

Cognitive Biases in Stochastic Coordination Games and Their Evolution

Daniel H. Wood*

December 2013

Preliminary and Incomplete

Abstract

I model the evolution of behavior in small groups whose members play pairwise 2×2 pure coordination games. Players follow a best-response dynamic with errors, in which they either best-respond to the true strategy group distribution or to beliefs about the distribution generated by the representativeness heuristic or availability heuristic. Mistakes in strategy choice impair successful coordination, but they also allow the group to occasionally switch between equilibria. If the equilibrium which generates higher payoffs also occasionally changes, then errors create a positive externality for other players because the group returns to high-payoff equilibria faster. I characterize group members' payoffs in this setting for varying error rates. Biases in belief formation produce larger externalities relative to simply higher error rates. I then analyze numerically how these biases would evolve in a setting with group structure. Group selection partially internalizes the positive externalities, stabilizing a population state in which some players have biases and others do not.

JEL Classification: C72, C73, D03, D83.

*Department of Economics, Clemson University, 228 Surrine Hall, Clemson, SC 29634. Email: dwood2@clemson.edu. Phone: 1-864-656-4740. Fax: 1-864-656-4192.

1 Introduction

A stochastic coordination game is a repeated pure coordination game in which the payoffs change periodically. An example is depicted in Figure 1. It is well-known that in coordination games, under myopic best response dynamics with errors, a population's behavior converges quickly on a Nash equilibrium of the game and then rarely but with positive probability switches between equilibria. I show that when players' behaviors are distorted by two cognitive biases – either the representativeness heuristic or the availability heuristic – that the distortions speed the transition of player groups from low-paying equilibria to high-paying equilibria. I apply this insight to argue that stochastic coordination environments could have influenced the evolution of cognitive biases. In a stochastic coordination environment where players interact in groups, biased individuals produce positive externalities for other group members. If groups compete in addition to competition between individuals within groups, groups with a larger number of biased individuals can replace groups with fewer biased individuals, leading to a stable state in which some members of a population suffer a bias while others do not.

	A	B
A	2, 2	0, 0
B	0, 0	1, 1

Game I

	A	B
A	1, 1	0, 0
B	0, 0	2, 2

Game II

Figure 1: Stochastic Coordination Game. The normal form being played alternates back and forth between I and II, with a constant small probability of switching every period.

In my model, there is an infinite population of players who split into fixed small groups every period. Within a period, each player interacts for many subperiods with other members

of their group. Group members are repeatedly randomly matched into pairs who then play a 2×2 coordination game with two strict Nash equilibria, A and B, with differing payoffs. Each subperiod, all agents play the game and then simultaneously update their strategies using a myopic best response dynamic with errors (KMR-style dynamics (Kandori et al. 1993)).

Players may be one of two cognitive types: rational, who form their beliefs about other players' behavior by observing the average behavior of all other group members and almost always optimize given their beliefs, or biased types, who form their beliefs in some other way. For example, they may sample the strategies of a few members of the group and extrapolate group behavior from that sample (i.e., use the representativeness heuristic) or overweight a common focal observation (i.e., the availability heuristic). Individual behavior and the prevalence of cognitive biases are determined by separate evolutionary processes, one embedded within the other.

One part of this paper characterizes the expected payoff of an individual player over a period for a simplified version of the model in which there is no correlation between errors. His payoff is decomposable into three parts. Because errors are rare, each group spends most subperiods coordinating on one of the equilibria of the coordination game. The frequency with which a player i successfully coordinates depends on i 's type and what fraction of players in i 's group are of each type. More bias in either of these two parts reduces the chance of coordination and through that, i 's payoff. However, errors also cause the group to occasionally switch equilibria. The third part of i 's expected payoff corresponds to how much time the group spends in the higher-payoff equilibrium. If the payoffs to coordinating on A and B sometimes change, errors persistently create positive externalities for other players, because the group returns to the payoff-dominant equilibrium more quickly.

Next I consider how the frequency of biased players in the population will change over time. I model the evolution of biases over a series of periods as occurring at the population level through replicator dynamics in which more successful cognitive types increase in frequency. Because the small group structure allows biased types to capture some of the positive externality they produce, selection for the presence of cognitive biases in the population can occur.

Under these dynamics, there are usually two stable steady states in the frequency of biased players. One is a state in which some members of the population are biased and others are rational. The other is a state in which all members of the population are rational. These two states occur because biased individuals are complementary to each other – so if biased individuals are infrequent, rational-type players do better than biased-type players – and because the benefit they produce is an externality – so “free-riding”-like problems mean that if all players were biased, some would do better as rational types.

Rational players by assumption make errors with small positive probability, but these errors are uncorrelated across individuals within a subperiod or across subperiods and are also equi-probable regardless of which equilibrium the group is in. Cognitive biases in general produce correlated errors across individuals and sometimes across time. For example, if agents suffer from the representativeness heuristic, they believe small observed samples of other players behavior are more representative of the population than they actually are.¹ These noisy beliefs produce more errors in bad states than in good states, because the marginal effect of noise on transition probabilities is higher when the equilibrium is unattractive relative to the

¹The representativeness heuristic, sometimes referred to as the “law of small numbers”, is a probability estimation heuristic where biased individuals treat small samples as more informative about the population than the sample would be treated in forming a posterior belief via Bayes’ Law. Tversky and Kahneman (1974) introduced the idea to economists, while Rabin (2002) models how the bias affects behavior. Kahneman (2011) is a popular exposition.

other equilibrium. Errors by biased individuals are therefore positively correlated, meaning that larger switching probabilities can be produced with fewer errors, improving within-period coordination relative to uncorrelated errors. Cognitive biases are especially beneficial for coordination dynamics because correlated errors usually produce state-dependent errors, clusters of errors, or both. In a companion paper I show these correlations cause fast convergence to the stochastically stable equilibrium regardless of population size for populations with high enough exogenous levels of bias (Wood 2013). This result qualifies the well-known objection to standard stochastic stability analysis that convergence is slow (Ellison (2000), but also see Kreindler and Young (2013)).

In addition to the representativeness heuristic I consider the availability heuristic. With the availability heuristic, people estimate the likelihood of events by the ease with which they can remember examples of the event. If the same group member's strategic decision is easily recalled by all players with this bias, they each overweight the same observation, leading again to correlated mistakes and clusters of errors.

The paper is organized as follows: Section 2 provides a brief example of within-group behavioral evolution and the general effects of cognitive biases in a stochastic setting. Section 3 relates this paper to the prior literatures in behavioral economics and evolutionary game theory. Section 4 and Section 5 then develop within-period models of behavioral evolution and across-period models of cognitive bias evolution, respectively. Finally Section 6 concludes.

2 Motivating Example

Consider a group of 10 players playing the coordination game on the left-hand side of Figure 1, marked Game I. Let $s_{i,t}$ be the strategy choice of player i at time t . Let $BR(s)$ be the

best response to a vector of strategies s . If the fraction of A strategies in s is greater than $1/3$, then $BR(s) = A$, and otherwise $BR(s) = B$. Let $WR(s) = S \setminus BR(s)$.

Players revise their strategies using the following rule:

$$s_{i,t} = \begin{cases} BR(s_{t-1}) & \text{with probability } 1 - \epsilon \\ WR(s_{t-1}) & \text{with probability } \epsilon \end{cases} \quad (1)$$

where $1 \gg \epsilon > 0$. Equation (1) defines a Markov process on a state space $\mathcal{A} = \{0, 1, \dots, 10\}$, whose state a_t is the number of players playing A at time t . a_t is distributed binomially:

$$a_t \sim \begin{cases} B(10, 1 - \epsilon) & \text{if } a_{t-1} \geq 4 \\ B(10, \epsilon) & \text{if } a_{t-1} < 4 \end{cases}$$

The probability of switching from a state where most group members play A to one where most group members play B is given by $\Pr(a_t < 4 | a_{t-1} \geq 4)$ while the probability of switching from B to A is given by $\Pr(a_t \geq 4 | a_{t-1} < 4)$. The probabilities are polynomials in ϵ . Because ϵ is small, higher-order terms in these polynomials are much smaller than the lower-order terms, so the probability of an A to B switch is proportional to ϵ^7 , while the probability of an B to A switch is proportional to ϵ^4 .²

In a stochastic environment, switches akin to that from Game II to Game I leave the group coordinating on a payoff-dominated strategy, so higher error rates are especially beneficial in stochastic environments. If at some time t , the group finds itself in a state $a_t < 4$, a higher ϵ for all members can benefit the group because the transition to the higher-payoff

²Kandori et al. (1993) and Young (1993) were the first papers to characterize the behavior of stochastic evolutionary processes of this nature, using Friedlin and Wentzell (1998)'s methods for determining the limiting distribution of perturbed dynamical systems.

A equilibrium will be more likely at each t . For example, a group with $\epsilon = 1/10$ spends approximately 80 more periods in B conditional on being in B, but increasing ϵ to $\epsilon = 3/20$ would reduce the expected time to approximately 20 periods.³ However, high error rates reduce the chances of successful coordination. If all other players play A at $t - 1$ with probability ϵ_{-i} , player i 's expected payoff is $2(1 - \epsilon_{-i})(1 - \epsilon_i)$. Increases in ϵ_i thus lower private payoffs but produce (on net) positive externalities for the other players.

In equation (1), errors are i.i.d., but cognitive biases produce positive correlation between a_{t-1} and a_t and within a time period, between individual $s_{i,t}$ and $s_{j,t}$. These correlations can lead to much higher error rates in the B than in the A equilibrium.

For example, agents using the representativeness heuristic treat a small sample as more informative than it actually is. If instead of best-responding to $s_{-i,t-1}$ as in (1), cognitively biased players sample two strategies from $t - 1$ and best respond to that sample, then A strategy choices are much more likely. If at least one of the players they sample is playing A in $t - 1$, a biased player will best respond with A because A is the best response if at least 1/3 of the other players are playing A . Figure 2 compares the probability of a biased player choosing A to that of an unbiased player choosing A , for given a_{t-1} . $\Pr(s_{i,t} = A)$ increases swiftly for biased players, so for instance at $a_{t-1} = 3$, $\Pr(s_{i,t} = A) = (3/10) + (1 - 3/10)(3/10) \approx 1/2$.

Representativeness allows errors to “build up” over several periods. Consider the following Markov transition matrix for a group of players who all follow the representativeness heuristic, with $\epsilon = 1/10$. Let T be a 5×5 matrix where $t_{ij} \equiv \Pr(a_{t+1} = j | a_t = i)$ for $i, j \in \{0, 1, \dots, 4\}$ and let the absorbing state 4 index any state $a_t \geq 4$. Then under the

³In contrast, the $\epsilon = 1/10$ group stays in A for an expected 100,000 periods, and increasing ϵ to $\epsilon = 3/20$ reduces the expected time to 7430 periods – in other words, payoff-reducing transitions almost never occur for either error rate.

Expected durations are calculated using standard results on time to absorption for absorbing Markov processes and treating and treating strategy combinations in the other basin of attraction as a single absorbing state.




Figure 2: Probability that representativeness heuristic player (blue) or unbiased player (green) adopts strategy A given number of A strategies at $t - 1$.

heuristic, T is

$$T = \begin{bmatrix} 0.35 & 0.39 & 0.19 & 0.06 & 0.01 \\ 0.05 & 0.17 & 0.27 & 0.25 & 0.26 \\ 0 & 0.04 & 0.12 & 0.21 & 0.63 \\ 0 & 0.01 & 0.03 & 0.09 & 0.87 \\ 0 & 0 & 0 & 0 & 1 \end{bmatrix} \quad (2)$$

Inspection shows that $\Pr(a_t + 1 > a_t) > 1/2$, and in fact a group with $\epsilon = 1/10$ that switches to the B equilibrium spends an average of 5 periods in it.⁴ The heuristic has two effects. Biased players make more errors in the B equilibrium because they only need to observe one A player to choose A , while in the A equilibrium they need to observe both players playing

⁴The transition matrix for players who do not use the heuristic but have $\epsilon = 1/10$ would be a matrix with every row equal to the first row of T , because that row is the transition probabilities for any $a_{t-1} < 4$ if errors are i.i.d.

B to choose B . In addition, this effect is magnified because errors feed back over time. More errors make it easier to observe enough errors to make an error oneself.

While biased players also increase transitions from high-payoff A to low-payoff B , qualitatively these “bad” transitions remain slow. When the group is coordinating on A , both samples of a representative agent must be B in order for the heuristic to cause him to switch, because the probability of B must be at least $2/3$'s for B to be a best response. Therefore increased mis-coordination from the representativeness heuristic in high-payoff states is also limited. Figure 3 shows expected times to transition from B to A (left) or from A to B (right) for groups where some players are biased and best-respond to a sample while others best-respond to the aggregate group behavior.




Figure 3: Expected additional time in B if $a_t = 0$ (left) or expected additional time in A if $a_t = 10$ (right) for possible group compositions.

Because biases speed transitions to high-payoff states, in stochastic coordination games where fast transitions are especially valuable, cognitive biases can be selected for by evolution. In large populations, the marginal effect of a biased individual on transitions is smaller than the increased mis-coordination from the bias, so selection would operate against biases. In a setting with small groups, however, the marginal effect on transitions can be large.

Consider forming a group of 10 players from unbiased and biased players. After the ninth player is added, the tenth player's players payoffs will depend on whether she is biased or unbiased. The left plot in Figure 4 shows expected per-interaction payoffs of either type depending on the composition of the group.⁵ Now consider a population of biased and unbiased players who only interact in groups of size 10, with non-assortative group formation. Denote the fraction of biased players in the population by λ . plot shows the difference in payoffs between unbiased and biased players by share of biased players in the population obtained by integrating the left payoff curves over the (binomial) distribution of group compositions for a given λ . For a population with few biased players – $\lambda < 0.2$ – biased players do worse than unbiased, but then for λ between 0.2 and 0.8, biased players do better than unbiased. If the more successful type's population share grows, there are two stable population states: $\lambda = 0$ and $\lambda = 0.8$.

Error-generating processes in which errors feed back over time are difficult to tractably analyze. In order to illustrate the within-group dynamics in my setting and provide intuition for how particular cognitive biases affect those dynamics, I characterize the more tractable case in which errors are uncorrelated over time, but for which analytic expressions for players' expected payoffs are attainable (Section 4). I then analyze the more general case of error correlation over time numerically in examining the conditions under which cognitive biases might be selected for (Section 5).

⁵Payoffs are calculated assuming that the group begins with $a_0 = 0$ and $a_0 = 10$ each half the time, and that the game repeats with probability $29/30$.

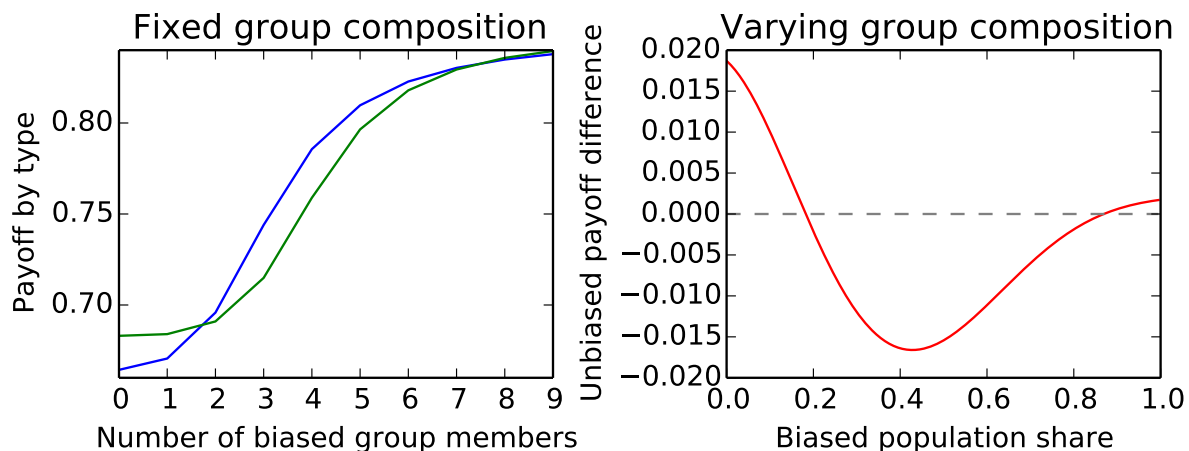


Figure 4: (Left) Payoff to marginal group member if unbiased (green) or biased (blue) by group composition. (Right) Difference in payoffs between unbiased and biased players in a population interacting in groups.

3 Relationship to Literature

Environments that resemble stochastic coordination games played repeatedly by members of a group within a larger population are a fixture of both the present and the distant past. Coordination is important and payoffs are stochastic in technology adoption or the evolution of social mores: interactions in small groups with network externalities or complementarities fit all the characteristics of a stochastic coordination game. One contribution of this paper is to elucidate the effects of cognitive biases in settings like social networking, technology adoption, or fashion.

A second contribution of this paper is to offer a rationale for why cognitive biases might exist despite violating Bayesian standards of rationality. In the evolutionary environment thousands of years ago, small groups were the unit of organization. Technology adoption

within a group or competition between groups both had strong coordination game elements. Group selection in the evolutionary past would have favored these biases because groups with moderate numbers of biased members are superior at coordination than groups without any biased members Bowles (2006) argues that in this environment, intergroup competition could produce altruism.

Competition between groups has most frequently been used to explain the prevalence of altruism. Altruism benefits a group collectively but harms individual altruists, so competition between groups can resolve the puzzle of why altruism would persist. Biased individuals in this paper are analogous to altruistic individuals.⁶ Salomonsson (2010) surveys and classifies various forms of group selection; in his nomenclature my model of group selection is a form of “reproductive externalities”, where group members make other group members more successful at reproduction. Herold (2012) is a recent model of the evolution of punishing or rewarding behavior in Prisoner’s Dilemma games that shares many similarities to this paper. He uses the same period and subperiod time structure and his results rely on the same contrast between within- and across-group payoff differences.

The paper’s analysis of the evolution of cognitive biases uses the indirect evolutionary approach pioneered by Güth and Yaari (1992). The majority of the indirect evolution literature explains the evolution of social preferences through their indirect effect on payoffs through selecting better equilibria (for example, Dekel et al. (2007) or Herold (2012)). Heller (2012) is an exception to this rule, applying the indirect evolutionary approach to explain heterogeneous sophistication in reasoning. He shows that in a finitely repeated Prisoner’s Dilemma environment, a polymorphic population in which some players engage in one round

⁶One difference between altruism and a cognitive bias is that altruistic preferences may allow more flexible behavior than a cognitive bias does; altruists can be reciprocal altruists but biased people cannot choose to “reciprocate” a bias.

of backwards induction and others engage in three rounds can be a neutrally stable state. Another exception is Heifetz and Spiegel (2001), who show that optimistic beliefs can emerge evolutionarily for strategic reasons.

This paper is also related to the literature on errors in stochastic evolutionary game theory. Most directly, group members' behavior evolves through a simple version of KMR-style dynamics (Kandori et al. 1993). A well-known critique of KMR and similar stochastic evolutionary models is that they take implausibly long amounts of time to reach the stable equilibrium (Ellison 2000). Cognitive biases are one reason that behavior might evolve more quickly than simple stochastic evolutionary models predict, which I have shown more systematically in Wood (2013).

Many authors have investigated how different assumptions about errors determine which outcomes are stable in stochastic evolutionary processes. Bergin and Lipman (1996) show that when error rates are state dependent, any equilibrium can be selected; Sawa (2012b) characterizes this effect in 2x2 symmetric coordination games. Robson and Vega-Redondo (1996) investigate the sensitivity of predictions of stochastic evolutionary models to their matching assumptions. These papers are focused on the striking equilibrium selection results of Young (1993) and Kandori et al. (1993), though, and do not consider a stochastic environment or relate errors to cognitive biases. Sawa (2012a) is one of the only papers to combine stochastic behavioral evolution with behavioral agents; he derives the stochastically stable state in a two-stage Nash demand game with outside option when players obey prospect theory.

Finally this paper is related to behavioral economists' modeling of cognitive biases. The "heuristics and biases" literature initiated by Kahneman and Tversky (Tversky and Kahneman 1974) has typically focused on negative effects of probability heuristics. I adopt existing

models of cognitive biases from this literature where possible. While these papers show how biases warp decision-making and reduce welfare, there is another line of research on “fast and frugal heuristics” initiated by Gigerenzer and coauthors that argues that heuristics are beneficial – they require fewer cognitive resources, require less information, and can be reasonably accurate in natural environments. Neither of these literatures focuses on the positive dynamic *interpersonal* effects of heuristics, the subject of this paper.

4 Evolution of Behavior Within Periods

In my setting, interactions between players involve three key elements: a coordination game structure, a small-group setting, and payoffs that change over time. Consider a coordination game as depicted in Figure 5 played repeatedly by members of an infinite population. One form of coordination always yields $\bar{\pi}$, while the other yields $\underline{\pi} < \bar{\pi}$.

Each player i is endowed with a cognitive type that determines how he chooses between actions. Players are of two types, low rationality or biased types (B-types) and high rationality (R-types). In choosing actions, both types myopically respond to the current behavior of the group, but otherwise R-types follow a Bayesian normative standard of rationality of conditioning correctly on all relevant data. Cognitive types will be fully specified in the next section.

	A	B
A	$\pi_{A,t}, \pi_{A,t}$	$0, 0$
B	$0, 0$	$\pi_{B,t}, \pi_{B,t}$

Figure 5: Coordination Game. $\pi_{A,t}, \pi_{B,t} \in \{\bar{\pi}, \underline{\pi}\}$ with $\bar{\pi} > \underline{\pi} > 0$.

Time is divided into periods and subperiods. A period is an indefinite but large number of

subperiods. During a period, players are separated into groups, which are randomly formed at the beginning of the period and dissolved at the end of the period. A subperiod is one play of the coordination game by every member of the population against a random partner in their group. Periods represent the longer time frame over which cognitive types evolves, and subperiods represent the shorter timeframe over which behavior evolves, with the evolution of behavior influenced by cognitive types.

For each group, the within-period evolution of behavior is a Markov process with transition probabilities that are functions of the group composition. Within a group, the two cognitive types have different average payoffs because B's coordinate less well than R's. Across groups, average payoffs also differ because groups with more B's have higher transition probabilities.

4.1 Within-Period Markov Process

At the start of each period, groups of G players are formed, with all the players being assigned to a group. Players' types and group assignment are uncorrelated. Throughout a period, players are endowed with strategies, with $s_{i,t}$ denoting i 's strategy in subperiod t . Assume that all players start with strategy A : $\forall i, s_{i,1} = A$.

During a subperiod, every player is matched with another member of his group and the pair then plays the coordination game. Next all players simultaneously update their strategies and finally the game payoffs may change or the period may end.

Let p_t be the fraction of players playing A at time t . Let $\phi_\tau : [0, 1] \times \{\pi_A, \pi_B\} \rightarrow \Delta S$ be an function describing a cognitive type τ 's strategy choice in response to the distribution of group members' strategies p . For all types, with probability $(1 - 2\epsilon)$ a player of type τ uses the type-specific update function, and with probability 2ϵ the player chooses a random

strategy.

Rational Type: Type-R players correspond to KMR-style players and have $\phi_r(s) = BR(s)$, i.e, they optimally respond to the other players' strategies, with

$$\phi_r(p, \pi_A, \pi_B) = \begin{cases} A & \text{if } \pi_A = \bar{\pi} \text{ and } p > \frac{\pi}{\bar{\pi} + \underline{\pi}} \\ A & \text{if } \pi_A = \underline{\pi} \text{ and } p > \frac{\bar{\pi}}{\bar{\pi} + \underline{\pi}} \\ B & \text{otherwise} \end{cases}$$

Type R's have $\epsilon > 0$ for two reasons: technically, this guarantees that the state space is ergodic for any group composition; practically, assuming that R's make errors, albeit rarely, makes comparisons to type B's easier. This assumption can also be thought of as a behavioral assumption that even rational people occasionally make mistakes.

Representativeness Heuristic Biased Type: The representativeness heuristic, sometimes referred to as the "law of small numbers", is a probability estimation heuristic where biased individuals take small samples as more representative of the population than the sample would be if used to form a posterior belief via Bayes' Law (Tversky and Kahneman 1974; Rabin 2002).

Let ϕ_B be such that a B-type samples $S = 2$ strategies from s_{t-1} at random with replacement, forms beliefs about the distribution of behavior in the group using that sample, and best responds to that belief. Let \hat{p}_i^t be biased player i 's belief about the fraction of players with $s = A$ at $t - 1$. These beliefs are distributed $B\left(\frac{a_{t-1}}{G}, S\right) / S$. Then ϕ_B assigns probabilities to each strategy choice according to the distribution of best responses to \hat{p}_i^t , i.e., it assigns $\phi_B(s) = A$ with probability $\Pr(\hat{p}_i^t > \frac{\pi}{\bar{\pi} + \underline{\pi}} | s)$ when $\pi_A = \bar{\pi}$. Because that

probability rises in a_{t-1} , errors create a positive feedback loop in which an increase in the number of errors leads to samples with more errors, causing B-types to make more errors, and so on. The feedback loop is stronger in state L because fewer group members need to be non-conforming in L in order for a revising B-type to choose the non-conforming strategy than in H .

Availability Heuristic Biased Type: Players following the availability heuristic overweight the most easy to recall data point when considering which strategy to adopt. Let k be a random “focal” player whose choice s_k is publically announced and is focal for every B-type, and let p_k be an indicator variable for $s_k = A$. Then an availability B-type player estimates the fraction of players who are playing A as

$$\hat{p} = \eta p_k + (1 - \eta)p. \tag{3}$$

In other words, for $\eta > 0$ type-B’s overweight the focal player’s choice with overweighting increasing in η . Let ϕ_B be such that availability heuristic player choose a best response to \hat{p} .

Availability type-B’s either overestimate the true coordination rate, if the focal player is coordinating, or underestimate it, if the focal player is not coordinating. If enough players are type-B’s, the block of biased players not conforming when $p_k = 1$ is the most likely way to transition out of the current state and determines γ_{LH} and γ_{HL} .

Assume for expository purposes that $\eta = 1/2$. Then for any $\underline{\pi} < \bar{\pi}$, if the focal player plays H in L , all type-B players play H with probability $1 - \epsilon$ in the next subperiod. Therefore if there are enough type-B players so that they can cause an equilibrium switch by themselves, then the probability of switching to H is a constant (approximately) ϵ , the

probability that a focal player does not conform.

After strategy updating has occurred, either a new subperiod occurs with the same payoffs, the payoffs in the game may change, or the period may end. With probability $1 - \delta$ the period ends and the subperiod's sequence of games is over. With probability $1 - \sigma$, the values $U(A, A)$ and $U(B, B)$ update. If payoffs update in subperiod t , then in the subperiod $t + 1$,

$$\begin{aligned} \pi_{A,t+1} &= \begin{cases} \bar{\pi} & \text{with probability } 1/2 \\ \underline{\pi} & \text{with probability } 1/2 \end{cases} \\ \pi_{B,t+1} &= \bar{\pi} + \underline{\pi} - \pi_{A,t+1} \end{aligned}$$

For high $\delta\sigma$, the situation the players face at t is very likely to be repeated at $t + 1$. Let $\psi \equiv (1 - \delta\sigma)/(\delta\sigma)$ be a measure of how inconstant the situation is over time.

The strategy and subperiod updating together define a Markov process on a state space consisting of the selected players' strategies. It is difficult to find closed-form expressions for expected period payoffs using this Markov process, but considering a simpler reduced Markov process can provide some intuition. The next section analyzes the special case of the process with two states corresponding to the basin of attraction the group is in at t .

4.2 Reduced Markov Process

Now consider representing the within-period system using two states:⁷

⁷ For a group consisting of R's and baseline B's, the group is in the basin of attraction of S in sub-period t if and only if $BR(s_t) = S$ for enough players so that in $t+2$ all players would play S if no errors occurred after subperiod $t - 1$, i.e., if group behavior converges under the deterministic dynamic to all members playing S .

Defining basins of attraction in terms of time $t + 2$ group behavior is necessary because the best response of i and j at t could differ if $s_{i,t} \neq s_{j,t}$ because their own strategy does not influence their best response. However, if $G > 2$, the impact of the strategy choice of each player is small enough so that if $BR(s_t) = A$, all players play A eventually if no errors occur.

- **High state (“H”)**: The group population is in the basin of attraction of A and $\pi_A = \bar{\pi} > \pi_B = \underline{\pi}$ or the group population is in the basin of attraction of B and $\pi_B = \bar{\pi} > \pi_A = \underline{\pi}$
- **Low state (“L”)**: The group population is in the basin of attraction of A and $\pi_B = \bar{\pi} > \pi_A = \underline{\pi}$ or the group population is in the basin of attraction of B and $\pi_A = \bar{\pi} > \pi_B = \underline{\pi}$

If error rates only depend on which basin of attraction the group is in and which type each player is, as is the case with type-R and baseline type-B players, then none of the information lost with this simpler representation is needed for ϕ_B or ϕ_R . For the representativeness heuristic and availability heuristic ϕ_B functions, error rates depend on exactly how many errors there are in period $t - 1$ and the dynamics of the reduced process do not fully capture the biases’ effects.

Let $\epsilon_{\tau H}$ be the probability of choosing the worst response for type τ when the best response is coordinating on $\pi = \bar{\pi}$ strategy (the high payoff strategy) and let $\epsilon_{\tau L}$ be defined analogously for the low payoff strategy. For rational types, $\epsilon_{RH} = \epsilon_{RL} = \epsilon$ always, while for baseline biased types, it is possible that $\epsilon_{BL} \neq \epsilon_{BH}$.

Denote the set of these reduced states by $\mathcal{S} \equiv \{H, L\}$. Let γ_{LH} be the probability of the underlying state switching to H if the current state is L, and γ_{HL} be the probability of the state switching to L if the current state is H. These probabilities are functions of group composition, ϕ , and the base error rates.⁸

⁸The transition probabilities of the reduced Markov process are

$$\begin{aligned}
\gamma_{HL} &\equiv \Pr(S_t = L | S_{t-1} = H) \\
&= \Pr(U_t(BR(s_t), BR(s_t)) = \underline{\pi} \text{ and } U_t(BR(s_{t-1}), BR(s_{t-1})) = \bar{\pi}) \\
\gamma_{LH} &\equiv \Pr(S_t = H | S_{t-1} = L) \\
&= \Pr(U_t(BR(s_t), BR(s_t)) = \bar{\pi} \text{ and } U_t(BR(s_{t-1}), BR(s_{t-1})) = \underline{\pi})
\end{aligned}$$

The utility player i earns through participation in a single period supergame can be approximated in two steps. First I calculate the expected payoff of a player for a sequence of subperiods in which no payoff-switch occurs. Using this expression, I then calculate the total expected payoff for an entire period, taking into account that payoff-switches occur.

Let N be a matrix with elements n_{ij} expressing the expected number of subperiods spent in state i if the system begins in state j , conditional on payoffs not switching, where H is state 1 and L is state 2. Let Q be a matrix of transition probabilities where q_{ij} is the probability that the state is i at $t + 1$ if the state is j at t :

$$Q = \delta\sigma \begin{pmatrix} 1 - \gamma_{HL} & \gamma_{LH} \\ \gamma_{HL} & 1 - \gamma_{LH} \end{pmatrix}.$$

Then $N = I + Q + Q^2 + \dots = (I - Q)^{-1}$ so

$$\begin{aligned} N &= \left(\frac{1}{1 - \delta\sigma} \right) \begin{pmatrix} \frac{1/\delta\sigma - 1 + \gamma_{LH}}{1/\delta\sigma - 1 + (\gamma_{HL} + \gamma_{LH})} & \frac{\gamma_{LH}}{1/\delta\sigma - 1 + (\gamma_{HL} + \gamma_{LH})} \\ \frac{\gamma_{HL}}{1/\delta\sigma - 1 + (\gamma_{HL} + \gamma_{LH})} & \frac{1/(\delta\sigma) - 1 + \gamma_{HL}}{1/\delta\sigma - 1 + (\gamma_{HL} + \gamma_{LH})} \end{pmatrix} \\ &= \left(\frac{\psi + 1}{\psi} \right) \begin{pmatrix} \frac{\psi + \gamma_{LH}}{\psi + (\gamma_{HL} + \gamma_{LH})} & \frac{\gamma_{LH}}{\psi + \gamma_{HL}} \\ \frac{\gamma_{HL}}{\psi + (\gamma_{HL} + \gamma_{LH})} & \frac{\psi + \gamma_{HL}}{\psi + (\gamma_{HL} + \gamma_{LH})} \end{pmatrix}. \end{aligned}$$

Each element of N is an average of long-run behavior and short-run behavior. This can be seen by taking the limit as $\delta \rightarrow 0$, where $\lim_{\delta \rightarrow 0} N = I$, or as $\psi \rightarrow 0$,

$$\lim_{\psi \rightarrow 0} \psi N = \begin{pmatrix} \frac{\gamma_{LH}}{\gamma_{LH} + \gamma_{HL}} & \frac{\gamma_{LH}}{\gamma_{LH} + \gamma_{HL}} \\ \frac{\gamma_{HL}}{\gamma_{LH} + \gamma_{HL}} & \frac{\gamma_{HL}}{\gamma_{LH} + \gamma_{HL}} \end{pmatrix}$$

if the payoff values do not update in subperiod t . Expressions for γ_{HL} and γ_{LH} are derived in Appendix B.

(the stationary distribution).

Pairs of players can coordinate in two ways: first, they can conform to the group equilibrium by playing the best reply to the population; alternatively, if neither plays the best response to the group population, they can achieve non-conforming coordination. Let M_τ be a matrix whose elements are the probabilities of achieving conforming or non-conforming coordination in a given state, given i 's type and the distribution of types in the group. m_{ij} is the probability of a type τ receiving payoff i (where $\bar{\pi}$ is indexed by 1) in state j , i.e.,

$$M_\tau = \begin{pmatrix} \Pr(u_i = \bar{\pi} \text{ in } H) & \Pr(u_i = \bar{\pi} \text{ in } L) \\ \Pr(u_i = \underline{\pi} \text{ in } H) & \Pr(u_i = \underline{\pi} \text{ in } L) \end{pmatrix} \quad (4)$$

for i such that $\tau = \tau(i)$.

Finally, let

$$\Pi = \begin{pmatrix} \bar{\pi} & \underline{\pi} \\ \bar{\pi} & \underline{\pi} \end{pmatrix}.$$

Then $(\Pi M_\tau N)_{ii}$ is the expected payoff in states L and H weighted by the expected number of subperiods spent in each if the group starts in i , for a constant payoff sequence.

Weighting the diagonal elements of the matrix by the expected of payoff resets will therefore give the total payoff over a period, and multiplying by $1 - \delta$ gives the average per-subperiod payoff. The expected number of payoff resets is

$$1 + \left(\frac{1 - \sigma}{2}\right) \left(\frac{\delta}{1 - \delta}\right)$$

because one occurs in subperiod 1 and then with probability $1 - \sigma$ in every other subperiod payoffs are updated, which causes a payoff switch half the time. The expected number of

payoff resets whose initial state is S is half this number.⁹ Therefore

Theorem 1. *The average per sub-period expected utility earned by a player of type τ in a group of size G with b type-B agents and type-specific behaviors ϕ is*

$$V(\tau, b) = \left(\frac{2 - \delta - \delta\sigma}{4} \right) \text{tr}(\Pi M_\tau N).$$

where M_τ and N are functions of τ , δ , σ , and the group's composition.

The technique of calculating $V(\tau, b)$ using matrix products can be generalized to the non-reduced Markov process in which the state is the number of players playing A , and I use this for calculating numeric results in Section 5.

The following lemmas use the result of Theorem 1 to precisely state the relationships between type B and type R individual payoffs and group compositions. Each player's type influences V in three ways: individually it affects her match rate, through M , and through group composition it changes how often the higher-value coordination is achieved in the group, with compositional effects on both M and N (i.e., the compositional effect dV/db is given by $dV(\tau)/db = \text{tr}(\Pi(\partial M_\tau/\partial b)N) + \text{tr}(\Pi M_\tau(\partial N/\partial b))$). The individual effect of being a B-type on match rate is always deleterious, but the net compositional effect is often beneficial, because the partial term with regard to N is positive and for many parameter values is larger than the other partial term, which is negative. The effect of b through M is approximately constant for any group composition, but for the effect through N , B-types are complementary to each other.

⁹Each time payoffs reset it is equally like that $\pi_A = \bar{\pi}$ or $\pi_B = \bar{\pi}$. Then if at time t the payoffs change, whether the state at $t - 1$ was H or L does not matter: the new state at t is equally likely to be either H or L.

Lemma 1. *If ϵ is low and type-B errors are state-independent, then for any b , $v(R) - v(B) \approx (\epsilon_B - \epsilon_R)v(R)$.*

Proof. In Appendix A. □

Lemma 1 establishes that when errors are state-independent, each individual type-B player would do better if he were a type-R player. The next two lemmas characterize the effect of type-B's through group composition, first through reducing match probabilities (Lemma 2) and then through speeding transitions between high and low payoff states (Lemma 3).

Lemma 2. *The effect of an increase in the number of B-types in a group through the increase's effect on the probability of conforming and non-conforming matches, $\frac{\text{tr}(\Pi(\partial M)N)}{\partial b}$, is*

- i) proportional to $\bar{\epsilon}/G$, where $\bar{\epsilon}$ is a convex combination of $\epsilon_{BL} - \epsilon_{RL}$ and $\epsilon_{BH} - \epsilon_{RH}$, and is*
- ii) larger in magnitude for R-types than for B-types.*

Proof. In Appendix A. □

Lemma 2 establishes that the part of the group composition effect on payoffs due to B-types reducing conforming matching is small.

I now turn to effect of a change in group compositions on payoffs through N , the transition probability component of V . From Lemma B1, $\gamma_{HL} \approx 0$ for ϵ small, in which case

$$N = \begin{pmatrix} 1 & \frac{\gamma_{LH}}{\psi + \gamma_{LH}} \\ 0 & \frac{\psi}{\psi + \gamma_{LH}} \end{pmatrix} \quad (5)$$

That is, when the group starts in H , it spends all remaining sub-periods in the good state, while when the group starts in L , it spends $\psi/(\psi + \gamma_{LH})$ fraction of the remaining sub-periods in the bad state. Using these expected frequencies in each state,

$$V(\epsilon) \propto \text{tr}(\Pi MN) = (\bar{\pi}m_{11} + \underline{\pi}m_{21}) \left(\frac{\psi + 2\gamma_{LH}}{\psi + \gamma_{LH}} \right) + (\bar{\pi}m_{12} + \underline{\pi}m_{22}) \left(\frac{\psi}{\psi + \gamma_{LH}} \right). \quad (6)$$

V is the fraction of time spent in H and in L , respectively, weighted by the expected payoff in each state. It follows that

Lemma 3. *The effect of an increase in the number of B-types in a group through the increase's effect on transition speeds, $\frac{\text{tr}(\Pi M(\partial N))}{\partial b}$, is*

$$\frac{\text{tr}(\Pi M(\partial N))}{\partial b} = \left(\frac{\partial \gamma_{LH}}{\partial b} \right) \left(\frac{\psi}{(\psi + \gamma_{LH})^2} \right) [(\bar{\pi}m_{11} + \underline{\pi}m_{21}) - (\bar{\pi}m_{12} + \underline{\pi}m_{22})]. \quad (7)$$

Proof. Differentiating (6) with respect to b holding M constant gives (7). \square

The middle term in equation (7) implies that b does not affect V if $\delta\sigma$ is either too small or too large. If $\delta\sigma$ is small (ψ is large) the starting state of the Markov process dominates. $\lim_{\delta\sigma \rightarrow 0} N = I$ so b has no effect on N . Likewise if $\delta\sigma$ is too large (ψ is small), the term $\psi/(\psi + \gamma_{LH})^2$ goes to zero, and b again has no effect on V . If the system is too stable, most of the time is spent in the long-run steady state, which is H , and γ_{LH} 's beneficial effect of speeding transitions to H becomes relatively small. Only for intermediate ψ are B's useful.

Lemma 3 also shows that the transition speed effect contains a $\partial\gamma_{LH}/\partial b$ term. Lemma B2 in Appendix B (Transition Probabilities) shows that this derivative is approximately proportional to γ_{LH} , implying that B-types are complements with each other in producing this last effect.

5 Selection of Biases Across Periods

The distribution of types within the population and within each group are fixed during a period, but between periods, the population evolves, with more successful cognitive types replacing less successful types. I numerically solve for which population states are asymptotically stable in the replicator dynamics (Taylor and Jonker 1978). The model is loosely adapted from Herold (2012).¹⁰

Consider an infinite population with two types of players, each a fully parameterized cognitive type ϕ . A particular B-type occurs with frequency λ and R-types occur with frequency $1 - \lambda$. There is not assortative matching, so the probability of a group with x players of type B for groups of size G is distributed binomially, with

$$p(x, \lambda, G) = \Pr(b = x) = \binom{G}{x} \lambda^x (1 - \lambda)^{G-x}. \quad (8)$$

I discuss these assumptions further later in the section.

Conditional on being a biased type, a player is in a group with one additional biased type because she herself is a biased type. Let $U(\tau, \lambda) = \mathbb{E}[V(\tau, b)|\lambda]$, i.e.,

$$U(B, \lambda) = \sum_{i=0}^{b=G-1} p(i, \lambda, G-1) V(B, b+1) \quad (9)$$

$$U(R, \lambda) = \sum_{i=0}^{b=G-1} p(i, \lambda, G-1) V(R, b). \quad (10)$$

¹⁰Herold models the indirect evolution of positive and negative reciprocity in groups by integrating payoffs for fixed group compositions, analogous to (11) over binomially distributed group sizes as in (8). Herold's equivalent of $V(\tau, b)$ yields simple closed-form solutions and is better-behaved for integration, allowing him to prove that there is positive selection for reciprocity in a variety of settings instead of providing numerical results.

and the payoff difference between the cognitive types in the population is

$$U(B, \lambda) - U(R, \lambda) = \sum_{i=0}^{b=G-1} p(i, \lambda, G-1) \{ [V(B, b) - V(R, b)] + [V(B, b+1) - V(B, b)] \} \quad (11)$$

The biased type population share λ evolves according to $\dot{\lambda} = \lambda(U(B, \lambda) - U(R, \lambda))$. A state λ^* is asymptotically stable if at that point, $\dot{\lambda} = 0$ and there exists $\delta > 0$ such that for any $\lambda_0 \in (\lambda^* - \delta, \lambda^* + \delta) \cap [0, 1]$, $\lim_{t \rightarrow \infty} \lambda(t) = \lambda^*$. I focus on asymptotically stable states as predictors of the long-run outcome of evolution of biases in this population.

Which states are stable depends entirely on $U(B, \lambda) - U(R, \lambda)$. This difference depends in turn on the difference in payoffs between the two types holding group composition fixed (the first pair of terms in the brackets) and the expected marginal effect of increasing b given the distribution induced by λ (the second pair of terms in the brackets). These two effects correspond to the comparative statics approximated in Lemma 1 and Lemmas 2 and 3.

These forces typically lead to two asymptotically stable λ under both the representativeness heuristic and the availability heuristic. In general, one stable population state is for the entire population to be rational ($\lambda^* = 0$), while the corresponding monomorphic state in which the entire population is biased is unstable. The stability of the monomorphic states occur because the payoff difference $V(B, b) - V(R, b)$ is always negative and the often positive group composition effect $V(B, b+1) - V(B, b)$ is relatively small for λ close to 0 or 1. As $\lambda \rightarrow 0$ the limiting group composition effect is $V(B, 1) - V(B, 0)$, which is a smaller effect than under intermediate λ because biased players are complements with each other in increasing transition speeds. It is small as $\lambda \rightarrow 1$ for the same reason: due to these complementarities, transition speeds are typically fast enough that the marginal effect

$V(B, G) - V(B, G - 1)$ is also small.

There is generally also a second stable population state in which R-types and B-types coexist. This interior stable state exists if the weighted $V(B, b + 1) - V(B, b)$ terms in (11) are large enough for B-types to offset the payoff difference between B's and R's.

Figure 6 illustrates the relationship between $V(\tau, b)$, $U(B, \lambda) - U(R, \lambda)$, b , and λ . The left figure shows the expected per sub-period payoffs for an R-type (green) in a group with b B-types and for a B-type (blue) in a group with b other B-types. For $b < 3$ and $b > 7$, $V(R, b) > V(B, b)$, while for between 3 and 7 biased group members, $V(R, b) < V(B, b)$. The right figure shows the induced $U(R, \lambda) - U(B, \lambda)$. There are two asymptotically stable states, an all-rational state and a state with λ around 0.8.




Figure 6: Expected per-subperiod payoff by type and group composition (left) and expected per-period payoff difference between R and B by R population share λ (right). $V(R, b)$ is green and $V(B, b)$ is blue. $G = 15$, $\psi = 1/199$, $\bar{\pi}$ is normalized to 1 and $\underline{\pi} = 6/9$.

While the dynamics in Figure 6 are typical, for some parameter values only the all-rational state is stable, while for a few the interior λ is the uniquely stable population state. In the within-period model there are four parameters – ψ , ϵ , $\bar{\pi}/\underline{\pi}$, and G – that affect $V(\tau, b)$

and hence the possibility that biases are selected for.¹¹ The rest of this section describes qualitatively how these parameters affect the evolution of errors and on how the dynamics vary depending on the cognitive bias. In general, the existence of a stable interior state is most sensitive to $\bar{\pi}/\underline{\pi}$ and to ϵ ; the necessary conditions for one to exist are that ϵ be small enough and $\bar{\pi}/\underline{\pi}$ be large enough. Stable interior states are also much more likely to exist for the availability heuristic than the representativeness heuristic.

Instability of payoffs (ψ): $\psi = \frac{1-\delta\sigma}{\delta\sigma}$ measures how likely it is that the payoffs change subperiod-to-subperiod or that the period ends. Lemma 3 showed that the transition portion of the group composition effect is maximized at an intermediate level of stability, and so theoretically both great instability and great stability are not conducive to a stable state in which biased types exist. In practice, for reasonable parameter values, the existence of the polymorphic equilibrium seems more sensitive to instability than to excess stability.

Figure 7 shows λ^* for a range of expected log-times. For both heuristics, there is no stable state in which the heuristic is used for expected durations less than approximately 20 sub-periods, while expected durations as long as 100,000 sub-periods still produce a stable state with $\lambda^* > 0$.

High versus low payoffs ($\bar{\pi}/\underline{\pi}$): The ratio $\bar{\pi}/\underline{\pi}$ is a sufficient measure of the effect of coordinating payoffs on existence of polymorphic stable population states. Payoffs have two effects: first, the number of errors necessary to switch from the low payoff state to the high payoff state is proportional to $\underline{\pi}/(\underline{\pi} + \bar{\pi})$, so an increase in $\bar{\pi}$ relative to $\underline{\pi}$ increases γ_{LH} , and is in that sense a substitute for biased group members. Second, in the reduced Markov process the difference in payoffs by type and effect of group composition on matching (Lemmas 1

¹¹In addition, the intensity of the availability heuristic η must be specified. I assume $\eta = 3/10$ throughout.




Figure 7: XXX

and 2) are both proportional to the average payoff achieved, while the effect on payoffs from increases in transition speed caused by biased players is proportional to an appropriately weighted payoff difference $\bar{\pi} - \underline{\pi}$. Hence an increase in $\bar{\pi}$ has a relatively larger effect on the force that favors cognitive biases. In sum, the theoretical comparative statics are ambiguous; in practice, polymorphic states are much more likely to be stable for larger $\bar{\pi}$.

Base error rate (ϵ): ϵ increases the speed of transitions and hence is a substitute for cognitive biases. There is some maximum level of ϵ , depending on the other parameter combinations, beyond which the only stable population is an all-rational one.




Figure 8: Maximum share of biased types λ in stable population state for representativeness heuristic by group size (row) and stability (column). For each subplot, X-axis is ϵ , y-axis is $\bar{\pi}/\underline{\pi}$, and color is λ^* for given parameters.

Figures 8 and 9 show how ϵ and $\bar{\pi}/\underline{\pi}$ affect the mixed steady state for representative combinations of ψ and G . Each subfigure is a “heat map” whose color at a point $(\epsilon, \bar{\pi}/\underline{\pi})$ corresponds to the maximum stable share of biased types in the population. Each row is a different group size and each column is a different degree of instability. In general, the representativeness heuristic is stable in a population for ϵ less than around 0.09, while the

availability heuristic is stable for ϵ less than 0.15.

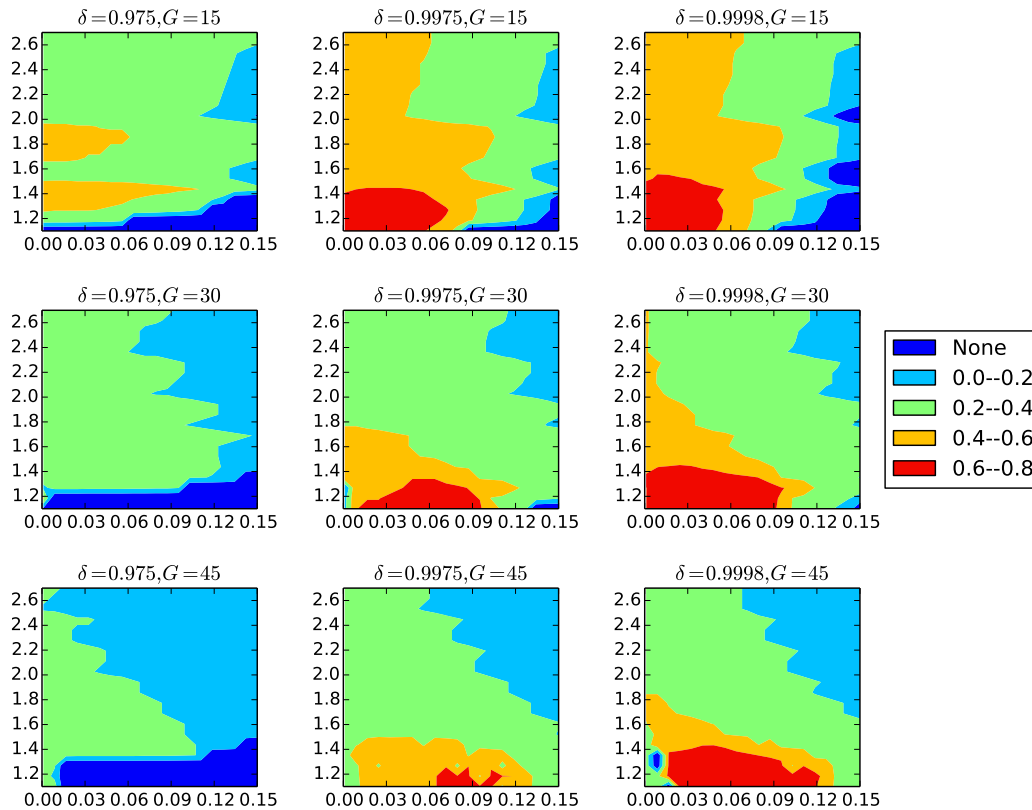


Figure 9: Maximum share of biased types λ in stable population state for availability heuristic by group size (row) and stability (column). For each subplot, X-axis is ϵ , y-axis is $\bar{\pi}/\pi$, and color is λ^* for given parameters.

Group size (G): In a prior paper, I showed that transition speeds for all-biased groups are independent of group size for large enough degrees of bias (Wood 2013). That suggests that the cognitive biases should have beneficial effects on transition speed regardless of group

size, in which case the stability of biased population states should not be very sensitive to group size. In practice, G does not affect the existence of mixed stable population states under the availability heuristic, but the stability of mixed states is more sensitive to G under the representativeness heuristic.

For both heuristics, as shown in Figure 6, for small numbers of biased types b in a group, $V(R, b) > V(B, b)$. For intermediate b this relationship reverses and marginal biased types outperform rational players, before returning to $V(R, b) > V(B, b)$ for high b . Under the representativeness heuristic, in the low- b region, biased players do significantly worse than non-biased players in a group because they coordinate much less effectively but are not prevalent enough to increase transition speed. In contrast, under the availability heuristic, while transition speeds are not increased for low b , biased players also play non-conforming strategies only slightly more frequently than rational players. Figure 10 illustrates these differences. Because group sizes are distributed binomially, the low- b payoffs of the biases affect the overall success of biased players.




Figure 10: Expected per-subperiod payoff by type and group composition for representativeness heuristic (left) and availability heuristic (right). $V(R, b)$ is green and $V(B, b)$ is blue. $G = 40$, $\bar{\pi}$ is normalized to 1, and $\underline{\pi} = 6/9$.

The three main assumptions I make in analyzing the evolution of cognitive biases are i) that there are only two cognitive types in a population, ii) that types' payoffs depend only on individual payoffs and not group payoffs, and iii) that there is no assortative matching when groups form. All of these assumptions are unfavorable for the stability of cognitive biases. (i) makes ϵ exogenous, while allowing selection for ϵ would cause lower ϵ to be present in the population because errors are generally deleterious for rational types. Because groups with large numbers of B-types collectively do much better than groups with few B-types (compare $b = 0$ and $b = 10$ in Figure 6), competition between groups would make biased types much more successful. A similar logic applies for (iii). Therefore the numeric results in this section support the the argument that in myopic populations, errors will not be i.i.d. but instead correlated in various ways.

6 Conclusion

I develop a theory of how cognitive biases affect the evolution of behavior in repeated pure coordination games in which the payoff-dominant equilibrium occasionally switches. Their stochastic nature makes the speed of behavioral adjustment especially important for groups playing the game. The importance of adjustment means that cognitive biased group members produce large positive externalities for other members of their group, because these biases cause faster movement between equilibria.

The specific biases I consider are the representativeness heuristic and the availability heuristic. These biases speed transitions to “good” equilibria because movement between equilibria require several simultaneous mistakes by group members. Cognitive biases have two beneficial effects: the errors of biased individuals are positively correlated, reducing the negative byproduct of errors on these individuals, and errors are state-dependent, with

biased individuals making more errors in “bad” states for the group, which causes faster switches to “good” states.

Characterizing the evolution of behavior in stochastic coordination environments is important because of their prevalence in modern economics, where network externalities naturally produce coordination-game-like strategic situations and where the pace of modernity naturally produces stochasticity of payoffs. While I build on a rich literature in stochastic evolutionary game theory, heretofore that literature has focused primarily on equilibrium refinements in unchanging environments.

Having characterized payoffs in stochastic coordination games where some group members exhibit biased cognition, I turn the main question of the paper around and ask why these biases might have survived if there is persistent evolutionary pressure to “make good decisions”. Using my analysis of the evolution of behavior I then look at the indirect evolution of the determinants of behavior. I show that with even weak group selective pressures, for a wide range of parameters, stable population states exist in which some members of the population are biased while others are not.

References

- Bergin, James and Barton T. Lipman**, “Evolution with State-Dependent Mutations,” *Econometrica*, 1996, *64*, 943–956.
- Bowles, Samuel**, “Group competition, reproductive leveling, and the evolution of human altruism,” *Science*, 2006, *314*, 1569–1572.
- Dekel, Eddie, Jeffrey C. Ely, and Okan Yilankaya**, “Evolution of Preferences,” *Review of Economic Studies*, 2007, *74*, 685–704.
- Ellison, Glenn**, “Basins of Attraction, Long-Run Stochastic Stability, and the Speed of Step-by-Step Evolution,” *Review of Economic Studies*, 2000, *67*, 17–45.
- Friedlin, M.I. and A.D. Wentzell**, *Random Perturbations of Dynamical Systems*, New York: Springer-Verlag, 1998.
- Güth, Werner and Menahem Yaari**, “Explaining Reciprocal Behavior in Simple Strategic Games: An Evolutionary Approach,” in Ulrich Witt, ed., *Explaining Process and Change: Approaches in Evolutionary Economics*, Ann Arbor: University of Michigan Press, 1992, pp. 23–34.
- Heifetz, Aviad and Yossia Spiegel**, “The Evolution of Biased Perceptions,” 2001. working paper.
- Heller, Yuval**, “Three Steps Ahead,” 2012. Working paper.
- Herold, Florian**, “Carrot or Stick? The Evolution of Reciprocal Preferences in a Haystack Model,” *American Economic Review*, 2012, *102*, 914–940.
- Kahneman, Daniel**, *Thinking, Fast and Slow*, New York: Farrar, Straus and Giroux, 2011.
- Kandori, Michihiro, George J. Mailath, and Rafael Rob**, “Learning, Mutation, and Long Run Equilibria in Games,” *Econometrica*, 1993, *61* (1), 29–56.
- Kreindler, Gabriel E. and H. Peyton Young**, “Fast Convergence in Evolutionary Equilibrium Selection,” *Games and Economic Behavior*, 2013, *80*, 39–67.
- Rabin, Mathew**, “Inference by Believers in the Law of Small Numbers,” *Quarterly Journal of Economics*, 2002, *117* (3), 775–816.
- Robson, Arthur J. and Fernando Vega-Redondo**, “Efficient Equilibrium Selection in Evolutionary Games with Random Matching,” *Journal of Economic Theory*, 1996, *70* (1), 65–92.

- Salomonsson, Marcus**, “Group selection: The quest for social preferences,” *Journal of Theoretical Biology*, 2010, *264*, 737–746.
- Sawa, Ryoji**, “An Analysis of Stochastic Stability in Bargaining Games with Behavioral Agents,” 2012. Unpublished manuscript.
- , “Mutation Rates and Equilibrium Selection under Stochastic Evolutionary Dynamics,” *International Journal of Game Theory*, 2012, *41*, 489–496.
- Taylor, P. and L. Jonker**, “Evolutionary Stable Strategies and Game Dynamics,” *Mathematical Biosciences*, 1978, *40*, 146–156.
- Tversky, Amos and Daniel Kahneman**, “Judgement Under Uncertainty: Heuristics and Biases,” *Science*, 1974, *185*, 1124–1131.
- Wood, Daniel H.**, “Cognitive Biases as Accelerators of Behavioral Evolution,” 2013. Working paper.
- Young, H. Peyton**, “The Evolution of Conventions,” *Econometrica*, 1993, *61* (1), 57–84.

A Proofs

A.1 Proof of Lemma 1

Proof. The proof is by decomposing M_τ into a product of simpler matrices and using the fact that trace is a linear operator (i.e., $\text{tr}(A + B) = \text{tr}(A) + \text{tr}(B)$).

Let M_ϵ be the matching matrix for a player who makes errors with probability ϵ , so

$$M_\epsilon \approx \begin{pmatrix} (1 - \epsilon) \Pr(c|H) & \epsilon(1 - \Pr(c|L)) \\ \epsilon(1 - \Pr(c|H)) & (1 - \epsilon) \Pr(c|L) \end{pmatrix}$$

M_ϵ can be written as

$$M_\epsilon = C[X_0 - \epsilon I] \tag{12}$$

where

$$C = \begin{pmatrix} \Pr(c|H) & -(1 - \Pr(c|L)) \\ -(1 - \Pr(c|H)) & \Pr(c|L) \end{pmatrix},$$

$$X_0 = \begin{pmatrix} \frac{\Pr(c|L) \Pr(c|H)}{\Pr(c|L) + \Pr(c|H) - 1} & \frac{\Pr(c|L)(1 - \Pr(c|L))}{\Pr(c|L) + \Pr(c|H) - 1} \\ \frac{\Pr(c|H)(1 - \Pr(c|H))}{\Pr(c|L) + \Pr(c|H) - 1} & \frac{\Pr(c|L) \Pr(c|H)}{\Pr(c|L) + \Pr(c|H) - 1} \end{pmatrix},$$

and I is the identity matrix. Now let $V(\epsilon)$ be the payoff to a player who makes errors with probability ϵ .

$$V(\epsilon) = \text{tr}(\Pi M_\epsilon N) = \text{tr}(\Pi C X_0 N) - \epsilon \text{tr}(\Pi C N)$$

The payoff difference between R and B players is

$$V(R) - V(B) = (\epsilon_B - \epsilon_R) \text{tr}(\Pi C N)$$

and

$$V(0) - v(R) = \epsilon_R \text{tr}(\Pi CN)$$

so the payoff difference can be expressed as

$$V(R) - V(B) = \left(\frac{\epsilon_B - \epsilon_R}{\epsilon_R} \right) [v(0) - v(R)]$$

Since $V(0) \approx (1 + \epsilon_R)v(R)$ for small ϵ , the claim holds. □

A.2 Proof of Lemma 2

Proof. $\text{tr}(\Pi(\partial M/\partial b)N) = \text{tr}(N\Pi(\partial M/\partial b))$ and

$$\frac{\partial M_\epsilon}{\partial b} \approx \begin{pmatrix} (1 - \epsilon) \frac{\partial \Pr(c|H)}{\partial b} & -\epsilon \frac{\partial \Pr(c|L)}{\partial b} \\ -\epsilon \frac{\partial \Pr(c|H)}{\partial b} & (1 - \epsilon) \frac{\partial \Pr(c|L)}{\partial b} \end{pmatrix}$$

so $\partial V/\partial b$ is

$$\begin{aligned} \partial V/\partial b &= \left(\frac{\psi + 2\gamma_{LH}}{\psi + \gamma_{HL} + \gamma_{LH}} \right) [(1 - \epsilon)\bar{\pi} - \epsilon\underline{\pi}] \left(\frac{\partial \Pr(c|H)}{\partial b} \right) \\ &+ \left(\frac{\psi + 2\gamma_{HL}}{\psi + \gamma_{HL} + \gamma_{LH}} \right) [(1 - \epsilon)\underline{\pi} - \epsilon\bar{\pi}] \left(\frac{\partial \Pr(c|L)}{\partial b} \right). \end{aligned}$$

This is smaller in magnitude for higher ϵ , establishing claim (ii).

Claim (i) follows because

$$\Pr(c|S) \approx \frac{(1 - \epsilon_{RS})(G - b) + (1 - \epsilon_{BS})b}{G},$$

so

$$\frac{\partial \Pr(c|S)}{\partial b} = -\frac{\epsilon_{BS} - \epsilon_{RS}}{G}. \quad \square$$

B Transition probabilities

Recall that γ_{HL} is the per-subperiod probability of switching from the high-paying strategy's basin of attraction to the low-paying strategy's basin of attraction given that the system is in the high-paying strategy's basin at the start of the sub-period, and γ_{LH} is the opposite switch. Because more errors are necessary to make the former switch than the later, $\gamma_{HL} < \gamma_{LH}$. This section derives expressions for these probabilities.

B.1 Homogenous groups with state-independent errors

Consider an homogenous group consisting of players who make errors with probability ϵ , where ϵ is low. Let x_t be the fraction of players that choose non-conforming strategies in sub-period t . Given the homogeneity of the group, x_t is approximately distributed $B(G, \epsilon)$.

In $t+1$ the state will be in L's basin of attraction if x_t is above some cutoff. In particular, let p be the fraction of group members playing L for which a player is indifferent between L and H. Then

$$\begin{aligned} U(L, p) &= U(H, p) \\ p\underline{\pi} &= (1-p)\bar{\pi} \\ p &= \frac{\bar{\pi}}{\bar{\pi} + \underline{\pi}} \end{aligned}$$

The probability of transition is

$$\gamma_{HL} = \Pr \left(x_t \geq \frac{\bar{\pi}G}{\bar{\pi} + \underline{\pi}} \right).$$

An analogous calculation shows

$$\gamma_{LH} = \Pr \left(x_t \geq \frac{\underline{\pi}G}{\bar{\pi} + \underline{\pi}} \right).$$

Ignoring integer complications, for ϵ low,

$$\gamma_{HL} = \binom{G}{pG} \epsilon^{pG} (1 - \epsilon)^{G-pG} + \text{h.o.t} \approx \binom{G}{pG} \epsilon^{pG} (1 - \epsilon)^{G-pG} \quad (13)$$

and

$$\gamma_{LH} = \binom{G}{(1-p)G} \epsilon^{G-pG} (1 - \epsilon)^{pG} + \text{h.o.t} \approx \binom{G}{(1-p)G} \epsilon^{G-pG} (1 - \epsilon)^{pG} \quad (14)$$

and note that the binomial coefficients are equal, so

$$\frac{\gamma_{HL}}{\gamma_{LH}} = \frac{\epsilon^{pG} (1 - \epsilon)^{G-pG}}{\epsilon^{G-pG} (1 - \epsilon)^{pG}} = \left(\frac{\epsilon}{1 - \epsilon} \right)^{(2p-1)G}. \quad (15)$$

Using equation (15), the fraction of time spent in H in the long-run can be calculated as

$$\frac{\gamma_{LH}}{\gamma_{LH} + \gamma_{HL}} = \frac{(1 - \epsilon)^{(2p-1)G}}{(1 - \epsilon)^{(2p-1)G} + \epsilon^{(2p-1)G}}. \quad (16)$$

For $G > 1/(2p - 1)$, the exponents in equation (16) are greater than one, implying $\gamma_{LH}/(\gamma_{LH} + \gamma_{HL}) > 1 - \epsilon$. It is thus clear that

Lemma B1. *For large G and small ϵ , $\gamma_{LH}/(\gamma_{LH} + \gamma_{HL}) \approx 1$.*

This lemma holds for state-dependent and heterogenous error probabilities as well.

B.2 State-dependent errors

Now let the probability of not conforming in state $S \in \mathcal{S}$ be ϵ_S . Ignoring integer complications,

$$\begin{aligned}\gamma_{HL} &\approx \binom{G}{pG} \epsilon_H^{pG} (1 - \epsilon_H)^{G-pG} \\ \gamma_{LH} &\approx \binom{G}{(1-p)G} \epsilon_L^{G-pG} (1 - \epsilon_L)^{pG}\end{aligned}$$

so

$$\frac{\gamma_{HL}}{\gamma_{LH}} = \frac{\epsilon_H^{pG} (1 - \epsilon_H)^{G-pG}}{\epsilon_L^{G-pG} (1 - \epsilon_L)^{pG}} = \left[\left(\frac{\epsilon_H}{1 - \epsilon_L} \right)^p \left(\frac{1 - \epsilon_H}{\epsilon_L} \right)^{(1-p)} \right]^G.$$

and

$$\frac{\gamma_{LH}}{\gamma_{LH} + \gamma_{HL}} = \frac{\epsilon_L^{G-pG} (1 - \epsilon_L)^{pG}}{\epsilon_H^{pG} (1 - \epsilon_H)^{G-pG} + \epsilon_L^{G-pG} (1 - \epsilon_L)^{pG}}. \quad (17)$$

B.3 Heterogenous groups

As in the previous sections, I only consider combinations of errors by B's and R's that produce the marginal number of errors necessary to switch states. Exact transition probabilities in this case are sums of products of two binomial PMF's, which are hard to work with analytically. Instead, I approximate the error rates through an assumption that is in the same spirit as using binomial PMF's instead of hypergeometric PMF's.

Consider a group made out of g individuals selected from a large population of four types: conforming R-types, conforming B-types, non-conforming R-types, and non-conforming B-types. Conditional on whether a player is conforming or non-conforming, his cognitive type

is unimportant for transitions in the reduced Markov process, so this population can in turn be thought of as an infinite population of two types:

- non-conforming players (with probability $\frac{G\epsilon_{RS}+(b-G)\epsilon_{BS}}{G}$)
- conforming players (with probability $\frac{1-G\epsilon_{RS}-(b-G)\epsilon_{BS}}{G}$)

where the frequencies are for a population in state S with b/G fraction of the population is a B-type.

Then let

$$\tilde{\epsilon}_S \equiv \epsilon_{RS} + \frac{b}{G}(\epsilon_{BS} - \epsilon_{RS})$$

be the effective probability of drawing a non-conformist, and

$$\gamma_{HL} \approx \binom{G}{pG} (\tilde{\epsilon}_H)^{pG} (1 - \tilde{\epsilon}_H)^{G-pG} \quad (18)$$

and

$$\gamma_{LH} \approx \binom{G}{(1-p)G} (\tilde{\epsilon}_L)^{G-pG} (1 - \tilde{\epsilon}_L)^{pG} \quad (19)$$

follow immediately.

The following lemma is also immediate:

Lemma B2. *The change in γ_{LH} as the group composition changes is approximately*

$$\frac{\partial \gamma_{LH}}{\partial b} = (\epsilon_{BL} - \epsilon_{RL}) \left(\frac{(1-p)}{\tilde{\epsilon}_L} - \frac{p}{1 - \tilde{\epsilon}_L} \right) \gamma_{LH}(b).$$

Because Lemma B2 depends on the approximation in (19), the expression will be most accurate for ϵ_{BL} close in magnitude to ϵ_R . This approximation is only used in discussing Lemma 3.