Cognitive Biases in Stochastic Coordination Games and Their Evolution

Daniel H. Wood*

4/15/2013

Preliminary and Incomplete

Abstract

I model the evolution of behavior in small groups whose members play pairwise 2x2 pure coordination games. Players follow a best-response dynamic with errors which are generated by a fixed cognitive bias, such as the representativeness heuristic. Mistakes in strategy choice impair successful coordination, but they also cause the group to occasionally switch between equilibria. If which equilibrium generates higher payoffs switches at a slower rate than the strategy-updating process, then errors create a positive externality for other players because the group returns to high-payoff equilibria faster. I characterize group members' payoffs generated by this dynamic for several error-generating biases, including limited cognition, the representativeness heuristic, the false-consensus effect, and loss aversion. Behavioral biases produce larger externalities relative to simply higher error rates. I then analyze 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.

*Department of Economics, Clemson University, 228 Sirrine 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-knowm that 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 a variety of cognitive biases – the representativeness heuristic and the false consensus effect – 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 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 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. 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, or biased, who form their beliefs in some other way. For example, they may sample the strategies of a few members of the group and over-extrapolate group behavior from that sample (i.e., use the representativeness heuristic). Individual behavior and the prevalence of cognitive biases are determined by separate evolutionary processes, one embedded within the other.

Agents interact within fixed groups during a period. 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 have a base rate of errors and might also make errors because they are a biased-type player.

One part of this paper characterizes the expected payoff of a player i during a period. 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 i successful 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 in the population will change over time. I model the evolution of biases over a series of periods as occuring both through selection within groups, where more successful members replace less successful ones, and also across groups, where more successful groups replace less successful ones. The frequency of biased players in the population at period t is a deterministic difference equation that is produced by these two forces. Within a group, rational players always do better than biased players, so intra-group evolution favors rational types. However, inter-group evolution generally favors biased types, because when comparing different groups, groups with a larger number of biased agents do better on average than groups with fewer biased agents, because they spend more time on average in the higher-payoff equilibrium.

Under these dynamics, intra-group and inter-group selection together often support 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. 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 equilirium is unattractive relative to the 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.

In addition to the representativeness heuristic, I consider the false consensus effect, and loss aversion. Relative to the representativeness heuristic, the false consensus effect produces smaller differences in average error rates between good and bad equilibria, but more variation in error rates across periods, which is particularly beneficial for transitions between equilibria. Loss aversion, in contrast to the other biases, reduces transition speeds and so is an example of a bias with only negative effects in stochastic coordination games. I consider several biases in order to make the general point that cognitive biases are usually beneficial in improving coordination dynamics because correlated errors usually produce state-dependent errors, clusters of errors, or both.

The paper is organized as follows: Section 2 provides a brief example of within-group behavioral evolution and relates it to the general effects of cognitive biases in a stochastic setting. Section 3 briefly discusses applications of this paper and the paper's relationship to several literatures; Section 4 and Section 5 then develop general within-period models of behavioral evolution and across-period models of cognitive bias evolution, respectively; Section 6 analyzes the particular dynamics stemming from the representativeness heuristic, availability heuristic, false consensus effect, and loss aversion; Section 7 argues in depth that the within-period model can be applied to real-world phenomena with network externalities; and finally Section 8 concludes.

2 Motivating Example

Consider a group of 10 players playing the coordination game on the left 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:

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

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

Let a_t be the number of A's at t. a_t is distributed binomially:

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

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.[1]

[1] Kandori et al. (1993) and Young (1993) were the first papers to characterize the behavior of stochastic evolutionary processes of this nature.

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 A equilibrium will be more likely at each t. For example, a group with $\epsilon = 1/10$ number of periods in B of approximately 80 periods, but increasing ϵ to $\epsilon = 3/20$ would reduce the expected time to approximately 20 periods.[2] 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 a stochastic environment, switches akin to that from Game II to Game I frequently leave the group coordinating on a payoff-dominated strategy, so higher error rates are especially beneficial in stochastic environments.

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 higher error rates in the B than in the A equilibrium, which magnifies the positive externality from the errors.

For example, agents using the representativeness heuristic treat a small sample as more informative than it actually is. If players have an underlying ϵ chance of making an error and a $1-\epsilon$ chance of sampling two strategies from $t-1$ and best responding to that sample, then B to A transitions 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.

Representativeness allows errors to "build up" over several periods. For example, in the

[2]In 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.

B equilibrium, with rational players who sample the entire population, each player has an ϵ chance of playing A, so $\Pr(a_t \geq 4 | a_t < 4) = \binom{10}{4}\varepsilon^4(1-\varepsilon)^6 \approx 210\epsilon^4$, while for a group of players all following the representativeness heuristic $\Pr(a_t > a_{t-1} | a_{t-1} < 4) \geq 1/2$: consider the following Markov transition matrix for the population with representativeness bias. Let T be a 5×5 matrix where $t_{ij} \equiv \Pr(a_{t+1} = j-1 | a_t = i-1)$ for $i, j \in \{0, 1, \ldots, 4\}$ and let the absorbing state 5 index any state $a_t \geq 4$. Then under the heuristic, T is

$$T = \begin{bmatrix} 0.35 & 0.39 & 0.19 & 0.06 & 0.01 \\ 0.04 & 0.16 & 0.26 & 0.26 & 0.28 \\ 0 & 0.03 & 0.10 & 0.19 & 0.68 \\ 0 & 0 & 0.02 & 0.07 & 0.91 \\ 0 & 0 & 0 & 0 & 1 \end{bmatrix}$$

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 has a greater than 50% chance of having left it after 4 rounds.

In contrast, A to B transitions require both samples to be B, 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 minor. While the heuristic does speed up transition from A to B, the expected number of rounds spent in the A equilibrium before leaving it is 2640 for $\epsilon = 1/10$ and group size 10.

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 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. Compared an iid error process, errors via the representativeness heuristic increase the speed of B to A transitions significantly without reducing the frequency of successful coordination in the good A equilibrium much.

Error-generating processes in which errors feed back over time are difficult to tractably analyze, so I analyze more tractable cases in which errors are uncorrelated over time, but for which analytic expressions for players' expected payoffs are attainable. This provides a lower bound on how beneficial particular cognitive biases are. In addition, I analyze the more general case of error correlation over time numerically (*work in progress*).

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. A modern example of an environment where stochastic coordination is important could be evolving internet technologies with network externalities such as Facebook, MySpace, and Google+. Here players are repeatedly deciding which site to post to and read. Friends who choose the same site share exciting life details and embarrasing pictures, generating utility for them. The arrival of new social networking websites can be seen as adding new, possibly superior, forms of coordination. More generally, 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 several cognitive biases in settings like social networking, technology adoption, or fashion.

A second contribution of this paper is to apply the analysis of how biases effect the evolution of behavior to understanding 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, while Boehm (1999) describes the environment more generally.[3]

Competition between groups has most frequently been used to explain altruism or as a source of equilibrium selection. Altruism benefits a group collectively but harms individual altruists, so competition between groups can resolve the puzzle of why altruism would persist. In this paper, biased individuals in this paper are analogous to altruistic individuals.[4] Salomonsson (2010) surveys and classifies various forms of group selection; in his nomenclature the group selection across periods in my model 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 some 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. However, Herold's group selection model does not include explicit group-level selection; instead he simply looks at average payoffs as the fraction of players engaging in punishment of the selfish or rewarding of the altruistic changes.

In my setup, competition between groups sometimes favors the survival of some individuals with higher error rates than other individuals. These errors are uncorrelated, but

[3]The formal group selection model that I develop is agnostic to the unit of selection and hence the evolution could be re viewed as cultural rather than biological. Stable individual behavior across periods would represent different mental habits that were more or less prone to biased thinking, instead of genetic variation. I believe selection at the genetic level is more convincing for cognitive biases than selection at the cultural level.

[4]One 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.

cognitive biases in general lead the biased to produce more errors in some situations and fewer in others. Therefore selection is much stronger when extended to the particular biases. 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. 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 equilirium is unattractive relative to the 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.

This paper is related to several literatures in evolutionary game theory. Most directly, group members' behavior evolves through a simple version of KMR-style dynamics (Kandori et al. 1993), a classic model in stochastic evolutionary game theory.[5] A well-known critique of stochastic evolutionary models is that they take implausibly long amounts of time to reach the stable equilibrium (Ellison 2000). Researchers have proposed spatial models as a way to make the speed of the evolution more reasonable (Ellison (1993); see also Young (2011)). Cognitive biases are another, complementary reason for why behavior may evolve more quickly than simple stochastic evolutionary models predict.

A few authors have investigated how the nature of errors affects 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

[5]Note that "stochastic evolutionary game theory" is about evolutionary processes with stochastic changes to behavior; a "stochastic coordination game" in contrast is one where the game payoffs change.

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.

The paper's analysis of the evolution of congitive biases uses the indirect evolutionary approach pioneered by Güth and Yaari (1992). The majority of the indirect evolution literature explains the evolution of preferences through the 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 heterogenous 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 of backwards induction and others engage in three rounds can be a neutrally stable state. This result suggests that there is some evolutionary basis for the level-K models used in behavioral game 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

Consider a coordination game as depicted in Figure 2 played repeatedly by members of a population. Which Nash equilibrium is payoff-maximizing for the population depends on the relative values of π_A and π_B, which change over time.

There is an infinite population of players. Each player i is endowed with a cognitive type that determines how he chooses between actions. Players are of two types, low rationality / biased types (B-types) and high rationality (R-types). R-types come closer to the Bayesian normative standard of rationality than B-types. I consider several different B-types, corresponding to different cognitive biases identified in the behavioral economics literature, while keeping R-types' behavior constant. Cognitive types will be described generally later in this section and fully specified in Section 6.

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

Figure 2: 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 assigned 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.

Within a period, a player's type is fixed. It influences how the player's strategy evolves.

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.

Players update their strategies simultaneously. Let $s_{-i,t}$ be a vector of the strategies of all members of group except for i. Let $\phi_\tau : S^{G-1} \to S$ be an function describing a cognitive type τ's strategy choice in response to other group member's strategies. For all types, with probability $(1-\epsilon)$ a player of type τ uses the type-specific update function, and with probability ϵ the player chooses a random strategy.[6]

Type-R players have $\phi_r(s) = BR(s)$, i.e, they optimally respond to the other players' strategies. The simplest type-B choice function is a noisier version of the type-R function, in which these players choose a random strategy with probability $\frac{\eta}{1-\epsilon}$ and the best response with probability $\frac{1-\epsilon-\eta}{1-\epsilon}$, so that the overall probability that they best respond is $1-\eta-\epsilon$. I

[6]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.

will call these type-B's "baseline type-B's".

After strategy updating has occured, 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 expected period payoffs using this Markov process. For that reason, I reduce this Markov process to one with two states:[7]

- **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 $\pi_B =$

[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 occured 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.

$\bar{\pi} > \pi_A = \underline{\pi}$ or the group population is in the basin of attraction of B and $\pi_A = \bar{\pi} > \pi_B = \underline{\pi}$

For the baseline type-B specification and the R-specification, none of the information thrown away by this simpler representation is needed for ϕ_B or ϕ_R.

Let $\hat{\epsilon}(\tau, H)$ 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 $\hat{\epsilon}(\tau, L)$ be defined analogously for the low payoff strategy. For the baseline type-B specification, $\hat{\epsilon}(B, H) = \hat{\epsilon}(B, L)$. In such cases I will sometimes write $\hat{\epsilon}(B)$ for the type-B's error rates.

Then the updating rule is

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

For rational types, $\hat{\epsilon}(R) = \epsilon/2$ always, while for baseline biased types, $\hat{\epsilon}(B, H) = \hat{\epsilon}(B, L) = (\epsilon + \eta)/2$.

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 rate ϵ. The transition probabilities of the reduced Markov process are

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

if the payoff values do not update in subperiod t. Because the updating rule for baseline type-B and for type-R players only depends on which basin of attraction s_{t-1} is in, γ_{HL} and γ_{LH} can expressed as functions of $\hat{\epsilon}(B, L)$, $\hat{\epsilon}(B, H)$, and $\hat{\epsilon}(R)$. Expressions for γ_{HL} and γ_{LH} are derived in Appendix B.

The reduced Markov process is equivalent to the general Markov process if error rates only depend on which basin of attraction the group is in and which type each player is. Often the ϕ_B functions for cognitive biases will not fully satisfy this requirement. In those cases, I can use different choices for $\hat{\epsilon}(B, H)$ and $\hat{\epsilon}(B, L)$ to provide upper and lower bounds on the average payoffs, and supplement the bounds with numeric calculations. These problems will be discussed further in Section 6.

4.2 Average Utility Earned During A Period

This section approximates the utility player i earns through participation in a single period supergame. First we calculate the expected payoff of a player for a sequence of subperiods in which no payoff-switch occurs. Using this expression, we 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 $N = I + Q + Q^2 + \cdots = (I-Q)^{-1}$ so

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

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

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

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

$$M_\tau = \begin{pmatrix} \Pr(u_i = \bar{\pi} \text{ in } H) & \Pr(u_i = \bar{\pi} \text{ in } L) \\ \Pr(u_i = \underline{\pi} \text{ in } H) & \Pr(u_i = \underline{\pi} \text{ in } L) \end{pmatrix} \tag{2}$$

for i such that $\tau = \tau(i)$. These probabilities are the chances of achieving conforming or non-conforming coordination in a given state, given i's type and the distribution of types in the group.

For any set of induced error rates, with a group of size G with r type-R players and b baseline type-B players, these matrices are

$$M_\tau = \begin{pmatrix} (1-\hat{\epsilon}(\tau))\left(\frac{b(1-\hat{\epsilon}(B))}{G-1} + \frac{(r-1)(1-\hat{\epsilon}(R))}{G-1}\right) & \hat{\epsilon}(\tau)\left(\frac{b\hat{\epsilon}(B)}{G-1} + \frac{(r-1)\hat{\epsilon}(R)}{G-1}\right) \\ \hat{\epsilon}(\tau)\left(\frac{b\hat{\epsilon}(B)}{G-1} + \frac{(r-1)\hat{\epsilon}(R)}{G-1}\right) & (1-\hat{\epsilon}(\tau))\left(\frac{b(1-\hat{\epsilon}(B))}{G-1} + \frac{(r-1)(1-\hat{\epsilon}(R))}{G-1}\right) \end{pmatrix} \tag{3}$$

Finally, let

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

Then the matrix $\Pi M_\tau N$ contains the expected payoffs for every initial state and match probability. In particular $(\Pi M_\tau N)_{ii}$ is the expected payoff for from a sequence with constant payoffs with initial state i. Weighting the diagonal elements of the matrix by the expected of payoff resets will therefore give the average payoff over a period.

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.[8]

Theorem 1 collects these results.[9]

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*

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

where M_τ and N are functions of τ, G, b, and ϕ.

4.3 Individual and Group Effects on Payoffs

This section precisely states the relationships between type B and type R individuals and group compositions. Player i's type influences $V(\tau)$ 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. In particular, a compositional change in the number of B-types in the group, b, holding i's type fixed, changes i's payoffs by

$$\frac{\partial V(\tau)}{\partial b} = \mathrm{tr}(\Pi(\partial M_\tau/\partial b)N) + \mathrm{tr}(\Pi M_\tau(\partial N/\partial b)) \quad (4)$$

The individual effect of being a B-type on match rate is always deleterious, but the net compositional effect is often beneficial, because the second RHS term is positive and for

[8] Each 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.

[9] It may be useful to note that the matrix N includes an additional common factor $1/(1-\delta\sigma)$.

many parameter values is larger than the first RHS 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.

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. XXX State-dependence in theory could reverse this relationship if B-types achieve non-conforming coordination frequently enough in L states. However, because $\hat{\epsilon}(B, L)$ is proportional to ϵ, for low ϵ this possibility will not occur. Assume that ϵ is low enough to rule out $V(B) > V(R)$ within groups.

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, $\frac{\mathrm{tr}(\Pi(\partial M)N)}{\partial b}$, is*

- *i) proportional to $\bar{\epsilon}/G$, where $\bar{\epsilon}$ is a convex combination of $\hat{\epsilon}(B, L)-\hat{\epsilon}(R, L)$ and $\hat{\epsilon}(B, H)-\hat{\epsilon}(R, H)$, 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 A1, $\gamma_{HL} \approx 0$ for ϵ small, in which case

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

so

$$V(\epsilon) \propto \mathrm{tr}(\Pi M N) = (\bar{\pi} m_{11} + \underline{\pi} m_{21}) \left(\frac{\psi + 2\gamma_{LH}}{\psi + \gamma_{LH}} \right) + (\bar{\pi} m_{12} + \underline{\pi} m_{22}) \left(\frac{\psi}{\psi + \gamma_{LH}} \right). \quad (5)$$

The parenthetical ratios including ψ are the fraction of time spent in H and in L, respectively, and they are weighted by the expected payoff in each state. Differentiating (5) with respect to b gives, through its effect on γ_{LH},

$$\frac{\mathrm{tr}(\Pi M (\partial N))}{\partial b} = \left(\frac{\partial \gamma_{LH}}{\partial b} \right) \left(\frac{\psi}{(\psi + \gamma_{LH})^2} \right) [(\bar{\pi} m_{11} + \underline{\pi} m_{21}) - (\bar{\pi} m_{12} + \underline{\pi} m_{22})] \quad (6)$$

Equation (6) shows 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. $\lim_{\delta\sigma \to 0} N = I$ 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. This reflects that if the system is too stable, most of the time is spent in the long-run steady state, which is H. γ_{LH}'s beneficial effect of speeding transitions to H becomes small relative to overall payoffs.

Only for intermediate ψ are B's useful. It is easy to see that the intermediate value $\psi = \gamma_{LH}$ maximizes the expression in equation (6) for a given γ_{LH}. Combining that fact with values for $\partial\gamma_{LH}/\partial b$ (derived in Appendix B) produces a lower bound on the effect of b through N for ψ approaching $\gamma_{LH}(b = G)$.

Lemma 3.

$$\lim_{\delta\sigma \to \frac{1}{1+\gamma_{LH}(b)}} \frac{\text{tr}(\Pi M \partial N)}{\partial b} > \left(\frac{\hat{\epsilon}(B,L) - \hat{\epsilon}(R,L)}{G} \right) \left(\frac{(1-p)G}{\hat{\epsilon}(B,L)} - \frac{pG}{1-\hat{\epsilon}(B,L)} \right) * \left(\frac{b}{G} + \frac{G-b}{G} \left(\frac{\hat{\epsilon}(R,L)}{\hat{\epsilon}(B,L)} \right) \right)^{G-pG} \left(\frac{\bar{\pi} m_{11} + \underline{\pi} m_{21} - \bar{\pi} m_{12} - \underline{\pi} m_{22}}{4} \right)$$

Proof. In Appendix A. □

The lower bound in Lemma 3 is proportional to the difference in error rates $(\hat{\epsilon}(B,L) - \hat{\epsilon}(R,L))$, as is the effect of b through M. However, the bound in Lemma 3 also contains a factor $1/\hat{\epsilon}(B,L)$, so that for ϵ low enough, the bound can be made large.

5 Selection of Biases Across Periods

I model evolution across periods as a non-linear difference equation and use two methods for evaluating the evolutionary equilibria of this process: a formal model of group selection (this section) and numerical calibrations (Section 6). The model illustrates the general conditions under which type B individuals can survive in the long run.

The distribution of types within the population and within each group are fixed during a period, but between periods, more successful types will replace less successful types. The frequency of each type determines how likely groups of different compositions are. Let the fraction of type-B players be λ. The probability of a group with b players of type B is distributed binomially, so

$$\Pr(b = x) = \binom{G}{x} \lambda^x (1-\lambda)^{G-x}. \tag{7}$$

There are an infinite number of groups. Let $F(x) = \Pr(b_j < x)$ be the cumulative distribution function for the number of type B's in each group.

Definition 1. For an evolutionary process $\lambda_{t+1} = f(\lambda_t)$, λ^* is an evolutionary equilibrium if

i) $\lambda^* = f(\lambda^*)$, and

ii) there exists $\delta > 0$ such that for any $\lambda_0 \in (\lambda^* - \delta, \lambda^* + \delta)$,

$$\lim_{t \to \infty} \lambda_t = \lambda^*.$$

Assume that λ is subject to two evolutionary forces: first, within a group, more successful individuals are more likely to reproduce; second, across groups, less successful groups are more likely to die off. Consider the following stylized dynamic. Each period an infinite number of groups of size G are formed with individual i in group $g(i)$ and compositions of R- and B-types as given by equation (7) above. At the end of the period, for each group, sometimes an individual with the minimum within-group payoff is replaced by a copy of the individual with the maximum within-group payoff. These intra-group replacements happen with probability α, drawn independently for each group. Let λ^{il} and λ^{ih} be the expected frequency with which B-types earn the minimum and maximum utilities in a group, respectively.

Then groups are ranked from lowest to highest average payoff. For every group in the bottom ν percent, with probability β, they are replaced by a copy of a random group in the top ν percent. Let $\lambda^{gh} = \mathrm{E}[b/G|F(b) > 1 - \nu]$ and $\lambda^{gl} = E[b/G|F(b) < \nu]$.

The frequency of type B individuals evolves as

$$\begin{aligned}
\lambda_{t+1} &= \lambda_t + \alpha(\lambda_t^{ih} - \lambda_t^{il}) + \beta\nu G(\lambda_t^{gh} - \lambda_t^{gl}) \\
\Delta\lambda_t &= \alpha(\lambda_t^{ih} - \lambda_t^{il}) + \beta\nu G(\lambda_t^{gh} - \lambda_t^{gl}) \\
\Delta^2\lambda_t &= \Delta\lambda_t - \Delta\lambda_{t-1}
\end{aligned}$$

A frequency λ is an evolutionary equilibrium if i) $\Delta\lambda = 0$ and $\Delta^2\lambda < 0$ (R's and B's are doing equally well and λ is stable), ii) $\Delta\lambda > 0$ and $\lambda = 1$ (all players are B and B's do better than R's given this), or iii) $\Delta\lambda < 0$ and $\lambda = 0$ (all players are R and R's do better than B's given this).

I will show that an equilibrium of type (iii) always exists, a type (ii) equilibrium never exists, while an equiibrium of type (i) exists for some parameter values. The characterizations of the M and N matrices in Section 4.3 are useful for these results.

Lemmas 2 and 3 allow the composition effect $d\,\mathrm{tr}(\Pi MN)/db$ to be positively signed if the game is neither too stable nor too unstable and the error rate is small enough:

Theorem 2. *If $\bar{V}(b) \equiv (bV(B) + (G-b)V(R))/G$ is the average utility within a group, then for $\delta\sigma$ close enough to $1/(1+\gamma_{LH}(g))$ and for small enough ϵ, $d\bar{V}(b)/db > 0$.*

Proof. In Appendix A. □

Theorem 3. *For any within-group selection rate α and between-group selection rate β, $\lambda = 0$ is an evolutionary equilibrium and $\lambda = 1$ is not an evolutionary equilibrium.*

In the limit as λ approaches 0 or 1, $\lambda^{gh} = \lambda^{gl}$, so $\Delta\lambda = -\alpha < 0$.

Theorem 4. *If $\bar{v}(b)$ is increasing in b, then for any between-group selection rate $\beta > 0$ and $\nu > 0$, there is an $\bar{\alpha}$ such that for any low enough within-group selection rate ($\alpha \in (0, \bar{\alpha})$) there is an evolutionary equilibrium for some $\lambda \in (1/2, 1)$.*

Theorem 2 provides a sufficient condition for the within-period conditions of Theorem 4 to be satisfied. If group selection is strong enough relative to within-group selection and the type B's produce a strong enough positive exterality, there are two evolutionary equilibria. One is a mixed population of R's and B's and the other is entirely R's. If group selection is weak or there are no positive externalities, then B's are always selected against. The reason that the equilibrium with more B's is polymorphic is that group selection requires variance in group composition to positively select B's.

6 Cognitive Biases

Conditional on group size and composition, each cognitive type induces an effective probability of conforming in state $S \in \mathcal{S}$, and transition probabilities between states in $\mathcal{S}$. For every ϕ_B type-B updating rule, I choose ϕ_B so that it nests ϕ_R as a special case, with a parameter η measuring the divergence between ϕ_B and ϕ_R and $\phi_B = \phi_R$ at $\eta = 0$ (as in the baseline case).

6.1 Baseline: Independent Errors

The baseline B specification does not introduce a directional bias to B behavior, but increases their variability relative to R's. In the baseline case,

$$\hat{\epsilon}(B, H) = \hat{\epsilon}(B, L) = (\epsilon + \eta)/2 > \epsilon/2 = \hat{\epsilon}(R, H) = \hat{\epsilon}(R, L),$$

so errors among B's are neither correlated nor state-dependent. In this section I parameterize the model and solve numerically for the type-specific and average payoffs for each number of type-B group members $b \in \{0, \dots, G\}$. The goals of this exercise are 1) to compare the strengths of intra- and inter-group selection, as measured by how payoffs vary within and across groups, and 2) to explore the comparative statics of the model further. For the parameterization I set $G = 10, \delta = 0.999, \sigma = 1, \hat{\epsilon}(R) = 0.01$, $\hat{\epsilon}(B) = 0.02, \bar{\pi} = 3, \underline{\pi} = 1$.

Figure 3 shows the expected period payoff for each type as a function of group composition while Figure 4 shows the expected group average payoff for each starting state as a function of group composition.

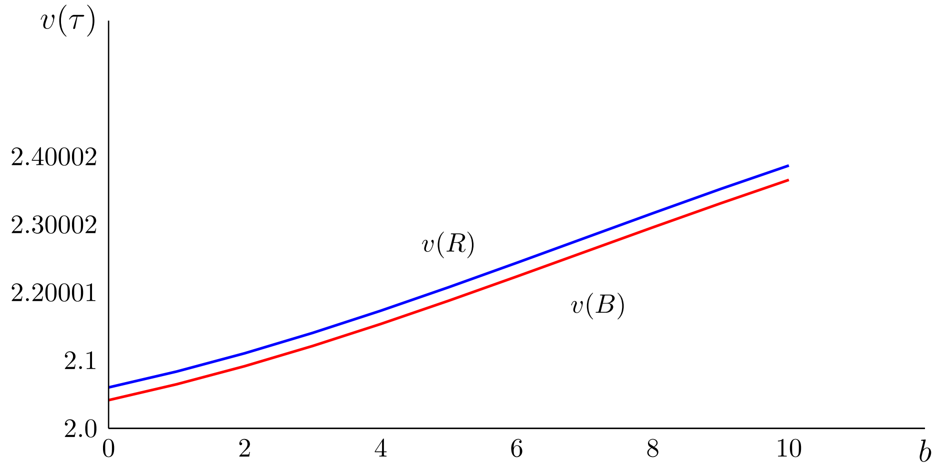

Figure 3: Average Payoff for Each Cognitive Type in Group With b Type-B's. Red line depicts $V(B)$, blue line depicts $V(R)$. $G = 10, \delta = 0.999, \hat{\epsilon}(R) = 0.01, \hat{\epsilon}(B) = 0.02$

Figure 3 shows the effect of M_B versus M_R, because the effect of group composition on N is held constant for fixed b. B's have slightly lower period payoffs than R's, because their higher induced error rates lead them to be slightly less effective at coordination in both states. The difference $v(B) - v(R)$ is approximately constant, reflecting the structure of M that lemmas 1 and 2 explored. The difference is also small because B's are only one percent more likely to make mistakes than A's under Case I.

Figure 3 shows that the within-group selective pressure against type B's is small because it is proportional to $\hat{\epsilon}(B, S) - \hat{\epsilon}(R, S)$. The across-group selective pressure for B's is larger.

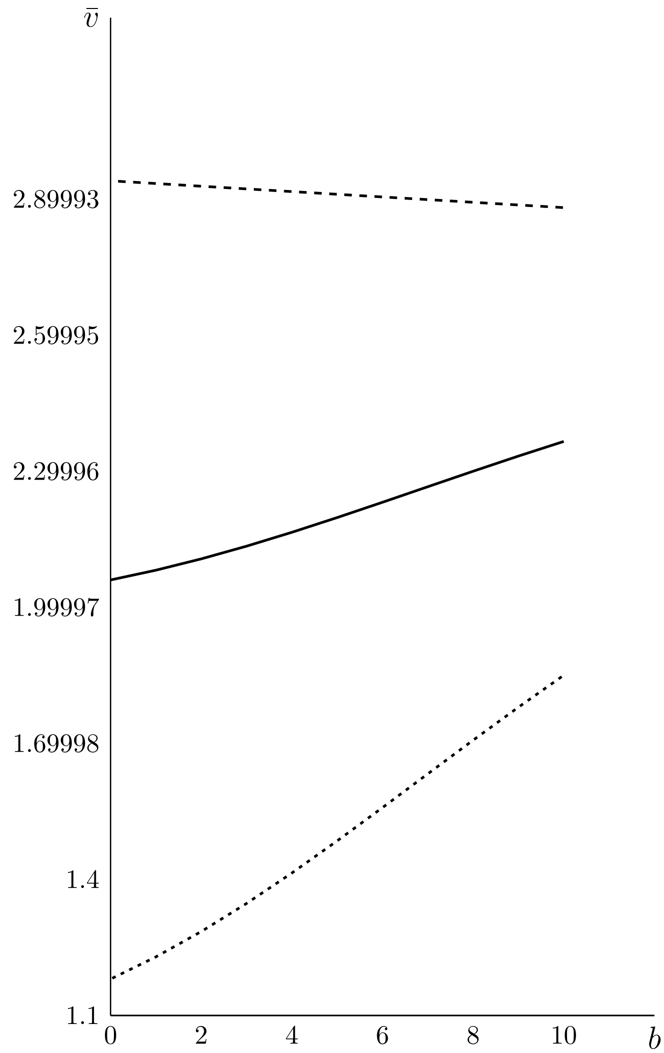

Figure 4: Average Group Payoff with b Type-B's. Dashed line is average in initial state H, dotted line is average in initial state L, and solid line is expected average.

Figure 4 shows the group average payoff when it starts the period in state H – the top line – and the payoff when it starts the period in state L – the bottom line. The expected average payoff – the middle line – is the average of these payoffs.

When the group starts in H, higher b slightly reduce the group's expected payoff, because B's increase the chance of switching to L. However, it is much less likely to switch to L than to stay in H, so the payoff reduction is small. In contrast to this small effect, when the group starts in L, increasing b significantly increase the group's expected payoffs, because B's have a larger effect on the transition L to H, which is an easier transition because the basin of attraction for H is large. Thus the presence of B's is much more beneficial for the group in L states than it is harmful in H states.

6.2 Representativeness Heuristic

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) introduced the idea to economists, while Rabin (2002) models how the bias affects behavior. Kahneman (2011) is a popular exposition.

In general, the representativeness heuristic will lead to errors being correlated across subperiods, because sampled strategies from a vector s_t containing more errors will lead to higher error rates in s_{t+1}. I therefore adopt a stylized model of the representativeness heuristic that preserves the Markov property of the reduced $\mathcal{S}$ state space and serves as a lower bound on the true effect of the heuristic. Players choose strategies in two stages: in stage one all players duplicate the R-type strategy choice method by choosing a best response to s_{t-1} with probability $1-\epsilon$ and a random strategy with probability ϵ. Then in stage two, type-B players who did not make a mistake in stage one may, each with probability η, use the heuristic to revise their strategy, using as input the stage one strategy vector of the group. If a B-type does revise their strategy, they sample two strategies from stage one at random, form beliefs about the distribution of behavior in the group using that sample, and best respond to that belief.

With a sample size of two, revising B-types' believe that either that no one is conforming, half the group is conforming, or the entire group is conforming. Let $\hat{p}$ be the fraction of the group they believe to be conforming. Then for revising B's

$$\hat{p} = \begin{cases} 1 & \text{with prob. } (1-\epsilon/2)^2 \\ 1/2 & \text{with prob. } 2(\epsilon/2)(1-\epsilon/2) \\ 0 & \text{with prob. } (\epsilon/2)^2 \end{cases} . \tag{8}$$

These beliefs feed into the strategy choice revision with

$$s_{i,t+1} = \begin{cases} BR(s_{-i,t}) & \text{if } s = H \text{ and } \hat{p} > 0 \\ WR(s_{-i,t}) & \text{if } s = H \text{ and } \hat{p} = 0 \\ BR(s_{-i,t}) & \text{if } s = L \text{ and } \hat{p} = 1 \\ WR(s_{-i,t}) & \text{if } s = L \text{ and } \hat{p} < 1 \end{cases} . \tag{9}$$

Using equations (8) and (9), the induced error rate for type-B's is

$$\begin{aligned} \hat{\epsilon}(B,H) &= \epsilon/2 + \eta(1-\epsilon)(\epsilon/2)^2 \\ &= \frac{4\epsilon + \eta(2\epsilon^2 - \epsilon^3)}{8} \\ \hat{\epsilon}(B,L) &= \epsilon/2 + \eta(1-\epsilon)(1-(1-\epsilon/2)^2) \\ &= \frac{4\epsilon + \eta(8\epsilon - 6\epsilon^2 + \epsilon^3)}{8} \end{aligned}$$

as compared to $\hat{\epsilon}(R,S) = \epsilon/2$. In H states, $\hat{\epsilon}(B,H)$ is approximately equal to $\hat{\epsilon}(R,H)$, while in L states, $\hat{\epsilon}(B,L) \approx (1+2\eta)\hat{\epsilon}(R,L)$. As the example in Section 2 showed, state-dependence would be strengthened considerably if stage one was removed and type-B players instead sampled directly from s_t: then they would sometimes sample type-B players already using the representativeness heuristic, so errors would be even less (more) frequently sampled in H (L) states. [10] [11]

Figure 5 shows the average group payoffs in states L and H and the average per-period payoff, corresponding to 3 for $\epsilon = 0.02$ and $\nu = 0.5$ (holding other parameters the same as in the baseline case). From the formulae above, $\hat{\epsilon}(B, L) = 0.0101 \approx \hat{\epsilon}(R)$ but $\hat{\epsilon}(B, H) = 0.198 \approx 2\hat{\epsilon}(R)$. Thus the per-period earnings starting in H states are approximately three, because matched occur 99% of the time regardless of b, while the per-period earnings in L increase from 1.18 at $b = 0$ to 1.9 at $b = G$. The representativeness heuristic's negative effects are negligible relative to those of iid errors.

6.3 False Consensus Effect

The false consensus effect refers to the empirical regularity that people overestimate how much other people agree with them (Ross 1977; Marks and Miller 1987). For example, liberals believe the population is more liberal than it is, while conservatives believe the opposite.

Again, assume strategy choice occurs in two stages.[12] In stage one, all players choose a best response to s_{t-1} with probability $1 - \epsilon$ and a random strategy with probability ϵ – i.e.,

[10] Removing stage one would make the reduction to a two-state Markov process impossible, because the average strategy choice would then need to be kept track of. This would make characterizing the expected behavior of the system more difficult. I avoided this issue in Section 2 by loosely talking about "steady states", but how quickly the system would reach a steady state would also be important to characterizations.

[11] Increases in $\hat{\epsilon}(B, H)$ relative to $\hat{\epsilon}(B, L)$ could be smaller if type-B players used larger sample sizes. Samples of larger than two strategies would lead to $\hat{\epsilon}(B, S)$ depending on the ratio of payoffs $\overline{\pi}/\underline{\pi}$, however, so fixing sample size at two simplifies the expressions for type-B behavior as well as the resulting dynamical systems.

[12] Again this two stage setep underestimates the full effect of the false consensus effect type B's but preserves the Markovian nature of the two-state within-period process.

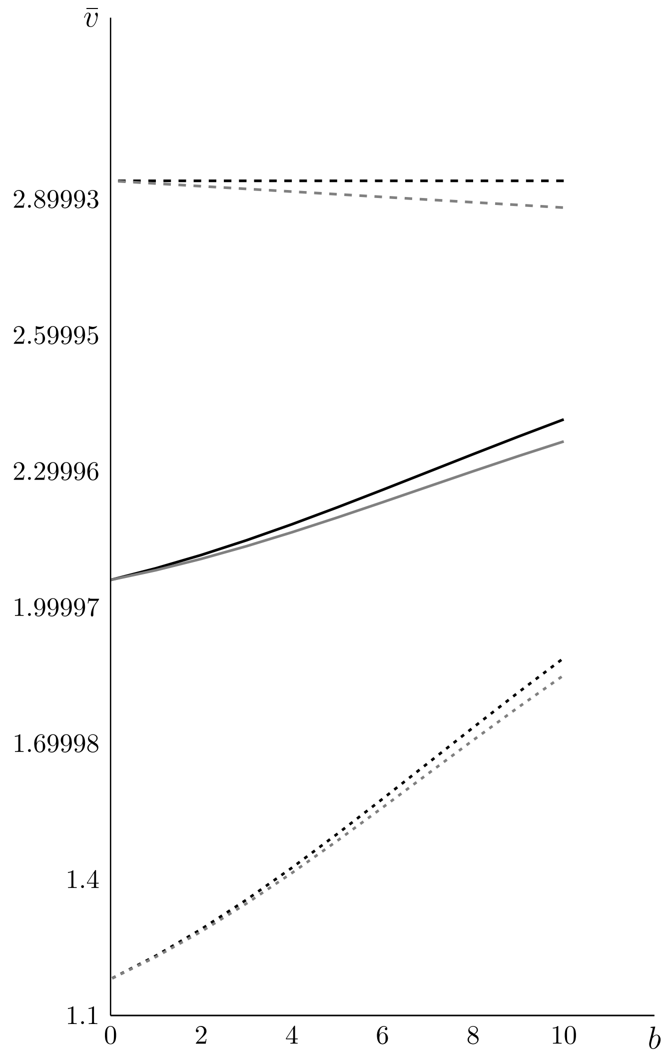

Figure 5: Average Group Payoff with b Type-B's under Representativeness Heuristic (Black lines) versus Baseline (Gray Lines). Dashed line is average in initial state H, dotted line is average in initial state L, and solid line is expected average.

they duplicate the R-type strategy choice method. Then in stage two, type-B players form new beliefs and reevaluate their stage one choice. Let p be the fraction of players (including i) choosing the conforming strategy in stage one and p_i be an indicator for whether biased player i chose a conforming strategy. Then i estimates the fraction of players who are playing a conforming strategy to be

$$\hat{p}_i = \nu p_i + (1-\nu)p \tag{10}$$

In other words, for $\eta > 0$ type-B's overweight their own choice, with overweighting increasing in η.

Player i's response to $\hat{p}_i$ is

$$s_{i,t+1} = \begin{cases} BR(s_{-i,t}) & \text{if } s = H \text{ and } \hat{p}_i > \underline{\pi}/(\bar{\pi}+\underline{\pi}) \\ WR(s_{-i,t}) & \text{if } s = H \text{ and } \hat{p}_i < \underline{\pi}/(\bar{\pi}+\underline{\pi}) \\ BR(s_{-i,t}) & \text{if } s = L \text{ and } \hat{p}_i > \bar{\pi}/(\bar{\pi}+\underline{\pi}) \\ WR(s_{-i,t}) & \text{if } s = L \text{ and } \hat{p}_i < \bar{\pi}/(\bar{\pi}+\underline{\pi}) \end{cases}. \tag{11}$$

Equation (11) shows that again the induced error rate will be state-dependent, because the basin of attraction of H is larger than the basin of attraction of L.

This state-dependence is much milder than that under the representativeness heuristic, though. For example, consider $\eta = 1/3$ and $\bar{\pi} = 3\underline{\pi}$. Then

$$\hat{p}_i = \frac{p_i + 2p}{3}$$

In state L, a type-B individual conforms after stage two if

$$\begin{aligned} U_i(L|\hat{p}_i) &> U_i(H|\hat{p}_i) \\ 1\hat{p}_i &> 3(1-\hat{p}_i) \\ \hat{p}_i &> (3/4) \end{aligned}$$

If i did make an error in stage one, he never conforms. If i did not make an error in stage one, he conforms in stage two unless $p < (3/4) - (1/3) = 5/12$, a low probability event, so with high probability he conforms.

This mild state-dependence does not mean that the false consensus effect type-B specification is similar to type R's, though. Although the state-dependence is mild, the correlation between type-B errors is highly positive when the group as a whole makes many errors in stage one. Return to the example of the last paragraph. In the rare cases where $p < 5/12$ in stage one, type B's are much more likely to be non-cornforming. But these events, while rare, are much more likely than the events necessary for type R's and/or baseline type B groups to switch equilibria. Speaking loosely, the false consensus effect induces errors when they are valuable for equilibrium switching, but it does not induce errors when they would only reduce coordination.

6.4 Loss Aversion

As the final bias I consider, I turn to loss aversion. Loss aversion is the idea that people care more about avoiding losses than achieving similar gains. It is unlike risk aversion because risk averse agents are approximately risk neutral for small gambles, but loss averse agents marginal utility is discontinuous between the loss and gain domain, causing them to avoid even small fair gambles (Rabin 2000). Loss aversion is not a cognitive bias *per se*, since it is a replacement for concave utility rather than Bayesian updating. However, the effect of loss aversion in my setting can be tractably analyzed and, in contrast the preceeding biases, it is unambiguously bad for agents. That is, it reduces biased agents' payoffs and imposes a *negative* externality on other agents. Thus it provides a nice contrast to the preceeding biases.

A type B agent with loss aversion evaluates his best response relative to a payoff-transformed version of the coordination game. Let r be this agent's reference utility. Then this agent solves

$$\max_{s_i \in A} \sum_{s' \in A} \Pr(s_{-i} = a')v(u_i(s_i, s'), r) \tag{12}$$

where

$$v(u,r) = \begin{cases} u - r & \text{if } u \geq r \\ (1+\eta)(u-r) & \text{if } u < r \end{cases}$$

Loss averse agents overweight losses by fraction η relative to type-R agents, who are expected-utility maximizing agents. I assume that loss averse agents solve equation (12) incorrectly with probability ϵ, in which case they are assigned a random strategy.

Many different assumptions about reference points are made in the literature. A reasonable choice of reference point in this context seems to be the median of the utilities earned by every agent in subperiod $t-1$, so in state H for example $r = \bar{\pi}$. This reference point and value function transforms the game's payoffs for a loss averse type B to the state-dependent payoffs given in Figure 6, where I normalize the payoffs to $U(A,B) = U(B,A) = 0$ by adding a constant payoff.

	A	B
A	$(1+\eta)\bar{\pi}$	0
B	0	$(1+\eta)\underline{\pi}$

State H

	A	B
A	$\bar{\pi} + \eta\underline{\pi}$	0
B	0	$(1+\eta)\underline{\pi}$

State L

Figure 6: Loss Averse Type-B Payoffs When $\pi_A = \bar{\pi}$. Row choice is type-B's; column choice is other player's strategy choice.

Loss aversion scales up payoffs in H by a constant, not affecting basins of attraction in H. However, it increases each payoff in state L by $\eta\underline{\pi}$. B coordination becomes relatively more attractive in L, and so the basin of attraction of the L state increases. This implies loss aversion in a group reduces the transition probabilities from L to H for that group.

The intuition for this result is that loss aversion makes failing to coordinate seem like an especially bad outcome, and so all forms of coordination become relatively more attractive. In state H, though, coordinating on B is also a loss relative to the $r = \bar{\pi}$ reference point, so the effect of the threat of non-coordination is eliminated. In state L however either A-coordination or B-coordination seems much better than non-coordination, and so the ratio $U(A, A)/U(B, B)$ falls for the loss averse. This ratio is what determines the basins of attraction of H and L.

7 Applications

In this section I argue that important modern commercial and social institutions can be understood as examples of the stochastic coordination framework in Sections 4.1 to 4.2. There are three important elements to this framework: a coordination game structure, a small-group setting, and payoffs that change over time. This section develops two examples with these characteristics: internet services, in which network externalities effectively produce a coordination game, and academic schools of thought, in which complementarity with colleagues produces a coordination game.

7.1 Internet Services

Consider an internet service such as social networking or online auctions. For many types of service, the payoffs for using website A are increasing in the share of others using that website. For example, if everyone in a social group uses Facebook rather than MySpace, the group members benefit because they can share information more efficiently. The presence of network externalities leads to a coordination game.

The primary use of these services within social networks (for social networking services) or within hobbies with thin markets (such as stamp or comic book markets on eBay) lead to these coordination games being played more often with fellow group members than with general members of the online populace.

Finally, the rapid evolution of internet services means that sometimes superior new service providers will arise, while other times new service providers inferior to existing ones will launch. Bankruptcies and dot-com bubble crashes mean that sometimes providers will disappear.

Formally, consider a population using a type of internet service, and let there be a sequence of websites $\{W_0, W_1, \ldots\}$ providing this service. For each website, let $\underline{t}(W_i)$ be the subperiod in which W_i is first available and $\bar{t}(W_i)$ be the subperiod in which W_i is last available. Assume $\underline{t}$ and $\bar{t}$ are such that a fixed number of websites compete in each subperiod. Each play of the coordination is a single exchange of information between a pair of friends.[13] If they choose different websites, they suffer disutility relative to choosing the same website, either because of a higher time cost of being on both websites, some chance of failure to communicate, or both.

[13]Note that while I describe subperiod interactins as between pairs of players throughout the paper, the same process would be generated by a G-player game. In this context, a multiplayer game formulation might make more sense.

Payoff shocks are of two kinds: the set of available websites may change, or one service may become more or less useful. Let $\pi(W,t) \in \{\bar{\pi}, \underline{\pi}\}$ be the payoff to a pair coordinating on site W in subperiod t. Normally, $\pi(W,t) = \pi(W,t-1)$, but with some small probability $\pi(W,t) = \bar{\pi} + \underline{\pi} - \pi(W,t-1)$.

$\underline{t}$, $\bar{t}$, and π together define a stochastic coordination game environment that generalizes the model in Section 4 in two ways: allowing an arbitrarily sized game and allowing $\pi_{A,t} = \pi_{B,t}$. The first generalization, increasing the size of the strategy space, would tend to make the difference between R's and B's more extreme, as it effectively makes coordination harder. The second generalization, allowing equal payoffs, would tend to make B's modestly worse off, as in subperiods where $\pi_A = \pi_B$, B's match less often but cannot improve group payoffs by speeding the transition from L to H. The qualitative predictions of the Section 4's model should apply to internet services.

7.2 Academia

As a second example, consider schools of thought or rival methodologies within an academic discipline. A philosophy department whose members all share the same general approach, such as specializing in analytic philosophy, will have more productive conversations with colleagues than if half the department were analytic philosophers and half were continental philosophers. Within an economics subfield, such as development economics, the field will be more productive if the techniques used by active researchers are similar. Researchers focusing on field experiments will have more productive conversations with each other than if half the researchers focused on field experiments and half on structural models or case studies, say. Subfields or departments naturally divide a discipline into smaller groups within which most interactions occur.

The stochastic nature of these coordination games is produced by the shifting frontier of cutting-edge research. A technique within a specialty that was very productive, such as analytic equilibrium refinements in game theory, may become relatively less so over time. In addition to waning, techniques in decline may wax again, such as persuasive models of advertising, which were worked on by top early 20th century economists such as Chamberlain and Kaldor, have been readopted and modernized by behavioral economists (see Bagwell (2007) for a historical overview).

As a stylized model of an academic stochastic coordination game, consider a philosophy department with G professors where strategy A is "focus on analytic philosophy" and strategy B is "focus on continental philosophy". A subperiod is a set of conversations between pairs in the department. Analytic philosophers get more utility from conversations with each other, and continental philosophers get more utility from conversations with other continental philosophers, than either tradition gets from cross-tradition talks. With probability $1 - \sigma$, which school of thought is more productive may switch, perhaps because of (unmodeled) events in the discipline as a whole.

8 Conclusion

I develop a theory of how cognitive biases affect the evolution of behavior in stochastic coordination games, repeated pure coordination games in which the payoff-dominant equilibrium occasionally switches. Their stochastic nature makes the speed of behavior 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, the false consensus effect, and loss aversion. Representativeness and false consensus are "good" 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. Loss aversion serves as a counterexample to the other "good" biases, because it slows movement between equilibria.

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, heretonow 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 group selective pressures, polymorphic evolutionary equilibria are likely to exist in which some individuals possess biases and others do not.

References

Bagwell, Kyle, "The Economic Analysis of Advertising," in Mark Armstrong and Robert Porter, eds., *Handbook of Industrial Organization*, Vol. 3 2007.

Bergin, James and Barton T. Lipman, "Evolution with State-Dependent Mutations," *Econometrica*, 1996, *64*, 943–956.

Boehm, Christopher, *Hierarchy in the Forest: The Evolution of Egalitarian Behavior*, Cambridge: Harvard University Press, 1999.

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, "Learning, Local Interaction, and Coordination," *Econometrica*, 1993, *61*, 1047–1071.

_ , "Basins of Attraction, Long-Run Stochastic Stability, and the Speed of Step-by-Step Evolution," *Review of Economic Studies*, 2000, *67*, 17–45.

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.

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*, 1 1993, *61* (1), 29–56.

Marks, Gary and Norman Miller, "Ten years of research on the false-consensus effect: An empirical and theoretical review," *Psychological Bulletin*, 1987.

Rabin, Mathhew, "Risk Aversion and Expected-utility Theory: A Calibration Theorem," *Econometrica*, 2000, *68* (5), 1281–1292.

— , "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.

Ross, Lee, "The false consensus effect: An egocentric bias in social perception and attribution processes," *Journal of Experimental Social Psychology*, 1977.

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.

Tversky, Amos and Daniel Kahneman, "Judgement Under Uncertainty: Heuristics and Biases," *Science*, 1974, *185*, 1124–1131.

Young, H. Peyton, "The Evolution of Conventions," *Econometrica*, 1993, *61* (1), 57–84.

— , "The Dynamics of Social Innovation," *Proceedings of the National Academy of Science*, 2011, *108*, 21285–21291.

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., $\text{tr}(A + B) = \text{tr}(A) + \text{tr}(B)$).

M_B and M_R are defined in equation (3). Let M_ϵ be the matching matrix for a player who makes errors with probability ϵ, so

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

M_ϵ can be written as

$$M_\epsilon = C[X_0 - \epsilon I] \tag{13}$$

where

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

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

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

and

$$V(0) - v(R) = \epsilon_R\,\mathrm{tr}(\Pi C N)$$

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

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

so $\partial V/\partial b$ is

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

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

Claim (i) follows because

$$\Pr(\mathrm{c}|S) \approx \frac{(1-\hat{\epsilon}(R,S))(G-b)+(1-\hat{\epsilon}(B,S))b}{G},$$

so

$$\frac{\partial \Pr(\mathrm{c}|S)}{\partial b} = -\frac{\hat{\epsilon}(B,S)-\hat{\epsilon}(R,S)}{G}. \qquad \square$$

A.3 Proof of Lemma 3

Proof. Fix $\psi = \gamma_{\bar{L}H} = \gamma_{LH}(b=G)$. Then consider

$$\frac{\psi}{(\psi+\gamma_{LH}(b))^2}$$

For $b \leq G$, $\gamma_{LH}(b) = x\gamma_{\bar{L}H}$ where $x < 1$. Therefore

$$\frac{\psi}{(\psi+\gamma_{LH}(b))^2} = \frac{\gamma_{\bar{L}H}}{(\gamma_{\bar{L}H}+x\gamma_{\bar{L}H})^2} = \frac{1}{\gamma_{\bar{L}H}(1+x)^2} \geq \frac{1}{4\gamma_{\bar{L}H}}.$$

for any b.

Next, from Lemma A2, $\partial\gamma_{LH}(b)/\partial b$ is

$$\frac{\partial\gamma_{LH}(b)}{\partial b} = \left(\frac{\hat{\epsilon}(B,L)-\hat{\epsilon}(R,L)}{G}\right)\left(\frac{(1-p)G}{\hat{\epsilon}(B,L)} - \frac{pG}{1-\hat{\epsilon}(B,L)}\right)\gamma_{LH}(b).$$

The ratio $\gamma_{LH}(b)/\gamma_{\bar{L}H}$ is

$$\left(\frac{\hat{\epsilon}(R,L)+(b/G)(\hat{\epsilon}(B,L)-\hat{\epsilon}(R,L)}{\hat{\epsilon}(B,L)}\right)^{G-pG}\left(\frac{1-\hat{\epsilon}(R,L)-(b/G)(\hat{\epsilon}(B,L)-\hat{\epsilon}(R,L)}{1-\hat{\epsilon}(B,L)}\right)^{pG} > \left(\frac{b}{G}+\frac{G-b}{G}\left(\frac{\hat{\epsilon}(R,L)}{\hat{\epsilon}(B,L)}\right)\right)^{G-pG}$$

Combining these expressions above and substituting them into equation (6) establishes the lemma. □

A.4 Proof of Theorem 2

Proof. Because

$$\begin{aligned}\bar{V}(B) &\equiv \frac{bV_B(b) + (G-b)V_R(b)}{G} \\ &= V(R,b) - \frac{b(\epsilon_B - \epsilon_R)v(R,b)}{G},\end{aligned}$$

so

$$d\bar{V}/dB = \frac{\partial V(R,b)}{\partial b} - \frac{(\epsilon_B - \epsilon_R)v(R,b)}{G}.$$

From lemmas 1, 2, and 3, these equations are positive under the theorem's conditions. □

A.5 Proof of Theorem 3

Proof. Lemma 1 establishes that when errors are state-independent, $\lambda^{ih} = 0$ and $\lambda^{il} = 1$. State-dependence in theory could reverse this relationship if B-types achieve non-conforming coordination frequently enough in L states. However, because $\hat{\epsilon}(B, L)$ is proportional to ϵ, for low ϵ this possibility will not occur. Therefore, in the limit as λ approaches 0 or 1, $\lambda^{gh} = \lambda^{gl}$, so $\Delta\lambda = -\alpha < 0$. □

A.6 Proof of Theorem 4

Proof. For there to be an evolutionary equilibrium with $\lambda^* \in (0,1)$, two conditions must hold:

i) $\alpha = \beta\nu G(\lambda^{gh}(\lambda^*) - \lambda^{gl}(\lambda^*))$, and

ii) $(\lambda^{gh}(\lambda) - \lambda^{gl}(\lambda)$ must be falling in λ.

From Theorem 2, average group payoffs are increasing in b and hence $\lambda^{gh} > \lambda^{gl}$. Thus for condition (i), for a given β, there is some α that will satisfy the condition.

For condition (ii), the composition of groups is binomially distributed (equation (7)). The variance of b is $g\lambda(1-\lambda)$, and it is evident that $\lambda^{gh} - \lambda^{gl}$ is decreasing when the variance of b is decreasing, so an interior λ can be an equilibrium in the range $(1/2, 1)$. □

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 $\mathrm{B}(G, \epsilon)$.

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,

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

and

$$\gamma_{LH} = \binom{G}{(1-p)G}\epsilon^{G-pG}(1-\epsilon)^{pG} + \text{h.o.t} \approx \binom{G}{(1-p)G}\epsilon^{G-pG}(1-\epsilon)^{pG} \tag{15}$$

and note that the binomial coefficients are equal, so

$$\frac{\gamma_{HL}}{\gamma_{LH}} = \frac{\epsilon^{pG}(1-\epsilon)^{G-pG}}{\epsilon^{G-pG}(1-\epsilon)^{pG}} = \left(\frac{\epsilon}{1-\epsilon}\right)^{(2p-1)G}. \tag{16}$$

Using equation (16), the fraction of time spent in H in the long-run can be calculated as

$$\frac{\gamma_{LH}}{\gamma_{LH}+\gamma_{HL}} = \frac{(1-\epsilon)^{(2p-1)G}}{(1-\epsilon)^{(2p-1)G}+\epsilon^{(2p-1)G}}. \tag{17}$$

For $G > 1/(2p-1)$, the exponents in equation (17) are greater than one, implying $\gamma_{LH}/(\gamma_{LH}+\gamma_{HL}) > 1-\epsilon$. It is thus clear that

Lemma A1. *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

$$\frac{\gamma_{LH}}{\gamma_{LH}+\gamma_{HL}} = \frac{\epsilon_L^{G-pG}(1-\epsilon_L)^{pG}}{\epsilon_H^{pG}(1-\epsilon_H)^{G-pG}+\epsilon_L^{G-pG}(1-\epsilon_L^{pG})}. \tag{18}$$

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 $\frac{G\hat{\epsilon}(R,S)+(b-G)\hat{\epsilon}(B,S)}{G}$)
- conforming players (with probability $\frac{1-G\hat{\epsilon}(R,S)-(b-G)\hat{\epsilon}(B,S)}{G}$)

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

Then let

$$\tilde{\epsilon_S} \equiv \hat{\epsilon}(R,S) + \frac{b}{G}(\hat{\epsilon}(B,S) - \hat{\epsilon}(R,S))$$

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

$$\gamma_{HL} \approx \binom{G}{pG}(\tilde{\epsilon_H})^{pG}(1-\tilde{\epsilon_H})^{G-pG} \tag{19}$$

and

$$\gamma_{LH} \approx \binom{G}{(1-p)G}(\tilde{\epsilon_L})^{G-pG}(1-\tilde{\epsilon_L})^{pG} \tag{20}$$

follow immediately.

The following lemma is also immediate:

Lemma A2. *The change in γ_{LH} as the group composition changes is approximately*

$$\frac{\partial \gamma_{LH}}{\partial b} = \left(\frac{\hat{\epsilon}(B,S) - \hat{\epsilon}(R,S)}{G} \right) \left(\frac{(1-p)G}{\tilde{\epsilon_L}} - \frac{pG}{1-\tilde{\epsilon_L}} \right) \gamma_{LH}(b).$$