w_2^{\bar{k}}(\hat{p}_1) < \epsilon$ and $\left(w_1^{\bar{k}}\right)^{-1}(\hat{p}_1) > p_2^{NE} - \epsilon > \hat{p}_1$. Observe that $w_i^{k_i \alpha_i \bar{k}}(p_j) > p_j$ in a right neighborhood of \hat{p}_1 . This, in turn, implies that $w_2^{k_2 \alpha_2 \bar{k}}(p_1) < p_1 > (w_1^{k_2 \alpha_2 \bar{k}})^{-1}(p_1)$ in a right neighborhood of zero, and $w_2^{k_2 \alpha_2 \bar{k}}(p_1) > p_1 < (w_1^{k_1 \alpha_1 \bar{k}})^{-1}(p_1)$ in a left neighborhood of \hat{p}_1 . Thus, there exists a stationary state \tilde{p} that satisfies $0 < \tilde{p}_1 < \hat{p}_1 < 1$, and $w_2^{k_i \alpha_i \bar{k}}(p_1) > (w_1^{k_i \alpha_i \bar{k}})^{-1}(p_1)$ (resp., $w_2^{k_i \alpha_i \bar{k}}(p_1) < (w_1^{k_i \alpha_i \bar{k}})^{-1}(p_1)$) in a left (resp., right) neighborhood of \tilde{p}_1 . This

implies, by Claim 2, that \tilde{p} is asymptotically stable. The argument in this case is illustrated in the left panel of Figure 6.1.

2. Observe that

 $w_1^{k_1}(p_2) = \Pr(a \text{ new agent in population 1 has at least one } a_2 \text{ in her sample})$

is the same for all values of $u_1 > k_1$, and, similarly,

 $w_{2}^{k_{2}}(p_{1}) = \Pr(a \text{ new agent in population } 2 \text{ has no } b_{1}\text{-s in the sample})$

is the same for all values of $u_2 < \frac{1}{k_2}$. Fix a sufficiently small $\epsilon > 0$. Let $\bar{u}_1 > k_1, \bar{u}_2 < \frac{1}{k_2}$ be payoffs satisfying $p_2^{NE} = \frac{1}{1+\bar{u}_1} < \frac{w_2^{k_2}(\frac{1}{2})}{2} - \epsilon$ and $1 - p_1^{NE} = \frac{\bar{u}_2}{1+\bar{u}_2} < \frac{1-w_1^{k_1}(\frac{1}{2})}{2} - \epsilon$. Let $r = \min\left(\frac{1}{1+\bar{u}_1}, \frac{\bar{u}_2}{1+\bar{u}_2}\right)$. Fix any payoff profile $u_1 > \frac{1}{r}$ and $u_2 < r$. Let \bar{k} be sufficiently large such that $w_i^{\bar{k}}\left(p_j^{NE} - \epsilon\right) < \epsilon, w_i^{\bar{k}}\left(p_j^{NE} + \epsilon\right) > 1 - \epsilon$. In what follows, we show that the environment $\left((u_1, u_2), (k_1 \frac{1}{2}\bar{k}, k_2 \frac{1}{2}\bar{k})\right)$ admits an asymptotically stable interior state (i.e., we let $\alpha_1 = \alpha_2 = \frac{1}{2}$).

By Proposition 1, the two pure equilibria are asymptotically stable and $w_2^{k_2 \frac{1}{2}\bar{k}}(p_1) < (w_1^{k_1 \frac{1}{2}\bar{k}})^{-1}(p_1)$ (resp., $w_2^{k_2 \frac{1}{2}\bar{k}}(p_1) > (w_1^{k_1 \frac{1}{2}\bar{k}})^{-1}(p_1)$) in a right (resp., left) neighborhood of zero (resp., one). Next observe that $w_2^{k_2 \frac{1}{2}\bar{k}}(\frac{1}{2}) > \frac{1}{2}w_2^{k_2}(\frac{1}{2}) > p_{NE}^2 + \epsilon > (w_1^{k_1 \frac{1}{2}\bar{k}})^{-1}(\frac{1}{2})$, which implies that there is an (unstable) interior stationary state \hat{p} that satisfies: (1) $\hat{p}_1 < 0.5$, (2) $p_{NE}^2 - \epsilon < \hat{p}_2 < p_{NE}^2 + \epsilon < 0.5$, and (3) $w_2^{k_2 \frac{1}{2}\bar{k}}(p_1) > (w_1^{k_1 \frac{1}{2}\bar{k}})^{-1}(p_1)$ in a right neighborhood of \hat{p}_1 . By an analogous argument, there is also an (unstable) interior stationary state \tilde{p} that satisfies: (1) $\tilde{p}_2 > 0.5$, (2) $p_1^{NE} - \epsilon < \tilde{p}_1 < p_1^{NE} + \epsilon < 0.5$, and (3) $w_2^{k_2 \frac{1}{2}\bar{k}}(p_1) < (w_1^{k_1 \frac{1}{2}\bar{k}})^{-1}(p_1)$ in a left neighborhood of \tilde{p}_1 . This implies that must be an interior stationary state \hat{p} between \hat{p} and \tilde{p} such that $w_2^{k_2 \frac{1}{2}\bar{k}}(p_1) > (w_1^{k_1 \frac{1}{2}\bar{k}})^{-1}(p_1)$ (resp., $w_2^{k_2 \frac{1}{2}\bar{k}}(p_1) > (w_1^{k_1 \frac{1}{2}\bar{k}})^{-1}(p_1))$ in a left (resp., right) neighborhood of \hat{p} . Finally, Claim 2 implies that \hat{p} is asymptotically stable.

References

- Arieli, I., Y. Babichenko, R. Peretz, and H. P. Young (2020). The speed of innovation diffusion in social networks. *Econometrica* 88(2), 569–594.
- Arigapudi, S., Y. Heller, and I. Milchtaich (2021). Instability of defection in the prisoner's dilemma under best experienced payoff dynamics. *Journal of Economic Theory 197*, 105174.
- Bilancini, E., L. Boncinelli, S. Ille, and E. Vicario (2022). Memory retrieval and harshness of conflict in the hawk-dove game. *Economic Theory Bulletin* 10(2), 333–351.
- Brunner, C., C. F. Camerer, and J. K. Goeree (2011). Stationary concepts for experimental 2× 2 games: Comment. *American Economic Review* 101(2), 1029–1040.
- Cárdenas, J., C. Mantilla, and R. Sethi (2015). Stable sampling equilibrium in common pool resource games. *Games* 6(3), 299–317.
- Crawford, V. P. (1989). Learning and mixed-strategy equilibria in evolutionary games. Journal of Theoretical Biology 140(4), 537–550.
- Danenberg, T. and R. Spiegler (2022). Representative sampling equilibrium. Working Paper.
- Fudenberg, D. and D. K. Levine (1995). Consistency and cautious fictitious play. Journal of Economic Dynamics and Control 19(5–7), 1065–1089.
- Genesove, D. and C. Mayer (2001). Loss aversion and seller behavior: Evidence from the housing market. *Quarterly Journal of Economics* 116(4), 1233–1260.
- Goeree, J. K., C. A. Holt, and T. R. Palfrey (2003). Risk averse behavior in generalized matching pennies games. *Games and Economic Behavior* 45(1), 97–113.
- Green, D. N. (1983). Fixed point properties of the binomial function. Journal of the Franklin Institute 316(1), 51–62.
- Harsanyi, J. C. and R. Selten. A General Theory of Equilibrium Selection in Games. MIT Press.

- Hauert, C. and J. Miekisz (2018). Effects of sampling interaction partners and competitors in evolutionary games. *Physical Review E* 98(5), 052301.
- Heller, Y. (2017). Instability of belief-free equilibria. Journal of Economic Theory 168, 261–286.
- Heller, Y. and E. Mohlin (2018). Social learning and the shadow of the past. Journal of Economic Theory 177, 426–460.
- Herold, F. and C. Kuzmics (2020). The evolution of taking roles. Journal of Economic Behavior & Organization 174, 38–63.
- Kandori, M., G. J. Mailath, and R. Rob (1993). Learning, mutation, and long run equilibria in games. *Econometrica* 61(1), 29–56.
- Kaniovski, Y. M. and H. P. Young (1995). Learning dynamics in games with stochastic perturbations. *Games and Economic Behavior* 11(2), 330–363.
- Koster, H. and J. Rouwendal (2021). Housing market discount rates: Evidence from bargaining and bidding wars. Working Paper.
- Kreindler, G. E. and H. P. Young (2013). Fast convergence in evolutionary equilibrium selection. *Games and Economic Behavior* 80, 39–67.
- Larsen, B. J. (2021). The efficiency of real-world bargaining: Evidence from wholesale used-auto auctions. *Review of Economic Studies* 88(2), 851–882.
- Lyu, J., Z. Li, and Y. Xu (2022). Inexact information, strategic sophistication and equilibrium transition: A quasi-continuous-time experiment. Working Paper.
- Mantilla, C., R. Sethi, and J. C. Cárdenas (2018). Efficiency and stability of sampling equilibrium in public goods games. *Journal of Public Economic Theory* 22(2), 355–370.
- Maynard Smith, J. (1982). *Evolution and the Theory of Games*. Cambridge University Press.
- Maynard Smith, J. and G. R. Price (1973). The logic of animal conflict. *Nature 246* (5427), 15–18.

- McKelvey, R. D., T. R. Palfrey, and R. A. Weber (2000). The effects of payoff magnitude and heterogeneity on behavior in 2× 2 games with unique mixed strategy equilibria. *Journal of Economic Behavior & Organization* 42(4), 523–548.
- Morris, S., R. Rob, and H. S. Shin (1995). p-dominance and belief potential. *Economet*rica 63(1), 145–157.
- Oprea, R., K. Henwood, and D. Friedman (2011). Separating the hawks from the doves: Evidence from continuous time laboratory games. *Journal of Economic Theory* 146(6), 2206–2225.
- Osborne, M. J. and A. Rubinstein (1998). Games with procedurally rational players. American Economic Review 88(4), 834–847.
- Osborne, M. J. and A. Rubinstein (2003). Sampling equilibrium, with an application to strategic voting. *Games and Economic Behavior* 45(2), 434–441.
- Oyama, D., W. H. Sandholm, and O. Tercieux (2015). Sampling best response dynamics and deterministic equilibrium selection. *Theoretical Economics* 10(1), 243–281.
- Perko, L. (2013). Differential Equations and Dynamical Systems, Volume 7. Springer Science & Business Media.
- Salant, Y. and J. Cherry (2020). Statistical inference in games. *Econometrica* 88(4), 1725–1752.
- Sandholm, W. H. (2001). Almost global convergence to p-dominant equilibrium. International Journal of Game Theory 30(1), 107–116.
- Sandholm, W. H. (2010). Population Games and Evolutionary Dynamics. MIT Press.
- Sandholm, W. H., S. S. Izquierdo, and L. R. Izquierdo (2019). Best experienced payoff dynamics and cooperation in the centipede game. *Theoretical Economics* 14(4), 1347– 1385.
- Sandholm, W. H., S. S. Izquierdo, and L. R. Izquierdo (2020). Stability for best experienced payoff dynamics. *Journal of Economic Theory* 185, 104957.

- Sawa, R. and J. Wu (2023). Statistical inference in evolutionary dynamics. Games and Economic Behavior 137, 294–316.
- Selten, R. and T. Chmura (2008). Stationary concepts for experimental 2x2-games. American Economic Review 98(3), 938–66.
- Sethi, R. (2000). Stability of equilibria in games with procedurally rational players. Games and Economic Behavior 32(1), 85–104.
- Sethi, R. (2021). Stable sampling in repeated games. Journal of Economic Theory 197, 105343.
- Tversky, A. and D. Kahneman (1971). Belief in the law of small numbers. Psychological Bulletin 76(2), 105.
- Weibull, J. W. (1997). Evolutionary Game Theory. MIT Press.
- Young, H. P. (1993). The evolution of conventions. *Econometrica* 61(1), 57–84.