

The Evolution of Behavior in Biased Populations

Daniel H. Wood*

April 2014

Abstract

I consider how cognitive biases affect the evolution of behavior. In my model, a population of non-Bayesians repeatedly are matched together to play 2×2 coordination games. Members of the population have systematically biased beliefs about the distribution of strategies in the population, to which they noisily best respond. Their cognitive biases lead players to make more simultaneous errors than Bayesian players would, changing the evolutionary dynamics. For a large class of biases, the long-run outcome is unchanged from the Bayesian outcome, but behavior can evolve much more quickly as a result of correlated errors.

JEL Classification: C72, C73, D03, D83.

Keywords: Stochastic stability, heuristics and biases, false consensus effect, conventions.

*Department of Economics, Clemson University, 228 Surrine Hall, Clemson, SC 29634. Email: dhwood@gmail.com. Phone: 1-864-656-4740. Fax: 1-864-656-4192. An earlier version of this paper circulated as “Cognitive Biases as Accelerators of Behavioral Evolution”. I thank Glenn Ellison, Tom Mroz, and Patrick Warren for helpful advice, and audiences at the 2013 Stony Brook Game Theory Festival, 2013 SABE/IAREP/ICABEEP conference, and Temple University for their feedback.

1 Introduction

Over the past thirty years much attention has been paid by economists to biases in probability judgement and how they affect decision-making, but less attention has been paid to how these biases affect how aggregate behavior changes over time. I analyze the evolution of behavior in populations in which errors stem from two sources: first, there is a small probability that players do not make optimal strategy choices given their beliefs about how other players are acting, and second, players are not Bayesians in forming these beliefs. Each player places too much weight upon one or more observed strategies – his focal observations – and consequently places too little weight upon the other strategy choices in the population. I show that the presence of these biases generally does not change the long-run evolutionary outcome, but that convergence to that outcome can occur much faster than in a population of Bayesians.

In my model n players repeatedly play a 2×2 coordination game with pairwise random matching, noisily best-responding to their beliefs about the other players' strategies, as in Kandori, Mailath, and Rob (1993). Players generate new beliefs each period based on observing the exact history of play that period, but treat r strategy observations differently than the remaining $n - r$ observations. They assign weight $(1 - \eta)/n$ to every strategy observation but also assign an additional weight η/r to their focal strategy observations.

For instance, some players could assign additional weight to their own strategy choices, believing the population to be behaviorally more similar to themselves than it actually is, corresponding to beliefs influenced by the false consensus effect. Alternatively or in addition, some players could overweight the same (random) focal observation each period because that observation “sticks out” for some reason. This overweighting corresponds to the availability heuristic, a heuristic in which people estimate probabilities of events by how easy it is to recall instances of the event. Some players could have idiosyncratic random focal observations,

corresponding to the representativeness heuristic, which causes people to treat small samples from a population as more representative of the population than they actually are. I describe these biases further in Section 2 and use them as examples throughout, but my results apply to all non-Bayesian belief formation processes that do not explicitly favor one strategy over the other.

Loosely speaking, the additional weights on focal observations cause strategy choices that are not the best response to the population’s aggregate behavior – “errors” – to be positively correlated either within a period or over time. Because players are playing a coordination game, when a player’s focal observations are errors, she is more likely to make one herself. If players share focal observations, then “clusters” of errors occur. If players have different focal observations, more errors in the current period make it more likely that players overweight errors, making errors more likely the next period as well. In that case, errors are serially correlated. Exactly how these correlations change the evolutionary dynamics is the focus of this paper.

My paper makes two contributions.¹ First, while there is considerable evidence that many people are biased in some ways, and by now many different models address how biases affect equilibrium predictions in various contexts, to my knowledge no research exists which addresses this question for stochastic evolutionary models. My first contribution is to present a tractable model that encompasses a variety of probability judgement biases and makes sharp predictions about the long-run behavior of biased populations. I show that as long as biases are not too severe and the focal strategy observations do not explicitly favor a particular strategy, then the risk-dominant equilibrium is uniquely stochastically stable.²

¹I defer relating my paper to the existing literature until after stating my main results.

²In 2×2 games, the risk-dominant equilibrium is to play the strategy which provides a higher payoff against an opponent who is equally likely to play either strategy (Harsanyi and Selten 1988). This stochastic stability result is well-known for models with Bayesian populations. Young (1993) and Kandori, Mailath, and

This result applies to a large category of unrelated biases and mixtures of biases, so if one doubts that people are perfect Bayesians, but is agnostic about exactly how people are biased, my results apply.

My second contribution is to show that optimization errors – the typical assumption about the source of noise in stochastic stability models – and errors in beliefs are qualitatively different. Populations with errors in beliefs converge much more quickly to the risk-dominant equilibrium as a result of their irrationality. Ellison (2000) shows that for small enough probability of optimization error ϵ , convergence to the stochastically stable state can take $O(\epsilon^{-cn})$ periods, where c is a constant derived from the payoffs of the game, which can be extremely slow.³ The speed of evolution is generally limited by how often players make sufficiently many simultaneous errors to move the system between basins of attraction of states corresponding to each pure-strategy equilibrium of the underlying game. The positive correlation between errors produced by biases in beliefs leads to lower order convergence times as η increases for fixed population sizes n . The asymptotic behavior as n grows large depends on the exact form the biases take; I show that, given moderate levels of bias, convergence times are bounded as $n \rightarrow \infty$ under the false consensus effect and the availability heuristic, while under the representativeness heuristic convergence times are bounded as n becomes large as long as optimization errors are frequent enough.

Ironically, the long-run evolutionary outcome for a Bayesian population is a better prediction of behavior for populations with moderate probability judgement biases than for Bayesian populations. The identical outcome is reached more quickly by a biased population.

Rob (1993) were the first to show that a population subjected to persistent small noise in choosing strategies settles in the long-run on the risk-dominant equilibrium, using techniques developed in Freidlin and Wentzell (1998) and initially applied to evolutionary dynamics by Foster and Young (1990). It is one of the more robust results in evolutionary game theory (see, e.g, Blume (2003)).

³If $f(x)$ is of order $g(x)$, written $O(g(x))$, as $x \rightarrow 0$, then $\exists C, \bar{x}$ such that $|f(x)| < C |g(x)|$ for all $x \in (0, \bar{x})$.

2 Biases in Probability Judgement

Before proceeding with my model, I describe three well-known probability-judgement biases which the model applies to. Of course there are other plausible – and many implausible – ways that people may fall short of the Bayesian ideal, such as confirmation bias (ignoring data not consistent with prior beliefs) or the false uniqueness effect (essentially the opposite of the false consensus effect). Most of my results would apply to these other biases as well.

Representativeness Heuristic: Kahneman and Tversky first introduced the representativeness heuristic, sometimes referred to as the “law of small numbers” (Tversky and Kahneman 1971, 1974). It is a probability estimation heuristic where biased individuals take small samples from a population as more representative of that population than the samples truly are. Rabin (2002) is a contemporary model of how the bias affects individual behavior while Kahneman (2011) is a popular exposition.

Most research on how the representativeness heuristic affects economic behavior are for settings where news arrives over time, such as investors evaluating stock performance. In that case, biased individuals under-react (relative to Bayesians) to short-run trends, because they under-estimate the probability of sequences with high variance and consequently expect mean-reversion. However, they over-react to longer trends, because they take these sequences as more informative than they are about the underlying direction of the market. This pattern is consistent with the real-world patterns of short-run momentum and long-run reversals in stock returns (Rabin and Vayanos 2010; Barberis et al. 1998). Similarly, Barber et al. (2009) show that individual investors’ trading is positively correlated across time and across investor, which they explain as (partially) produced by overextrapolation of past returns due to common employment of the representativeness heuristic by these investors. My model

predicts similar positive correlation stemming from the representativeness heuristic.

False consensus effect: The false consensus effect refers to an egocentric bias in which people overestimate how much other people’s judgements are similar to their own (Ross 1977; Marks and Miller 1987). For instance, students take their own performance on a test of ‘social sensitivity’ as more informative about overall pass rates on the test than another student’s performance (Alicke and Largo 1995).⁴

Many studies have found that the the false consensus effect plays an important role in simple games where social preferences influence behavior. Blanco, Engelmann, Koch, and Normann find that in sequential prisoner’s dilemmas, first-movers who cooperate are more likely to cooperate reciprocally as second-movers. Much of this behavior is caused by a subject’s second-mover decisions influencing their beliefs about what other subjects would choose in the same situation, in addition to the direct influence of subject altruism on both decisions (Blanco et al. 2009). Other researchers have found similar belief-driven correlations in subjects’ choices in trust games (Vanberg 2008b; Ellingsen et al. 2010) and in a sequential voluntary contributions game (Gächter et al. 2012).

Proto and SgROI (2013) directly test the false consensus effect by asking student subjects to predict the distribution of such decisions as mobile phone purchase or restaurant choice in the overall student population. Their subjects perceive themselves as closer to the middle of the distribution than they in fact are. Similarly, Engelmann and Strobel find that in predicting other subjects’ behavior, subjects overweight their own decisions relative to observed other subjects’ decisions, as long as effort is required to calculate the other subjects’ decisions from the available information (Engelmann and Strobel (2012); see also Engelmann and Strobel

⁴In many circumstances, treating one’s own choice as informative may be rational, and so many of the early papers in the psychology literature do not conclusively document that the false consensus effect is a *bias* (Dawes 1989; Vanberg 2008a). Subsequent papers corrected this problem.

(2000)).

Availability Heuristic: People following the availability heuristic judge how probable an event is based on how easy it is to recall examples from memory (Tversky and Kahneman 1974). In a representative study, Schwarz et al. asked subjects to list 12 examples of assertive behaviors on their part. These subjects rated themselves as less assertive than subjects who were only asked to list 6 examples. Constructing a list of 12 examples is difficult for most people, and subjects treated the difficulty of the task as providing information that they were not assertive. Indeed, subjects who are asked to list 12 unassertive examples judge themselves to be more assertive than subjects who are asked to list 12 assertive examples (Schwarz et al. 1991).

Research on the availability heuristic in economic decision-making is more limited. Kuran and Sunstein argue that the availability heuristic is important to understanding which risks to consumers will be regulated, and that “availability entrepreneurs” understand this and try to manipulate public discourse to take advantage of the heuristic (Kuran and Sunstein 1999).

3 Model

Players in a population $N = \{1, 2, \dots, n\}$ are repeatedly randomly matched together to play the simultaneous-move coordination game in Figure 1. (A, A) and (B, B) are Nash equilibria along with a mixed-strategy equilibrium in which A is played with probability

$$p^* = \frac{d - b}{a - c + d - b}.$$

I assume throughout that $a - c > d - b$ so that $p^* < 1/2$, and (A, A) is risk-dominant.

At the start of each period t , every player is matched with another player, with every match being equally likely, and plays the game once. At the end of the period, each player then forms a belief $p_{i,t}$ about the population frequency of period- t A choices. That player then chooses a strategy for the next period, $s_{i,t+1}$, based on that belief. Players are myopic in that they respond to current behavior, not taking into account that it might change.

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

where $BR(p)$ is the best response to an opponent playing A with probability p and $WR(p)$ is the other strategy.

All players observe the vector s_t but form $p_{i,t}$ using a process that places too much weight on r players' observed strategy choices and too little weight on the true frequency of players playing A . Let \mathcal{F} , the set of all possible focal strategy combinations, be the set of all r -combinations of N . Then $F_{it} \in \mathcal{F}$, player i 's focal observation set, records which players player i overweights in period t . All $j \in F_{it}$ are overweighted equally. Let $a_{i,t}^f = |\{j \in F_{it} : s_{j,t} = A\}|$ be the number of A -players in i 's focal set, and let a_t denote the number of players playing A in period t .

	A	B
A	a, a	b, c
B	c, b	d, d

Figure 1: Coordination Game. $a > c, d > b, a - c > d - b$.

Definition 1. A player i who forms his beliefs about the fraction of players playing A with a *focal-observation* belief-formation process has beliefs

$$\hat{p}_{i,t} = \eta \left(\frac{a_{i,t}^f}{r} \right) + (1 - \eta) \left(\frac{a_t}{n} \right)$$

where $\eta \in [0, 1]$ parameterizes the severity of the players's bias.

Different members of the population can have different focal-observation biases, but I focus on biases that increase the variance of players' beliefs without explicitly biasing players towards a particular strategy by requiring them to satisfy a “no directional bias” condition under which each focal observation is an A observation with probability p_t .

Definition 2. A focal-observation bias has no *directional bias* if the focal observations are assigned such that the probability of a focal strategy observation being strategy S is equal to the true probability that a strategy is S : $E[a_{i,t}^f | a_t] = ra_t/n$.

Example 1 (Focal assignment processes without directional bias). Under the *false consensus effect* (FCE), each player's own strategy is focal for herself, so $r = 1$ and

$$F_{i,t} = \{i\}, \tag{FCE}$$

If $p_t \equiv a_t/n = 1/3$ and player i has $s_{i,t} = A$ then $\hat{p}_{i,t} = \eta + (1 - \eta)(1/3) = (1 + 2\eta)/3$, while for $s_{i,t} = B$, $\hat{p}_{i,t} = (1 - \eta)/3$. Therefore $E[a_{i,t}^f | a_t] = (1/3)(1 + 2\eta)/3 + (2/3)(1 - \eta)/3 = 1/3 = p_t$ satisfying no directional bias.

Under the *availability heuristic* (AH) and the *representativeness heuristic* (RH), focal observation sets are random draws from \mathcal{F} . Let Z be a probability distribution assigning equal probability to any element of \mathcal{F} . Imagine one strategy choice is broadcast to the

entire population; players with the availability heuristic then overweight the easily-recalled broadcast. Then $r = 1$ and

$$F_{i,t} = k_t \quad \text{where} \quad k_t \sim Z. \quad (\text{AH})$$

Under RH each player overweights a different small sample (of size $r \geq 1$):

$$F_{i,t} = k_{i,t} \quad \text{where} \quad k_{i,t} \sim Z. \quad (\text{RH})$$

In both cases, there is no directional bias because Z assigns equal probability to all possible focal observation sets, so the probability of each focal observation being A is p_t .

The strategy updating function together with a belief formation process define a Markov process with state a_t on a state space $\mathcal{A} = \{0, \dots, n\}$. This Markov process is ergodic for $\epsilon > 0$ and hence has a unique stationary distribution placing probability $\mu(a)$ on each state a , which the system converges to as $t \rightarrow \infty$ from any initial state. An equilibrium is stochastically stable if for the state a corresponding to that equilibrium, $\mu^*(a) \equiv \lim_{\epsilon \rightarrow 0} \mu(a) > 0$. A basic result of the literature is that in 2×2 coordination games with Bayesian players (i.e., $\hat{p}_{i,t} = p_t$), if A is risk-dominant, then $\mu^*(n) = 1$ and A is the unique stochastically stable equilibrium.

In many cases μ^* can be completely determined from properties of the unperturbed ($\epsilon = 0$) Markov process. A limit state of the unperturbed process is a state for which $a_{t+1} = a_t$, and I denote limit state $a = i$ by ω_i . The basin of attraction of a limit state ω , $D(\omega)$, is the set of states from which ω is eventually reached with probability 1 under the unperturbed dynamic. Finally, let $W_A(n, \epsilon)$ be the expected wait time before the system reaches the basin of attraction of the all- A state, $D(\omega_n)$, if there are n players and the optimization error

probability is ϵ .⁵

4 False consensus and availability in 2×2 coordination games

This section focuses on two biases for which focal observation sets are singletons – the false consensus effect and the availability heuristic – both because they are of independent interest and as a means of providing intuition for my more general results. For these $r = 1$ cases, let $f(i, t)$ be i 's focal observation at time t , simplifying notation. It is useful to think of player i 's focal observation as shifting p^* , the minimum fraction of players playing A such that A is a best response: either reducing it if $s_{f(i,t),t} = A$ causes $p_{i,t}^{\hat{}} > p$, or increasing it if $s_{f(i,t),t} = B$ causes $p_{i,t}^{\hat{}} < p$. Let \underline{p}^* be the minimum p such that $BR(p_{i,t}^{\hat{}} | s_{f(i,t)} = A) = A$ and let \overline{p}^* be the minimum p such that $BR(p_{i,t}^{\hat{}} | s_{f(i,t)} = B) = A$. These cutoffs are

$$\underline{p}^* \equiv \max\left\{0, \frac{p^* - \eta}{1 - \eta}\right\} \quad \text{and} \quad \overline{p}^* \equiv \min\left\{\frac{p^*}{1 - \eta}, 1\right\}. \quad (1)$$

They form an interval around p^* within which the movement of the system is influenced by the realized focal observations as well as by optimization errors.

False consensus effect: Players with this bias place undue weight on their own strategy choice, producing a “solipsistic region” around p^* where $BR(p_{i,t}^{\hat{}} | s_{f(i,t),t} = A) = A$ and $BR(p_{i,t}^{\hat{}} | s_{f(i,t),t} = B) = B$. These states $a \in (np_{i,t}^{\hat{}}, n\overline{p}^*)$ are never exited except through

⁵The Appendix contains formal definitions and proofs. $W_A(n, \epsilon)$ is defined in terms of reaching the basin of attraction of $a = n$ because once that basin of attraction is reached, there is a high probability that the system reaches a state close to $a = n$ the next period, and is unlikely to leave the basin of attraction for many periods.




Figure 2: Best-Response Regions Under the False Consensus Effect. In shaded region, p drifts towards to $1/2$; to left of shaded region, p jumps to close to 0; and to right of shaded region, p jumps to close to 1.

optimization errors, which cause the state to drift towards $a = n/2$.⁶ Outside of the solipsistic region standard myopic best response dynamics operate. Figure 2 depicts this adjusted dynamic.

For $\eta \in [p^*, 1 - 2p^*]$, the solipstic region encompasses $D(\omega_0)$, the entire basin of attraction of the all- B state, and so the state drifts towards $D(\omega_n)$ due to optimization errors. Because $\eta < 1 - 2p^*$, $\bar{p}^* < 1/2$, so the system drifts into $D(\omega_n)$ rather than drifting all the way to $a = n/2$ and remaining there, which would occur if $p = 1/2$ were in $D(\omega_0)$. For this η range, numerical solution for $W_A(n, \epsilon)$ finds that $W_A(n, \epsilon) \approx \epsilon^{-1}$ for any n .

For $\eta \notin (p^*, 1 - 2p^*)$, the fastest transition path from B to A combines (fast) linear increases in p in the solipsistic region with (slow) waiting for enough simultaneous errors to enter it from $D(\omega_0)$ or to exit it into $D(\omega_n)$. For low biasedness $\eta < p^*$, the fastest transition path from B to A involves first a jump from $D(\omega_0)$ to $p > \underline{p}^*$, the point at which linear growth takes over, after which p grows quickly until it reaches the basin of attraction of A . For high biasedness, the fastest transition path begins with quick growth to $p = 1/2$ and then requires a jump from around $p = 1/2$ into $D(\omega_n)$. By replacing a large jump with a smaller

⁶There are approximately $\epsilon(n - a)$ switches from $s_{i,t} = B$ to $s_{i,t+1} = A$ each period and ϵa switches from $s_{i,t} = A$ to $s_{i,t+1} = B$, so if $a \lesssim n/2$, $E[a_{t+1} | a_t] \gtrsim a_t$.

jump plus linear growth, a population's bias reduces $W_A(n, \epsilon)$. However, as n increases, the smaller jump still takes more and more time in expectation to occur, so $\lim_{n \rightarrow \infty} W_A(n, \epsilon) = \infty$.

While $W_B(n, \epsilon)$ also falls as η increases, it remains of higher order than $W_A(n, \epsilon)$ as long as $\eta \leq 1 - p^*$. However, if $\eta > 1 - p^*$, then the entire state space is a solipsistic region. The stationary distribution puts mass on all $a \in \mathcal{A}$ with greatest weights on states near $a = n/2$.⁷

Availability heuristic: Players following the availability heuristic place too much weight on the same focal observation $f(t) = f(1, t) = \dots = f(n, t)$ when considering which strategy to adopt. They either overestimate p , if $s_{f(t)} = A$, or underestimate it, if $s_{f(t)} = B$. If $s_{f(t)} = A$ and $p \geq \underline{p}^*$, then for all players $\hat{p}_{i,t} \geq p^*$, so all members of the population switch to A unless they make an optimization error. The most likely transition path out of $D(\omega_0)$ is for random optimization errors to lead to $p_t \geq \underline{p}^*$ and then for the entire block of biased players to switch to A when the focal observation is A .

Instead of changing part of the transition process to a fast linear one, the availability heuristic instead makes large numbers of simultaneous errors more likely. Because it does not produce serially correlated errors, there is no intermediate region in the state space in which the state can “get stuck” around $p = 1/2$. For degree of bias $\eta > p^*$, $W_A(n, \epsilon)$ is low even for large n , although for $\eta > 1 - p^*$, the A equilibrium is no longer the long-run outcome, for similar reasons to the false consensus effect.

⁷ The stochastic process is completely symmetric with regards to strategy labels, due to the symmetry of optimization errors and $BR(\hat{p}_{i,t})$ being independent of p_t . As $\epsilon \rightarrow 0$, the stochastic process resembles a logistic birth-death process with $\mu \sim B(n, 1/2)$.

5 General Asymptotic theory

As is typical in the stochastic stability literature, I analyze the limiting behavior of the system as $\epsilon \rightarrow 0$ or $n \rightarrow \infty$. I employ Ellison (2000)'s radius-coradius technique to establish the ϵ -limit results and a deterministic approximation of the evolutionary process as a difference equation for the n -limit results.⁸ The small- ϵ results apply to any mixture of focal-observation biases with no directional bias, while the large- n results depend on the exact bias or biases the population suffers from.

Two conditions simplify the small- ϵ theory without substantially limiting the results:

$$(C1) \quad \eta < 1 - p^*.$$

$$(C2) \quad \text{Let } Y_t \subset N. \text{ Then for all } i \in N, \text{ (i) if } Y_t \neq \emptyset, \forall y \in Y_t, \Pr(y \in F_{it}) > 0, \text{ and (ii) } \\ \forall k \in F_{it} \setminus Y_t, \text{ either } k = i \text{ or } \Pr(k \in F_{it}) > 0.$$

The first condition rules out cases in which the degree of bias η is large enough that the payoff advantages of strategy A are too small to overcome the additional noise introduced by the bias, in which case the A equilibrium is no longer the unique long-run outcome.⁹ It seems likely that this assumption is satisfied for all applications of the theorems: if $\eta > 1 - p^* > 1/2$, the evidence on cognitive biases would be much more conclusive than it is. (C2) rules setting focal observation sets to generate fixed networks and seems likewise to be satisfied for most applications of interest. It essentially requires that focal observation sets be formed in a random way, either being drawn from a set of shared observations Y_t , as with the availability

⁸Sandholm (2010) devotes considerable attention to treating the large-population limits of stochastic evolutionary processes as deterministic evolutionary processes. See also Sandholm (2001) or Benaim and Weibull (2003).

⁹(C1) is a necessary condition in Theorems 1 and 2 for $r = 1$ and sufficient for $r > 1$. If the bias is too severe, then the increased noise is large enough that the ergodic distribution μ is symmetric ($\mu(a) = \mu(n-a)$), with the exact form dictated by the particular focal-agent assumption made. I discuss this issue further following my proof of Theorem 1.

heuristic, or being drawn from the population as a whole.

Finally for small n , integer problems cause the risk-dominant equilibrium to not be stable, so I assume the population size is $n > \underline{n} \equiv \max\{\frac{r(1-\eta)}{1-p^*-\eta}, \frac{1-\eta}{1-2p^*}\}$. Provided these conditions are satisfied, the long-run outcome duplicates that of a Bayesian population.

Theorem 1. *For 2×2 coordination games played by a population of size $n > \underline{n}$ with every member forming his belief through a (possibly different) focal-observation process without directional bias that satisfies (C1) and (C2), then the risk-dominant strategy is uniquely stochastically stable.*

In addition transitions between equilibria are faster in biased populations:¹⁰

Theorem 2. *For 2×2 coordination games played by a population of size $n > \underline{n}$ with every member forming his belief through a (possibly different) focal-observation process without directional bias that satisfies (C1) and (C2), then as $\epsilon \rightarrow 0$,*

- i) there is a $c_A > 0$ such that the expected wait time to reach the risk-dominant equilibrium $W_A(n, \epsilon) < c_A \epsilon^{-\lceil p^* n \rceil}$ if $\eta < p^*$ or $W_A(n, \epsilon) < c_A \epsilon^{-r}$ otherwise; and*
- ii) there is a $c_B > 0$ such that the expected wait time to leave the risk-dominant equilibrium $W_B(n, \epsilon) > c_B \epsilon^{-\lceil (1-p^*) n \rceil}$.*

Biases increase transition speeds because they cause correlation in errors. As η increases, the speed of exits from $D(\omega_0)$ are still faster than the speed of exits from $D(\omega_n)$ as long as the increased noise from the biases does not fully counteract the advantage of the A equilibrium in the best-response dynamics. Hence the system spends almost all time in $D(\omega_n)$, and for low ϵ most of the time it is spent at $a = n$. Figure 3 shows calibrated $W_A(N, \epsilon)$ as η varies for my illustrative biases.

¹⁰Recall that $\lceil x \rceil$ denotes the smallest integer greater than x . In the following theorem, because a_t is an integer, conditions like $p > p^*$ become $a > \lceil p^* n \rceil$ after accounting for the state space.




Figure 3: Speeds of Transition to Risk-Dominant Equilibrium by Bias for $r = 1$, $p^* = 1/3$, $n = 50$, and $\epsilon = 1/12$. With Bayesians ($\eta = 0$) the transition speed $W_A(50, 1/12) \approx 6$ million periods, while for biases with $\eta > 2/10$, $W_A(50, 1/12) < 150$ periods. In contrast, $W_B(50, 1/12) > 6$ million periods for any $\eta < 0.58$.

With Bayesian populations all states within a basin of attraction, share the same best response, so players' strategy choices are not correlated between players during a period or over time.¹¹

$$\text{corr}(a_{i,t}, a_{j,t} | a_{t-1} \in D(\omega_0)) = \frac{\mathbb{E}[(a_{i,t} - \epsilon)(a_{j,t} - \epsilon) | a_{t-1} \in D(\omega_0)]}{\sigma_{a_{i,t}} \sigma_{a_{j,t}}} = 0.$$

Likewise, conditional on which basin of attraction the system is in at $t - 1$, errors are uncorrelated over time. Because movement between basins of attraction requires sufficiently many simultaneous errors, uncorrelated errors cause the population to converge slowly to the long-run distribution. Focal-observation beliefs cause correlated errors in the region $a \in$

¹¹Here with a slight abuse of notation, $a_{i,t}$ is an indicator taking the value 1 if $s_{i,t} = A$ and 0 otherwise.

$(\underline{p^*n}, \overline{p^*n})$. For lower η , the evolutionary system blends these dynamics and the uncorrelated dynamic, making exits from both $D(\omega_0)$ and $D(\omega_n)$ easier.

The transition speeds in Theorem 2 are for fixed n . Even when $W_A(n, \epsilon)$ is less than $c_A \epsilon^{-r}$, the expected wait time can still increase sharply with population size if c_A increases with n , so I also consider large-population convergence times. Here the exact form of focal-observation bias affects the limiting behavior, so in what follows I focus on the conditions under which a homogenous population sharing one of the three biases described in Section 2 will have “small” wait times even for large population sizes. For these theorems, let $W_A(\epsilon) \equiv \sup \{W_A(n, \epsilon) : n \in \mathbb{N}\}$ denote an upper bound on the expected wait time to reach the A basin of attraction, if the wait time is bounded.¹² The idea behind these theorems is that for populations with serial-correlation-producing biases and large n , $p_{t+1}(p_t) \approx f(p_t) \equiv \mathbb{E}[p_{t+1} | p_t]$. If outside of $D(\omega_n)$, $f(p_t) - p_t > \delta > 0$, then $D(\omega_n)$ will be reached in finite time regardless of n .

Theorem 3. *For in 2×2 coordination games played by a population with a focal-observation belief-formation process with sample size $r = 1$, if*

- i) the process is (RH) or (FCE), $W_A(\epsilon)$ exists if and only if $p^* \leq \eta < 1 - 2p^*$, while if*
- ii) the process is (AH), $W_A(\epsilon)$ exists if and only if $p^* \leq \eta < 1 - p^*$.*

With $r = 1$, numerical calculations find that $W_A(n, \epsilon) \approx \epsilon^{-1}$ regardless of n if the above conditions hold, so transitions to A are also fast in practical terms for intermediate η .

Theorem 4. *In 2×2 coordination games played by a population with the (RH) belief-formation process with sample size $r > 1$, $W_A(\epsilon)$ exists if and only if*

- (i) $\eta \geq p^*$, and*

¹²Kreindler and Young (2013), who look at $n \rightarrow \infty$ behavior in an environment similar to that in my Theorem 4, say an evolutionary system displays “fast selection” if $W_A(\epsilon)$ exists.

(ii) for some $\underline{\epsilon}(r, \eta) \in [0, 1/2)$, $\epsilon > \underline{\epsilon}(r, \eta)$, and

(iii) if r is odd, $\eta \leq r(1 - 2p^*)$.

Like the false consensus effect, the representativeness heuristic produces serial correlation in error rates in states $a \in (np^*, n\bar{p}^*)$. For instance, with $r = 1$, in this region player i 's best response is $s_{f(i,t),t}$, so

$$\Pr(s_{i,t+1} = A \mid p_t) = (1 - \epsilon) \Pr(s_{f(i,t),t} = A \mid p_t) + \epsilon \Pr(s_{f(i,t),t} = B \mid p_t) = \epsilon + p_t(1 - 2\epsilon) \quad (2)$$

and $f(p_t) = \Pr(s_{i,t+1} = A \mid p_t) = \epsilon + p_t(1 - 2\epsilon)$. With $r > 1$, players' best responses depend on the exact composition of their focal-observation sets. There is a threshold $\underline{a}^f(p)$ such that when $p_t = p$, if $a_i^f \geq \underline{a}^f(p)$, then $BR(p_{i,t}) = A$. $\underline{a}^f(p)$ is


$$\underline{a}^f(p) = \min \left\{ a \in \{0, \dots, r\} \cup \{n\} : \frac{\eta a}{r} + (1 - \eta)p \geq p^* \right\}, \quad (3)$$

where I adopt the convention that $\underline{a}^f(p) = n$ if $BR(p_{i,t}) = B$ even when $a_{i,t}^f = r$. Then equation (2) generalizes to

$$f(p_t) = \epsilon + (1 - 2\epsilon) \Pr(a_i^f \geq \underline{a}^f(p_t)). \quad (4)$$

Typically $\underline{a}^f(p^*) = r$, but as p increases $\underline{a}^f(p)$ falls.

Theorem 4 enumerates three possible causes for slow convergence to the A equilibrium under the representativeness heuristic for large n , all of which are illustrated in Figure 4. For $r = 1$, because $p_{t+1} = \epsilon + (1 - 2\epsilon)p_t$, in the region $(\underline{p}^*, 1/2)$, $p_{t+1} > p_t$. However, if $\underline{p}^* > 0$ because $\eta < p^*$, as in Figure 4(a), the system gets stuck around $p = \epsilon$ for arbitrarily long



4(a): Dynamics with $\eta < p^*$ and $r = 1$. In this case, as n grows large, the system remains stuck at $p = \epsilon < \underline{p}^*$. If for some t , $p_t > \overline{p}^*$, the share of A players grows linearly (because $E(p_{t+1}|p_t) > p_t$) until the A basin of attraction is reached. Here $\epsilon = 0.05, \eta = 0.2$, and $p^* = 0.3$.

4(b): Dynamics with $\eta > p^*$ and varying number of focal observations. For $r = 1$ (blue line), the share of A players grows quickly to $p = 1/2$ but never reaches the A basin of attraction. For $r = 3$ (brown line), the system converges to the marked point. Focal Observations. For $r = 4$ (dashed red line) p_t grows linearly and reaches the basin of attraction of A in finite time for any n . In all cases $\epsilon = 0.125, \eta = 0.5$, and $p^* = 0.3$.

Figure 4: Dynamics under Representativeness Heuristic for Large Population Size n . In the region $(\underline{p}^*, \overline{p}^*)$ surrounding p^* , s_{t+1} is influenced by the expected composition of players' focal observation sets. The distribution of $a_{i,t}^f$ conditional on p_t consequently influences the shape of $E[p_{t+1}]$, and there is positive feedback, reducing the expected time in B , $W_A(N, \epsilon)$. If for all p in $[0, \overline{p}^*]$, $E[p_{t+1} | p_t] > p_t$, then the feedback everywhere is strong enough for $D(\omega_0)$ to exited in finite time for arbitrarily large n .

times as n becomes large.

Even when $\eta \geq p^*$ the system can become stuck, though, for two reasons. First, p converges to $1/2$ for $r = 1$, so if $\overline{p^*} > 1/2$, or equivalently $\eta > 1 - 2p^*$, the system takes arbitrarily long to reach $D(\omega_n)$ as n becomes large. See the blue curve in Figure 4(b). This phenomenon can occur with every odd $r > 1$ if at $p = 1/2$, $\underline{a}^f(1/2) = (r + 1)/2$.¹³ For $r > 1$, $\underline{a}^f(1/2) = (r + 1)/r$ if η is large and p^* close to $1/2$. Condition (iii) of Theorem 4 guarantees that at $p = 1/2$, $\Pr(a_i^f \geq \underline{a}^f(p))$. This condition is an artifact of the modeling assumptions I make; if players' focal observation set sizes varied, then this condition would be unnecessary.

The final reason that wait times can grow large for large n is if positive feedback in $(\underline{p}^*, \overline{p}^*)$ is not strong enough to produce $E[p_{t+1} | p_t] > p_t$ over the entire interval. In the $r = 3$ case in Figure 4(b), there is a small interval when p is slightly below p^* in which $E[p_{t+1} | p_t] < p_t$. As n grows large the system spends an arbitrarily long time at the first root of $E[p_{t+1} | p_t] - p_t$, the point marked in the figure, instead of reaching $D(\omega_n)$. Generally, however, for high enough ϵ the positive feedback is sufficiently strong that this second reason for unbounded $w_A(n, \epsilon)$ does not apply, which is the rationale behind condition (ii). Intuitively, an increase in ϵ causes the $E[p_{t+1} | p_t]$ curve to rotate clockwise,¹⁴

¹³ $f(1/2) = 1/2$ in this case because $\Pr(a_i^f \geq (r + 1)/2) = \Pr(a_i^f < (r + 1)/2)$:

$$\Pr(a_i^f \geq \frac{r+1}{2} | p = 1/2) = \sum_{k=\frac{r+1}{2}}^r \binom{r}{k} (1/2)^k (1/2)^{r-k} = \frac{1}{2} = \sum_{k=0}^{\frac{r-1}{2}} \binom{r}{k} (1/2)^k (1/2)^{r-k} = \Pr(a_i^f < \frac{r+1}{2} | p = 1/2).$$

¹⁴The exact ϵ cutoff is sensitive to the particular form that optimization errors take and has no closed-form solution given my trembles specification. Kreindler and Young (2013) analyze a similar evolutionary system with logit errors, under which fast convergence occurs for a larger range of optimization noise levels than my trembles specification.

6 Discussion

I first discuss how robust the results in the previous section are to alternate modeling assumptions. I then discuss what my paper adds to related theory literatures.

Additional heterogeneity in the population: The assumption that all members of the population have a focal-observations belief-formation process can be weakened in two ways. First, players could vary in the number of observations r which they overweight. The order of $W_A(n, \epsilon)$ would be reduced for $\eta > p^*$, because there the transition speed away from ω_0 is determined by speed at which the first few players switch to $s = B$, and $r = 1$ players switch to $s = B$ faster than $r > 1$ players. Second, the assumption that the entire population is biased can be weakened. My results follow as long as at least p^* fraction of the population has focal-observation biases with $\eta \geq p^*$. If that is the case, enough biased players can switch swiftly to playing B to take the system out of $D(\omega_0)$ with $W_A(N, \epsilon)$ larger but of the same order.¹⁵

Directional bias in focal-agent choice: In the base model the no directional bias assumption ensures that focal observations are representative of the current population distribution, rather than biased towards either A or B . If instead B strategies are β times more likely to be sampled than A strategies, then $E[a_{i,t}^f] = \frac{rp_t}{\beta + p_t(1-\beta)}$. This form of directional bias does not change the basins of attraction of the unperturbed system, so analysis of the $\epsilon \rightarrow 0$ behavior is unaffected. However, although the stochastically stable state does not change, convergence speeds in the large-population limit in are affected. For serial correlation-producing biases, if $\beta > 1$ the $E[p_{t+1}|p_t]$ curves in Figure 4 shift downwards, so that wait times are unbounded

¹⁵Oyama et al. (forthcoming) model of sampling best response dynamics is a deterministic model in which there are similar kinds of heterogeneity.

for more parameter combinations. Conversely, if A strategies are more likely to be sampled, then wait times will be bounded for more parameter combinations.

Relation to other papers: While theoretical models of social learning with biased agents exist (for instance, Eyster and Rabin (2010) or Guarino and Jehiel (2013)). I am unaware of any models that address how cognitive biases affect the evolution of behavior in strategic situations as opposed to decision-theoretic ones. Sethi (2000) examines the dynamic stability of $S(1)$ equilibria, an equilibrium concept based on an model of procedural rationality in which players sample the effect of playing each action exactly once and the action they perceive to be best based on the sample of payoffs (Osborne and Rubinstein 1998). However, that model is about the dynamic effects of cognitive limitations rather than psychological biases. Blume (2003) probably comes closest to investigating how psychological biases impact evolution, though never in those terms. He shows that noisy strategy revision processes that are skew-symmetric – roughly, noise processes where strategy labels do not matter – lead to risk-dominant equilibria being stochastically stable. That result implies my Theorem 1 but is silent on the convergence speed results in this paper.

Several varieties of stochastic evolutionary model address the problem of long wait times. Many look at evolution on networks or where populations have some form of local interaction, which often increases transition speeds (for instance Ellison (1993) or Montanari and Saberi (2010)). For some applications these structural assumptions are hard to justify, though, and this paper instead focuses on evolution in populations that interact globally.

Given that focus, the papers closest to this one are Binmore and Samuelson (1997) and Kreindler and Young (2013) which show in different ways that adding additional noise leads to waiting times of lower order than the standard model. Binmore and Samuelson combines optimization errors with a learning model that is an additional source of noise. The noise

from the learning process makes convergence to the stationary distribution fast as $\epsilon \rightarrow 0$. My paper analyzes an opposite source of noise, but mechanically the biases in my model are akin to learning in Binmore and Samuelson.

Kreindler and Young instead consider a logit response dynamic where ϵ is small but non-vanishing and show that convergence is fast if p^* (in their model parameters, $1/(2 + \alpha)$) is low enough and ϵ ($1/\beta$) is high enough. The mechanism generating fast transitions is essentially serial correlation similar to the representativeness heuristic in this paper.

My work complements these papers on both substantive and technical levels. Substantively, I address a much different question than either of these papers: how do well-documented behavioral biases change the evolution of behavior. Technically, both Binmore and Samuelson (1997) and Kreindler and Young (2013) are birth-death models that allow simple closed-form solutions for the steady-state distribution but limit the analysis from being extended to non- 2×2 games, which would be possible in my framework.

In addition, several papers consider evolution when players best-respond to a sample of observations of play. In Young (1993)'s adaptive play dynamics, players respond to a sample of the strategies that other players recently played. Oyama et al. (forthcoming) consider deterministic best response dynamics where a continuum of players randomly samples a random number of observations and show that when samples of no more than size k are likely enough, the population converges quickly to an iterated $\frac{1}{k}$ -dominant equilibrium. Focal-observation belief formation generalizes sampling in these papers.

Finally, Norman (2009) analyzes evolution when players incur a switching cost for changing strategies. These costs produce intermediate limit states similar to those that occur with the false consensus effect in my model. In both cases, the intermediate limit states cause faster transitions due to the faster nature of step-by-step evolution initially described by Ellison

(2000). Norman does not analyze transition times in the large population limit, while I show that considering the large population limit is important in my model environment.

7 Conclusion

I show that populations following perturbed best response dynamics whose members hold biased beliefs about the distribution of strategies they face converge to playing the stochastically stable equilibrium more quickly. Under best response dynamics the limiting factor on the speed of the convergence is that achieving enough simultaneous errors to switch equilibria is a rare event. Biases accelerate behavioral evolution because they introduce positive correlation in errors across players or over time, which increases the likelihood of these events.

Whether probability judgement heuristics are on net beneficial or harmful in decision-making is much debated – see for instance Gigerenzer et al. (2000) – but my work suggests they might also be beneficial due to their effect on aggregate behavior. In a separate paper (Wood 2013) I explore conditions under which the “speed-up” effect of heuristics lead to positive selection for non-Bayesian reasoning. Biased players can be favored in situations such as coordinating on fads or or coordinating on adoption of technological innovations. In these environments, members of small groups repeatedly play coordination games with occasionally changing payoffs. The small-group interaction allows the positive externality to be partially internalized, in a manner similar to reciprocal preferences in Herold (2012), while the changing payoffs makes faster transitions persistently useful.

In other settings, biases can have intermediate-run equilibrium selection implications in addition to their speed-up effect. Consider a 3×3 pure coordination game played by multiple types, where $s_{i,1} = A$ for all players but one type prefers coordinating on B to coordinating on C and a second type has the opposite preferences. This situation might

arise in the evolution of a political party's platform. If η is exogenous and type-specific, a long-lasting intermediate-run outcome is likely to be the one that the higher η type favors, as the system will reach the favored equilibrium of the more biased type relatively quickly, and that equilibrium is more difficult to leave than the initial equilibrium.

References

- Alicke, Mark D. and Edward Largo (1995), “The role of the self in the false consensus effect.” *Journal of Experimental Social Psychology*, 31, 28–47.
- Barber, Brad M., Terrance Odean, and Ning Zhu (2009), “Systematic noise.” *Journal of Financial Markets*, 12, 547–569.
- Barberis, Nicholas, Andrei Shleifer, and Robert Vishny (1998), “A model of investor sentiment.” *Journal of Financial Economics*, 49, 307–343.
- Benaïm, Michel and Jorgen W. Weibull (2003), “Deterministic approximation of stochastic evolution in games.” *Econometrica*, 873–903.
- Binmore, Ken and Larry Samuelson (1997), “Muddling through: Noisy equilibrium selection.” *Journal of Economic Theory*, 74, 235–265.
- Blanco, Mariana, Dirk Engelmann, Alexander K. Koch, and Hans Theo Normann (2009), “Preferences and beliefs in a sequential social dilemma: A within-subjects analysis.” IZA DP. No. 4624.
- Blume, Lawrence E. (2003), “How noise matters.” *Games and Economic Behavior*, 44, 251–271.
- Dawes, Robin M. (1989), “Statistical criteria for establishing a truly false consensus effect.” *Journal of Experimental Social Psychology*, 25, 1–17.
- Ellingsen, Tore, Magnus Johannesson, Sigve Tjøtta, and Gaute Torsvik (2010), “Testing guilt aversion.” *Games and Economic Behavior*, 68, 95–107.
- Ellison, Glenn (1993), “Learning, local interaction, and coordination.” *Econometrica*, 61, 1047–1071.
- Ellison, Glenn (2000), “Basins of attraction, long-run stochastic stability, and the speed of step-by-step evolution.” *Review of Economic Studies*, 67, 17–45.
- Engelmann, Dirk and Martin Strobel (2000), “The false consensus effect disappears if representative information and monetary incentives are given.” *Experimental Economics*, 241–260.
- Engelmann, Dirk and Martin Strobel (2012), “Deconstruction and reconstruction of an anomaly.” *Games and Economic Behavior*, 678–689.
- Eyster, Eric and Matthew Rabin (2010), “Naïve herding in rich-information settings.” *American Economic Journal: Microeconomics*, 2, 221–243.

- Foster, Dean and H. Peyton Young (1990), “Stochastic Evolutionary Game Dynamics.” *Theoretical Population Biology*, 38, 219–232.
- Freidlin, M.I. and A.D. Wentzell (1998), *Random Perturbations of Dynamical Systems*. Springer-Verlag, New York.
- Gächter, Simon, Daniele Nosenzo, Elke Renner, and Martin Sefton (2012), “Who makes a good leader? cooperativeness, optimism, and leading-by-example.” *Economic Inquiry*, 50, 953–967.
- Gigerenzer, Gerd, Peter M. Todd, and the ABC Research Group (2000), *Simple Heuristics that Make Us Smart*. Oxford University Press, Oxford.
- Guarino, Antonio and Phillippe Jehiel (2013), “Social learning with coarse inference.” *American Economic Journal: Microeconomics*, 5, 147–174.
- Harsanyi, John C. and Reinhard Selten (1988), *A General Theory of Equilibrium Selection in Games*. MIT Press, Cambridge.
- Herold, Florian (2012), “Carrot or stick? The evolution of reciprocal preferences in a haystack model.” *American Economic Review*, 102, 914–940.
- Kahneman, Daniel (2011), *Thinking, Fast and Slow*. Farrar, Straus and Giroux, New York.
- Kandori, Michihiro, George J. Mailath, and Rafael Rob (1993), “Learning, mutation, and long run equilibria in games.” *Econometrica*, 61, 29–56.
- Kreindler, Gabriel E. and H. Peyton Young (2013), “Fast convergence in evolutionary equilibrium selection.” *Games and Economic Behavior*, 80, 39–67.
- Kuran, Timur and Cass R. Sunstein (1999), “Availability cascades and risk regulation.” *Stanford Law Review*, 683–768.
- Marks, Gary and Norman Miller (1987), “Ten years of research on the false-consensus effect: An empirical and theoretical review.” *Psychological Bulletin*.
- Montanari, Andrea and Amin Saberi (2010), “The spread of innovations in social networks.” *Proceedings of the National Academy of Sciences*, 107, 20196–20201.
- Norman, Thomas W.L. (2009), “Rapid evolution under inertia.” *Games and Economic Behavior*, 66, 865–879.
- Osborne, Martin J. and Ariel Rubinstein (1998), “Games with procedurally rational players.” *American Economic Review*, 88, 834–847.

- Oyama, Daisuke, William H. Sandholm, and Olivier Tercieux (forthcoming), “Sampling best response dynamics and deterministic equilibrium selection.” *Theoretical Economics*.
- Proto, Eugenio and Daniel SgROI (2013), “Self-centered beliefs on individuals’ choices and characteristics.” Working paper.
- Rabin, Matthew (2002), “Inference by believers in the law of small numbers.” *Quarterly Journal of Economics*, 117, 775–816.
- Rabin, Matthew and Dimitri Vayanos (2010), “The gamblers and hot-hand fallacies: Theory and applications.” *Review of Economic Studies*, 77, 730–778.
- Ross, Lee (1977), “The false consensus effect: An egocentric bias in social perception and attribution processes.” *Journal of Experimental Social Psychology*.
- Sandholm, William H. (2001), “Almost global convergence to p-dominant equilibrium.” *International Journal of Game Theory*, 107–116.
- Sandholm, William H. (2010), *Population Games and Evolutionary Dynamics*. MIT Press.
- Schwarz, Norbert, Herbert Bless, Fritz Strack, et al. (1991), “Ease of retrieval as information: Another look at the availability heuristic.” 61, 195–202.
- Sethi, Rajiv (2000), “Stability of equilibria in games with procedurally rational players.” *Games and Economic Behavior*, 32, 85–104.
- Tversky, Amos and Daniel Kahneman (1971), “Belief in the law of small numbers.” *Psychological Bulletin*, 76, 105–110.
- Tversky, Amos and Daniel Kahneman (1974), “Judgement under uncertainty: Heuristics and biases.” *Science*, 185, 1124–1131.
- Vanberg, Christoph (2008a), “A short note on the rationality of the false consensus effect.” Unpublished manuscript.
- Vanberg, Christoph (2008b), “Why do people keep their promises? an experimental test of two explanations.” *Econometrica*, 76, 1467–1480.
- Wood, Daniel H. (2013), “Cognitive biases in stochastic coordination games and their evolution.” Working paper.
- Young, H. Peyton (1993), “The evolution of conventions.” *Econometrica*, 61, 57–84.

A Proofs

Several proofs use the radius-coradius technique of Ellison (2000), which bounds transition speeds by comparing how difficult it is for the system to move between different sets of states defined by the unperturbed process, the Markov process for parameter $\epsilon = 0$ (but note $\eta > 0$ is assumed under the unperturbed process). A limit state ω of a Markov process is a state such that $\Pr(a_{t+1} = \omega \mid a_t = \omega) = 1$. The basin of attraction of ω is the set of states from which ω is eventually reached with probability 1: $D(\omega) = \{a \in \mathcal{A} \mid \Pr(\exists T \text{ s.t. } \forall t > T, a_t = \omega \mid a_o = a) = 1\}$. I adopt the notation that ω_k is a limit state such that $a_{t+1} = a_t = k$.

Let $c(x, y)$ be the minimum number of errors necessary to reach state $a = y$ from state $a = x$ and ω be a limit state of the unperturbed process given by some belief-formation function defined in Section 3 and the strategy-updating function (S) where $\epsilon = 0$. Then the radius of ω , $R(\omega)$, is the minimum cost to leave $D(\omega)$:

$$R(\omega) = \min_{a \notin D(\omega)} c(\omega, a),$$

and the coradius of ω , $CR(\omega)$, the cost of reaching ω from the most costly $a \in \mathcal{A}$:

$$CR(\omega) = \max_{a \in \mathcal{A}} c(a, \omega).$$

Finally, consider a path $a = y$ to $a = x$ that passes through intermediate limit states $l(1), l(2), \dots, l(r) \in \{x, \dots, y\}$. Then the modified cost of that path is $c^*(x, y) = c(x, y) - \sum_{i=2}^{r-1} R(l(i))$ and the modified coradius is $CR^*(\omega) = \max_{a \in \mathcal{A}} c^*(a, \omega)$. These modified costs take into account that a series of small jumps is faster than a single large jump.¹⁶ Then

¹⁶An analogy is that the evolution of wings is much more likely if it can proceed via small mutations that produce increasingly wing-like structures that are each evolutionarily stable, that if it must occur through a single large evolutionary step requiring many mutations.

(Ellison (2000), Theorem 2)

Lemma A1. *Suppose that limit state ω is such that $R(\omega) > CR^*(\omega)$. Then*

- i) ω is the long-run stochastically stable state of the system, and*
- ii) $\forall a \neq \omega$, the expected wait to reach ω from a is $O(\epsilon^{-CR^*(\omega)})$ as $\epsilon \rightarrow 0$.*

Two wait times of interest are the expected wait time before the system reaches either the basin of attraction $D(\omega_0)$ or $D(\omega_n)$ from the state state $a = n$ or $a = 0$ respectively. unperturbed dynamic, which depends on parameters n and ϵ . Let $W_A(n, \epsilon)$ be the expected wait time to reach the A basin of attraction, so

$$W_A(n, \epsilon) \equiv E[\min\{t \mid a_t \in D(\omega_n), a_0 = 0\}],$$

and let $W_B(n, \epsilon)$ be defined accordingly.

A.1 Proof of Theorem 1

First, the set of limit states is ω_0 , ω_n , and possibly some or all of the intermediate states ω_a such that $a \in (np^*, n\bar{p}^*)$. Transitions for intermediate limit states under the unperturbed dynamic are entirely determined by each player's focal observation set $F_{i,t}$.

I first show that these intermediate limit states, if they exist, are relatively easy to escape.

Lemma A2. *If ω_k is an intermediate limit state, then $R(\omega_k) = 1$.*

Proof. Assume towards contradiction that $R(\omega_k) > 1$. Then under the unperturbed dynamic, from state $a = k - 1$, state $a = k$ is reached almost surely as $t \rightarrow \infty$. Clearly this cannot be the case given (C2) if $r > 1$. If $r = 1$, at all times t such that $a_t = k - 1$, $n - k$ players must

have focal sets such that $s_{f(i,t),t} = A$ and k players must have focal sets such that $s_{j(i,t),t} = A$, so $E[a_{i,t}^f] = k/n \neq a_t/n = (k-1)/n$. Hence $R(\omega_k) > 1$ contradicts the no directional bias assumption for $r = 1$. \square

A lower bound on $R(\omega_n)$ given errors in beliefs can be calculated by letting all players i have $F_{i,t}$ such that $a_{i,t}^f = 0$. This bound may not be tight depending on the belief-formation process. Given $a_{i,t}^f = 0$, $p_{i,t} = (1-\eta)p$, and so all states $a \geq np^*/(1-\eta)$ are still in $D(\omega_n)$, so $R(\omega_n) \geq \lceil n - np^*/(1-\eta) \rceil = \lceil n(1 - \bar{p}^*) \rceil$ as long as $p^* < 1 - \eta$. Now, to bound $R(\omega_0)$, consider first $\eta < p^*$ and assume towards contradiction that $\bar{a} \equiv \lceil n(p^* - \eta)/(1 - \eta) \rceil \in D(\omega_0)$. $\bar{a} \geq r$ because $n > \bar{n}$, so if at $a_t = \bar{a}$, then the probability that each $s_{it} = A$ player has $a_{it}^f = r$ is positive under (C2), in which case

$$p_{i,t}(s_{it} = A \mid a_{it}^f = r) = \eta + (1-\eta)\frac{\bar{a}}{n} \geq \eta + \left(\frac{1-\eta}{n}\right) \left(\frac{n(p^* - \eta)}{(1-\eta)}\right) \geq p^*$$

so $a_{t+1} = a_t$, implying $\bar{a} \notin D(\omega_0)$. Therefore $R(\omega_0) \leq \lceil n(p^* - \eta)/(1 - \eta) \rceil = \lceil n\underline{p}^* \rceil$, as long as $\eta < p^*$. Next, consider $\eta \geq p^*$ and assume towards contradiction that $a = \bar{a} \equiv \lceil rp^*/\eta \rceil \in D(\omega_0)$. The probability that each $s_{it} = A$ player has $a_{it}^f = \bar{a}$ is positive under (C2), so

$$p_{i,t}(s_{it} = A \mid a_{it}^f = r) > \frac{\eta}{r} \left(\frac{rp^*}{\eta}\right) + (1-\eta) \left(\frac{rp^*}{\eta}\right) \geq p^*$$

in which case $a_{t+1} = a_t$, implying $\bar{a} \notin D(\omega_0)$. For $\eta \geq p^*$, $R(\omega_0) \leq \lceil rp^*/\eta \rceil$.

Now, consider $c^*(a, n)$ for any $a \in \mathcal{A}$. The cost of intermediate limit states between $a + 1$ and n is both added and subtracted from $c^*(a, n)$, so $c^*(a, n) = R(a)$. Together with Lemma A2, this implies that $\max_a c^*(a, n) = c^*(0, n)$, so $\text{CR}^*(\omega_n) = R(\omega_0)$.

Comparing $R(\omega_n)$ and $CR^*(\omega_n)$, if $\eta < p^*$, then

$$\begin{aligned} R(\omega_n) - CR^*(\omega_n) &\geq \lceil n(1 - \bar{p}^*) \rceil - \lceil np^* \rceil \\ &\geq \frac{n(1 - p^* - \eta)}{1 - \eta} - \frac{n(p^* - \eta)}{1 - \eta} - 1 = \frac{n(1 - 2p^*)}{1 - \eta} - 1 \end{aligned}$$

so $R(\omega_n) > CR^*(\omega_n)$ if $n > \underline{n} \geq \frac{1-\eta}{1-2p^*}$. If $\eta \geq p^*$, then $CR^*(\omega_n) \leq \lceil rp^*/\eta \rceil \leq r$, while $R(\omega_n) \geq n(1 - p^* - \eta)/(1 - \eta) > r$ provided $n > \underline{n} \geq \frac{r(1-\eta)}{1-p^*-\eta}$. Because $R(\omega_n) > CR^*(\omega_n)$, ω_n is uniquely stable from Lemma A1. \square

I now show why condition (C1) is necessary for $r = 1$. In the case of $r > 1$, there are $\eta > 1 - p^*$ such that A is not uniquely stable, but $\eta > 1 - p^*$ such that A is stable, and the exact η is a non-monotonic function of r and p^* , due to integer problems.

Theorem A1. *If $r = 1$ and (C1) is not satisfied ($\eta > 1 - p^*$), then ω_n is not uniquely stochastically stable.*

Proof. If $\eta > 1 - p^*$, then

$$s_{i,t+1} = BR(p_{i,t}^\wedge) = s_{f(i,t),t}. \quad (5)$$

Because the focal strategy sample is so overweighted, it follows that the stationary distribution is symmetrical: consider the state space $\mathcal{B} = \{0, \dots, n\}$, where $b = n - a$. Then due to (5) and $[E[a(f(i,t))]] = p_t$, the Markov transition matrix M_B under \mathcal{B} is identical to the transition matrix M_A under \mathcal{A} . As the stationary distribution μ_A under the original state space $\mu_A = M_A \mu_A$ is unique and $\mu_B = M_B \mu_B = M_A \mu_B$, $\mu_A = \mu_B$ for any ϵ . Hence $\mu^*(a) = \mu^*(n - a)$ and A cannot be uniquely stochastically stable. \square

It is apparent that the distribution μ^* will depend on the exact focal-observation bias.

A.2 Proof of Theorem 2

For $W_A(n, \epsilon)$, the claim is an immediate consequence of Lemma A1's claim (ii), applying the same logic as the proof of Theorem 1 above. For $W_B(n, \epsilon)$, an analogous calculation to that claim's is valid for any limit state (Ellison (2000), Lemma 6). $CR^*(\omega_0) = R(\omega_n) \geq \lceil n(1 - \bar{p}^*) \rceil$ so there is some $c_B > 0$ such that $W_B(n, \epsilon) > c_B \epsilon^{-\lceil n(1 - \bar{p}^*) \rceil}$ if $n(1 - \eta - p^*) / (1 - \eta) > r$, which is the case given (C1). \square

A.3 Proof of Theorem 3

I first prove that for large n , a deterministic approximation holds.

Lemma A3. *If $\text{Var}(a_t)$ is $O(n)$, then $\lim_{n \rightarrow \infty} p_{t+1}(p_t) = \mathbb{E}[p_{t+1}|p_t]$.*

Proof. $\text{Var}(a_t) < cn$ so $\text{Var}(p_t) < c/n$. From Chebyshev's inequality, for any e ,

$$\Pr(|p_{t+1} - \mathbb{E}[p_{t+1}|p_t]| > e) < \frac{\text{Var}(p_t)}{e^2} = \frac{c}{ne^2}$$

For any $e > 0$, as $n \rightarrow \infty$, the probability approaches zero. \square

Now I separate the theorem's claims into several parts.

For all biases and $\eta < p^*$: in an interval around ω_0 , all \hat{p} 's are such that $BR(\hat{p}) = B$. Hence from any $a_t < \bar{a} \equiv \min\{a \mid a \neq D(\omega_0)\}$, $a_{t+1} \sim B(n, \epsilon)$. From application of Chebyshev's inequality, $\lim_{n \rightarrow \infty} \Pr(a_{t+1} > \bar{a}) = 0$. Because $W_A(n, \epsilon) \geq \mathbb{E}[\min\{t \mid a_{t+1} > \bar{a}\}] = \Pr(a_{t+1} > \bar{a})^{-1}$, $\lim_{n \rightarrow \infty} W_A(n, \epsilon) = \infty$, so for this case $W_A(\epsilon)$ does not exist.

Availability heuristic for $\eta \geq p^*$: in this high biasedness case, $\hat{p}_{i,t} > p^*$ if $s_{f(t),t} = A$, so if shared focal strategy observation had $s = A$, then $\forall i, BR(\hat{p}_{i,t}) = A$. $\Pr(s_{f(t),t} = \epsilon \mid a_t \in D(\omega_0)) \geq \epsilon$ and $W_A(\epsilon)$ is the inverse of that probability.

False consensus effect for $\eta \geq p^*$: For the false consensus effect and representativeness heuristic, I show that a deterministic difference equation approximates p_t well for large n and then work with the difference equation. Consider a_{t+1} . Let $a_{t+1}^a \sim B(a_t, 1 - \epsilon)$ be the number of players with $s_t = A$ who have $s_{t+1} = A$ as well, and $a_{t+1}^b \sim B(n - a_t, \epsilon)$ be the number of players with $s_t = B$ who have $s_{t+1} = A$. Then $a_{t+1} = a_{t+1}^a + a_{t+1}^b$ and $E[a_{t+1} | a_t] = n\epsilon + a_t(1 - 2\epsilon)$. while $\text{Var}(a_{t+1}) = \text{Var}(a_{t+1}^a) + \text{Var}(a_{t+1}^b) = n\epsilon(1 - \epsilon)$.

Now applying Lemma A3, if at $t = 0$, $p_0 = 0$, then $p_1 = \epsilon$, $p_2 = \epsilon(1 - \epsilon) + \epsilon(1 - \epsilon)$, and in general $p_t = \epsilon(1 - p_{t-1}) + (1 - \epsilon)p_{t-1}$. Alternately, $p_t = \frac{1}{2}(1 - (1 - 2\epsilon)^t)$. For any n the system reaches $D(\omega_n)$ in finite time as long as $1 - 2p^* > \eta > p^*$: taking logs of $\bar{p}^* < \frac{1}{2}(1 - (1 - 2\epsilon)^t)$, for $t > \frac{\ln(1-2\bar{p}^*)}{\ln(1-2\epsilon)}$, $p_t > p^*$. To show that $W_A(n, \epsilon)$ is unbounded for $1 - 2p^* \leq \eta$, assume that some $W_A(\epsilon)$ exists. But then for n large enough, at $t = W_A(\epsilon)$, $p = \frac{1}{2}(1 - (1 - 2\epsilon)^{W_A(\epsilon)}) < 1/2$, a contradiction.

Representativeness heuristic for $\eta \geq p^*$: consider a_{t+1} . This a binomial random variable constructed out of n Bernoulli trials with probability of success $\Pr(s_{i,t+1} = A)$

$$\Pr(s_{i,t+1} = A) = (1 - \epsilon)\Pr(s_{f(i,t),t} = A) + \epsilon\Pr(s_{f(i,t),t} = B) = \epsilon + p(1 - 2\epsilon). \quad (6)$$

i.e., $a_{t+1} \sim B(n, \epsilon + p_t(1 - 2\epsilon))$. Since $E[p_{t+1} | p_t] = \epsilon + (1 - 2\epsilon)p_t$, the argument for the false consensus effect applies to this case as well. \square

A.4 Proof of Theorem 4

Let $g(p) \equiv \Pr(a_i^f \geq \underline{a}^f(p))$. As in the proof of Theorem 3, I show that a deterministic difference equation, $p_{t+1} = f(p_t) = \epsilon + (1 - 2\epsilon)g(p_t)$, approximates the evolutionary process well for large n and then work with the difference equation. The following characteristic of

$g(p)$ is used in Theorem 4's proof:

Lemma A4. *There is a point $\bar{p} \in [0, 1/2)$ such that $g(\bar{p}) = p$ and $g(p) > p$ for $p > \bar{p}$.*

Proof. Note that $g(p) = 1 - F(\underline{a}^f(p), r, p)$ where $F(z, n, p)$ is the binomial cumulative distribution function.

The binomial CDF with parameters z and r can be written as

$$F(z, r, p) = (r - z) \binom{r}{z} \int_0^{1-p} t^{r-z-1} (1-t)^z.$$

It follows that $\partial F / \partial p < 0$ and

$$\frac{\partial^2 F}{\partial p^2} = -(z - (r - 1)p)(r - z) \binom{r}{z} (1-p)^{r-z-2} p^z$$

so $\partial^2 F / \partial p^2 < 0$ for $p < \pi \equiv z / (r - 1)$ and $\partial^2 F / \partial p^2 > 0$ for $p > \pi$. Because $1 - g(z, r, p)$ is strictly convex for $p < \pi$, strictly concave for $p > \pi$, $1 - F(z, r, 1) = 1$, and $1 - F(z, r, 0) = 0$, it follows that $1 - g(z, r, p) > p$ for $p > \pi$.

Because $g(0) = 0$ and $g(1) = 1$, and $g(\cdot)$ is monotonically increasing, there is exactly one point \bar{p} such that

$$\lim_{p \rightarrow \bar{p}^-} g(p) \leq \bar{p} \leq \lim_{p \rightarrow \bar{p}^+} g(p).$$

Then for $p > \bar{p}$,

$$g(p) = 1 - F(\underline{a}^f(p), r, p) \geq 1 - F(\underline{a}^f(\bar{p}), r, p) > p,$$

using the preceding property of the binomial CDF, so $g(p) > p$ for $p > \bar{p}$.

If $g(\underline{a}^f(1/2), r, 1/2) > 1/2$ then $\bar{p} < 1/2$. If r is even then because $p^* < 1/2$ that condition is necessarily satisfied. For r odd, if $\underline{a}^f(1/2) = \frac{r+1}{2}$, then $g(p) = 1 - g(1-p)$, and $\bar{p} = 1/2$.¹⁷

¹⁷This follows from the symmetry of the binomial probability mass function $m(z, n, p)$, for which

If on the other hand $\underline{a}^f(1/2) \leq \frac{r-1}{2}$, then it will be satisfied. That condition is

$$\eta \left(\frac{r-1}{r} \right) + (1-\eta) \left(\frac{1}{2} \right) \geq p^*$$

or equivalently $\eta \leq r(1-2p^*)$. □

Now for the main proof, consider $a_{t+1} \in [0, n\overline{p}^*]$. For $z \in \{0, \dots, r\}$, let n_t^z be the number of players with $a_{i,t}^f = z$ and a_{t+1}^z be the number of players with $a_{i,t}^f = z$ and $s_{i,t+1} = A$; then $a_{t+1}^z \sim B(n_t^z, \pi_{zt})$ where $\pi_{zt} = \Pr(s_{i,t+1} = A \mid a_i^f = z, p_t)$. Finally $a_{t+1} = a_{t+1}^0 + \dots + a_{t+1}^r$ so

$$\begin{aligned} E[a_{t+1}] &= \sum_{z=0}^r \pi_{zt} n_t^z \\ \text{Var}(a_{t+1}) &= \sum_{z=0}^r \text{Var}(a_t^z) = \sum_{z=0}^r n_t^z \pi_{zt} (1 - \pi_{zt}) < n \left(\sum_{z=0}^r \pi_{zt} (1 - \pi_{zt}) \right) < \frac{nr}{4} \end{aligned}$$

Now let $p_{t+1} = a_{t+1}/n$. $E[p_{t+1} \mid p_t]$ is

$$\begin{aligned} E[p_{t+1} \mid p_t] &= \frac{1}{n} \sum_{z=0}^r \pi_{zt} n_t^z \\ &= \frac{1}{n} \left(\epsilon \sum_{z=0}^{\underline{a}^f(p_t)-1} n_t^z + (1-\epsilon) \sum_{z=\underline{a}^f(p_t)}^r n_t^z \right) \end{aligned}$$

From the strong law of large numbers $\lim_{n \rightarrow \infty} (\sum_{z=0}^{\underline{a}^f(p_t)-1} n_t^z)/n = \Pr(a_i^f < \underline{a}^f(p_t))$, so

$$\lim_{n \rightarrow \infty} E[p_{t+1} \mid p_t] = (1-\epsilon)g(p_t) + \epsilon(1-g(p_t)) = \epsilon + (1-2\epsilon)g(p_t).$$

Let $\Delta(p) \equiv E[p_{t+1} \mid p_t = p] - p = \epsilon + (1-2\epsilon)g(p) - p$ and $\delta = \min\{\Delta(p) : p \leq \overline{p}^*\}$. Then

$m(z, n, p) = m(r-z, n, 1-p)$, so $m(z, n, 1/2) = m(r-z, n, 1/2)$, combined with the additional symmetry of $\underline{a}^f(1/2) = r - \underline{a}^f(1/2)$. For lower η , the second symmetry does not occur, and for even r , the first symmetry does not occur.

from Lemma A4 there is some ϵ such that $\delta > 0$: consider $\epsilon = \bar{p} + \delta$ where $0 < \delta < 1/2 - \bar{p}$.

Then for $p \leq \bar{p}$, because $\bar{p} < 1/2$,

$$\Delta(p) = \bar{p} + \delta + (1 - 2\bar{p} - 2\delta)g(p) - p > \bar{p} + \delta - p > \delta.$$

For $\delta > 0$, if $p_0 = 0$, $p_1 \geq \delta$, and in general $p_t \geq t\delta$. Then for t such that $p_t > \bar{p}^*$, the system is in $D(\omega_n)$ almost surely as $n \rightarrow \infty$. Hence $W_A(n, \epsilon) \leq W_A(\epsilon) \equiv \bar{p}^*/\delta + 1$. \square