Penn Institute for Economic Research
Department of Economics
University of Pennsylvania
3718 Locust Walk
Philadelphia, PA 19104-6297
pier@econ.upenn.edu
http://economics.sas.upenn.edu/pier

PIER Working Paper 17-009

Fads and imperfect information

by

Nicholas Janetos

http://ssrn.com/abstract=2968483

Fads and imperfect information

Nicholas Janetos*

May 2017

Abstract

A fad is something that is popular for a time, then unpopular. For example, in the 1960s tailfins on cars were popular, in the 1970s they were not. I study a model in which fads are driven through the channel of imperfect information. Some players have better information about past actions of other players, and all players have preferences for choosing the same actions as well-informed players. In equilibrium, better informed (high-type) players initially pool on a single action choice. Over time, the low-type players learn which action the high-type players are pooling on, and start to mimic them. Once a tipping point is reached, the high-type players switch to a different action, and the process repeats. I explicitly compute equilibria for a specific parameterization of the model. Low-type players display instrumental preferences for conformity, choosing actions which appear more popular, while high-type players sometimes coordinate on actions which appear unpopular. Improving the quality of information to low-type players does not improve their payoffs, but increases the rate at which high-type players switch between actions.

Keywords: Fads, social norms

JEL Classification Numbers: D83

*538 McNeil building, 3718 Locust Walk, Philadelphia, PA 19104. Email njanetos@econ.upenn.edu.
Thanks to Aislinn Bohren, Selman Erol, Dan Hauser, Yunan Li, George Mailath, Andrew Postlewaite, Alvaro Sandroni, Kris Shaw, Jan Tilly, and especially Steven Matthews, for helpful feedback and suggestions.

1 Introduction

An important feature of consumer choice is that it is observed to shift over time, and in a way so that choices are correlated among individuals. For example, in the United States in the 1990s, consumers tended to choose loose-fitted clothing over tight-fitted clothing. In the 2000s, consumers tended to choose tight-fitted clothing instead. This back and forth in Western culture can be traced back centuries, to the tight and baggy breeches of Rennaissance Europe.

As a positive model of this phenomenom, the standard economic model of consumer choice is unsatisfactory because it requires us to accept time-varying and correlated preferences. Such a model is both intractable, and can explain too much behavior. In this paper, I take a different approach. I analyze a dynamic game between two types of short-lived players, high and low, each choosing between two equally costly actions, and who differ from each other in two ways. First, all of the players have preferences for matching the action of the high type players, and not matching the actions of low type players. I take these preferences as a primitive, but they might be generated by, for example, introducing a stage of the game in which players match with other players who chose the same action and wish to match with high type players. Second, the players differ in the quality of information they possess about the actions of players in the past. I mainly consider the case in which high type players have better information than low type players about the actions of past players, and consider other cases in extensions. Players have no special ability to distinguish which types of players choose which actions, rather, they only observe some information informative about the relative fractions of players who choose each action in the past. More precisely, once a player makes an action choice, their choice (but not their type) is visible for some (stochastic) period of time to other players. Technically, this corresponds to assuming that players observe a time-weighted average of past actions. High type players see an average of past actions which places more weight on more recent actions.

To fix an example, consider the high type players as being those who live in or near major population centers, and low type players as being those who live in or near rural or surburban regions, and consider the action choice as being between loose or baggy clothing. The interpretation, then, is that players prefer to dress like those who live in the city, and that those in the city have more up-to-date information about the sorts of clothing other players are wearing. It is not difficult to invent other examples. High type players may be those who have many friends on Facebook, low type players are those who do not have many friends, and the choice set is possible news articles to share. Or, high types may be white, low types may be Asian, and the choice set is possible extra-curricular activities to engage in before applying to a college which wishes to covertly discriminate against Asians.

I focus on stationary equilibria of this game, and I show they may be classified into two groups. In the first, players mix independently and identically between actions. No non-trivial dynamics arise from this group of equilibria. Players may all coordinate on one action, or another, or mix equally between both. In the second, players periodically switch between action choices. These dynamics are driven by the following intuition. Initially, the high type players coordinate on an action. Over time, the low type players learn which action the high type players are coordinating on, and start to choose that action. At some endogenously determined 'tipping point', the high type players switch to coordinate on a different action, while the low type players continue to coordinate on the 'old' action. The high type players are the first to switch because they are the first to percieve that the trend has 'played itself out'. The cycle then repeats, with high type players periodically switching between actions, and low type players following.

Equilibria of the model have features, broadly consistent with stylized facts, that might otherwise seem counter-intuitive. Less-informed players mimic the actions they see others taking, even though they may be mimicking the actions of other less-informed players. In this sense, less-informed players display instrumental preferences for conforming to the majority action, even though they do not have direct preferences for conformity. Well-informed players sometimes mimic the actions they see others taking, and sometimes do not, so that well-informed players sometimes appear to have preferences for conformity, and sometimes appear to have preferences for anti-conformity. These distinctions are best illustrated by an example: When I go to buy clothing, I tend to buy the sorts of clothing that I see others around me wearing. An economist observing my choices might reasonably conclude that I have a preference for conforming to what others are doing, i.e., I prefer to look like those around me. However, if I go to buy clothing with a fashionable friend, who recommends I buy a pair of pants that I haven't seen anyone else wearing, I strictly prefer his recommendation over any choice I would have made. An economist observing my choices in this case might reasonably conclude the opposite, that I have preferences for anti-conformity, i.e., I prefer *not* to look like those around me. In both cases, however, my choices are driven not by any intrinsic preferences for conformity or anti-conformity, but rather the information available to me about what certain classes of players are wearing. In this sense, the game considered in this paper is related to Bernheim (1994), in providing a rationale for the dynamics of conformist and anti-conformist behavior that is not directly grounded in direct preferences for conformity. In general, we draw negative inferences about people who are perceived to be conformist; the model suggests a rationale for why, since only less-informed players consistently act as if conformist.

I show also that the rate at which players switch between actions is driven by the speed of learning of less-informed players. If less-informed players only learn which action the well-informed players are coordinating on slowly, then the well-informed players switch less rapidly. The additional information available to well-informed players, however, has no effect on the rate at which players switch between actions. It serves purely as a coordination device. A broad stylized fact of the past few centuries is that the rate at which fads begin and end has increased. The shift from Renaissance-era baggy breeches to tight breeches took centuries, the shift from loose fitting pants in the 2000s to tight fitting pants took a decade. The model suggests that this phenomenon is driven by two factors. The first is the increasing visibility of well-informed elites. In the 1600s, this might have been the nobility, in 1980, it was popular musicians and actors. The second is the increasing democratization of information. In the 1800s, it may have taken weeks for information to spread among elites. In 1950, it still took weeks for information to spread among elites. But due to the invention of widely available sources of information such as broadcast television or the radio, what took decades to spread among non-elites in the 1800s might have only taken a few years in 1950. Accordingly, the model predicts that the length of fads in the 1800s should have been on the order of decades, while the length of fads in 1950 should have been on the order of years.

This democratization of information might have been expected to improve the welfare of less-informed non-elites. I show that contrary to this intuition, it has no effect on welfare. On one hand, the less-informed learn faster which action the well-informed are coordinating on, but on other hand, the well-informed switch more rapidly between actions. Together, these two effects exactly cancel out.

The model suggests that policies intended to help less-informed players have no effect on welfare if broadly targeted. For example, some people know that it is customary to wear a suit to a white-collar job interview. We do not wear suits to job interviews because we want to signal we can afford a suit, we wear suits to job interviews to signal that we are the sort of person who understands that the sort of thing one does to get a job at a job interview is wear a suit, and that therefore, we are the sort of person who also understands the other sorts of things that one does in an office environment to be successful. Telling one person, who otherwise would have worn a t-shirt, to wear a suit to a job interview might improve her payoff, telling everyone who might otherwise have worn a t-shirt to a job interview to wear a suit will have no effect, since the value of wearing a suit to a job interview came only because it credibly signaled that some people are well-informed.

1.1 Previous literature

Previous literature on the switching dynamics considered in this paper (Karni and Schmeidler, 1990; Matsuyama, 1992; Frijters, 1998; Caulkins et al., 2007) focus on models in which cyclical behavior is driven by differences in preferences, or technology ('conformists vs. anti-conformists' or 'predator / prey' models).

The paper conceptually closest to this one is Corneo and Jeanne (1999). There, as in this paper, one sort of player has access to better information than the other, such as the right restaurant to eat at. Gradually, the βs learn which restaurant is cool, and the authors analyze the dynamics of this learning process. However, they stop short of considering how players might switch to other restaurants once everyone has learned where to eat, and how this might impact the dynamics of the game; in the limit, players end up all pooling on one action. The major difference to this paper is that in Corneo and Jeanne (1999), players private information is about some exogenous state variable, here, the relevant private information of players is about the actions of other players, which I show gives rise to equilibrium switching dynamics.

The question of why we observe fashion and fashion trends is an old one in economics (Foley, 1893). Previous literature on fashion cycles focuses on Veblen goods and conspicuous consumption. In Pesendorfer (1995), a monopolist periodically releases new, expensive clothing lines, giving the wealthy an opportunity to buy expensive clothing to signal their wealth, then gradually lowers the price of the clothing to sell to more people, before eventually releasing a new line of expensive clothing and beginning the process again. Here, the monopolist is a 'norm entrepreneur' (Sunstein, 1995), strategically manufacturing social as-sets for profit. It is true that there are examples of monopolist fashion brands at the high end of the market—but fashion, and fashion cycles, are a much broader phenomenon, and not limited to high-end clothing. For example, no norm entrepreneur decided that car tailfins, which serve no aerodynamic purpose and are no more expensive than conventional styling, should be popular in the 1950s, and Plutarch (187) describes Cato the Younger wearing a subdued shade of purple, in reaction to what was then a trend among the Romans of wearing a bright shade of red. Nobody profits from the recent trend towards using 'Emma' as a girl's name, and many fashion trends involve clothing which is deliberately inexpensive.

Other related papers include the literature on social learning (Bikhchandani et al., 1992), and more specifically the literature on social learning with bounded memory, Kocer (2010). There, players do not observe (or remember, if players are interpreted as being long-lived) the full history of actions; here, players see only a summary statistic of past actions, and furthermore, there is no underlying state of the world which players draw inferences about.

The idea behind this paper is the same one behind the literature on supporting correlated equilibria in static games by modeling them as the result of a dynamic game in which players condition in some way on the actions of players in the past (Aumann, 1987; Milgrom and Roberts, 1991; Foster and Vohra, 1997). I apply the same concept, but require player's learning process to be Bayesian. For any dynamic (Nash) equilibrium in my model, in any period, the resulting distribution over action profiles is a correlated equilibria of the static game, however, the set of correlated static equilibria which can be supported by dynamic equilibria in this way is much smaller than the set of all correlated equilibria in the static stage game.

Although the empirical study of fads began much earlier, there is a recent interest in applying modern econometric techniques to identifying fads (Yoganarasimhan, 2012a,b). The ability to identify the 'next big thing' is of obvious interest to firms which sell consumer products. An industry of 'coolhunters' revolves around identifying what will be in fashion and what will be out of fashion.

In Section 2, I illustrate the main idea with a simple example model. In Section 3, I describe the stage game, and prove some basic results about equilibria in the static environment. In Section 4 I describe the full dynamic game, and I use analogous results to those in the static environment to characterize equilibria in the dynamic game. I show that a feature of all equilibria in the dynamic game is that β players show an instrumental preference for conformity. In Section 5, I apply the results from Section 4 to a parameterization of the game to explicitly compute equilibria and derive comparative statics results on the period length of the game and the strategies played by α players. In Section 6 I consider a generalization of the model in which α players are allowed to have access to older information about the actions of other players, and players may prefer to mimic β players. In Section 7 I conclude.

2 Illustrative example

I begin with a simple discrete time example to illustrate the two major mechanisms driving equilibrium dynamics in this paper: First, the preferences of players to match high type, but not low type actions, and second, the better information of high type players. The game analyzed in this section is simple, but not easily extendable, and so in the main body of the paper, I consider a richer continuous time game.

Model

Consider a discrete time ($t = 0, 1, 2, \ldots$) game. In each period, a continuum of players enters, each makes a once-and-for-all binary action choice $a \in \{0, 1\}$, and then each exits.[1] With equal probability, each player is either a high or a low type, denoted $\theta \in \{\alpha, \beta\}$. Once each player has chosen an action, a single player is sampled uniformly from the set of players, and

[1] That there be a continuum of players is not here necessary, but instead convenient for formulating payoffs, which depend on the action choices of other players in the same period.

his action choice, denoted a_t, is made visible to future players. Before a player chooses an action, he sees a truncated history of past sampled actions. Specifically, a low type player in period t sees the past action, a_{t-1}. A high type player, on the other hand, sees the past $N > 1$ actions, $(a_{t-1}, a_{t-2}, \ldots, a_{t-N})$. (The initial players see no actions, and the second through $N-1$st players, if the high type, see the entire history.) Hence, the information sets of a player of type θ in period t, H_θ^t, are

$$H_\alpha^t = \{(a_0, a_1, \ldots, a_{t-1})\} \, \forall t < N$$

$$H_\alpha^t = \{(a_{t-N}, \ldots, a_{t-1}\} \, \forall t \geq N$$

$$H_\beta^0 = \{\}$$

$$H_\beta^t = \{a_{t-1}\} \, \forall t \geq 1.$$

elements of which are denoted h_θ^t.

Player strategies are $\sigma_t^\theta : H_\theta^t \to [0, 1]$, denoting the probability that player t, of type θ, chooses action $a = 1$ in period t. Player beliefs are probability distributions over both the type of the player whose action was selected to be visible, and the action chosen by that player, denoted $\mu \in \Delta(\Theta^\infty \times \{0, 1\}^\infty)$.

Once each player has chosen an action, players randomly and uniformly match with the other players who chose the same action. Players who choose an action which no other player has chosen receive a payoff of 0. Players who match with a high type player receive a payoff of 1. Players who match with a low type player receive a payoff of 0. Players therefore prefer to choose actions which are more likely to be chosen by high type players, and less likely to be chosen by low type players. Formally, a player's payoff is 1 if and only if he is matched with a high type player, and 0 otherwise. Note that a high type player who chooses an action expects, at some interim stage with $t \geq N$, a payoff of

$$U_t(a, \mu, h^t_\alpha, h^{)}_\beta := \mathrm{P}_\mu(\theta_t = \alpha \mid a_t, a_{t-1}, a_{t-2}, \ldots, a_{t-N}),$$

that is, a player's interim expected payoff is the probability that a player randomly selected from the group of players who choose the same action is a high type player. Low type players have the same preferences as high type players, but do not see h^t_α, and so must form expectations over the type of previous players. A low type player's interim expected payoff, after choosing a, is therefore $\mathrm{E}_\mu[U_t(a, \mu, \tilde{h}^t_\alpha, \tilde{h}^t_\beta) \mid h^t_\beta]$.

The definition of equilibrium is standard:

Definition $\langle \sigma^\theta_t, \mu \rangle$ is an *equilibrium* iff

1. Each player is best responding, so that $\sigma^\theta_t(h^t_\theta)$ has support contained within

$$\arg\max_{a \in \{0,1\}} \mathrm{E}_\mu[U_t(a, \mu, \tilde{h}^t_\alpha, \tilde{h}^t_\beta) \mid h^t_\theta],$$

2. and beliefs, μ, are consistent with σ^θ_t wherever possible.

Analysis

This simple model captures the two important components of the richer game. First, low type players have worse information than high type players, in the strong sense that their information is a coarsening of the high type player's information: High type players see the actions of the past N players, low type players see only the action of the past player. Second, players have preferences for matching the action which would be taken by high type players, and not matching the action which would be taken by low type players.

Consider the following two candidate strategy profiles and beliefs. The first, which we might call a 'pooling' strategy profile, is one in which all players pool on the last action taken. The second, which we might call a 'periodic' strategy profile, is one in which players pool on the last action taken—with the occasional exception of high type players, who will choose a different action than the last action taken, as long as they percieve that the last N players all pooled on the same action. There are many more equilibria in the discrete time game which I do not consider here because they are not intuitively compelling. Both of 'pooling' and 'periodic' strategy profiles have analogues in the continuous time version of the game. In fact, it will happen that all equilibria in the continuous time game can be classified into these two categories, unlike the discrete time game.

Example 1 ('Pooling' strategy profile)**.** *Under this strategy profile, players coordinate on action* 1*. Formally,*

$$\sigma_t^\theta(h_\theta^t) = 1\,\forall \theta, t, h_\theta^t,$$

and μ *satisfies*

$$P_\mu(a_0, \theta_0, a_1, \theta_1, \ldots, a_t, \theta_t) = \begin{cases} 0 & \exists i \mid a_i = 0 \\ \left(\frac{1}{2}\right)^t & \textit{otherwise.} \end{cases}$$

This pooling strategy profile is one in which players coordinate on action $a = 1$, and deviators, unmmatched to any other player, believe that they will receive the lowest possible payoff, 0. Perhaps unsurprisingly, it is also an equilibrium strategy profile.

Proposition 1. *The pooling strategy profile described in Example 1 is an equilibrium strategy profile.*

Proof. Since in equilibrium all players choose $a = 1$, the probability that a player is an α

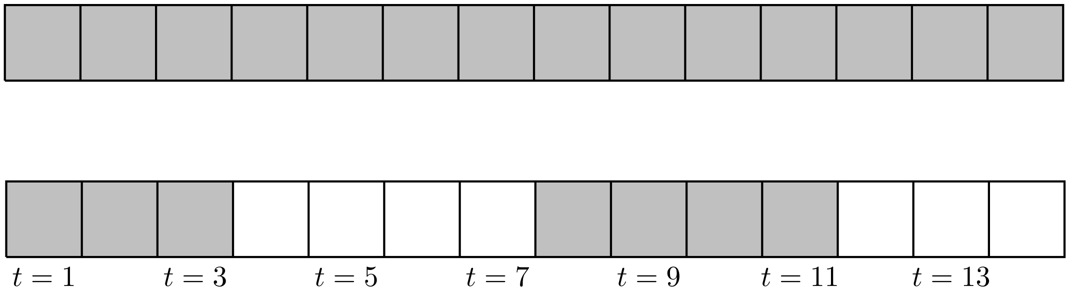

Figure 1: Two equilibrium path realizations of the pooling and periodic equilibria in the example.

Note: On the top, all the players coordinate on one action (represented by black). On the bottom, with $N=3$, the players coordinate on an action until $N \geq 3$, and a high type player comes along, at which point they switch. (Here, at $t=3, t=7$, and $t=10$.)

type, conditional on $a=1$, is always 1/2, and so for all on-path histories h^t_α, h^t_β,

$$U_t(1, \mu, h^t_\alpha, h^t_\beta) = \tfrac{1}{2}$$
$$U_t(0, \mu, h^t_\alpha, h^t_\beta) = 0,$$

and so on-path, $a=1$ is trivially a best-reply. □

Figure 1 contains an illustration of the pooling strategy profile.

A more interesting strategy profile than the pooling equilibrium is the following:

Example 2 ('Periodic' strategy profile)**.** *Under this strategy profile, low type players mimic the past action chosen, while high type players mimic the past action, unless all N observed actions are the same, in which case, they choose a different action. Formally,*

$$\sigma^\beta_t(h^t_\beta) = a_{t-1}$$
$$\sigma^\alpha_t(h^t_\alpha) = \begin{cases} 1 - a_{t-1} & a_{t-1} = \cdots = a_{t-N} \\ a_{t-1} & \textit{otherwise.} \end{cases}$$

It remains to specify the actions of the first N players. The first player uniformly chooses an *action, independently of type. The next $N-1$ players choose actions in the following way:*

1. *With probability*

$$2\left(\frac{N-1}{N+1}\right)^{\frac{1}{N}}-1$$

   the tth player, $1 \leq t < N-1$, mimics the previous action, $a_t = a_{t-1}$.

2. *Otherwise, the tth player mixes uniformly between actions.*

(The purpose of this somewhat artificial specification of the first N player's actions is to generate stationary beliefs. It is not strictly necessary, but assumed for tractability and simplicity.)

Proposition 2. *The periodic strategy profile described in Example 2 is an equilibrium strategy profile when $N \geq 5$.*

Proof. First, since players $t = 0, 1, \ldots, N-1$ choose actions with the same type-independent frequency, the expected payoffs from choosing either action are the same, and so any strategy is a best reply. Second, I state without proof that the players at $t = N$ believe

$$\mathrm{P}_\mu(a_{N-1} = a_{N-2} = \cdots = a_i) = \begin{cases} \frac{2}{N+1} & i = 0 \\ \frac{1}{N+1} & \text{otherwise.} \end{cases} \tag{1}$$

(In fact, the actions of the first N players were constructed specifically so that (1) holds! Verifying this fact is a straightforward matter of algebra and is omitted.)

Consider first whether the high type players are best responding at $t = N$. In the event that the past N players are observed to be choosing the same action, each high type player knows that each other high type player will be choosing $1 - a_{N-1}$, and so optimally each high type player prefers to choose $1 - a_{N-1}$, receiving an expected payoff of 1 (since only high type players choose $1 - a_{N-1}$ at this history) instead of 0 (since only low type players choose a_{N-1} at this history). In the event that the past N players are not observed to be choosing the same action, each high type player expects that no other player will be choosing a different action, and so optimally each high type player prefers to continue to mimic action a_{N-1}, knowing that if they choose $1 - a_{N-1}$, they will receive a payoff of zero, instead of $\frac{1}{2}$ (the probability that a player choosing action a_{N-1} is a high type player.) High type players, then, are trivially best responding at $t = N$, because choosing the same action as high type players is a best response.

Now consider whether the low type players are best responding at $t = N$ by mimicking a_{t-1}. From (1), with probability $\frac{2}{N+1}$, the past N actions are the same, and by mimicking a_{t-1}, a low type player receives a payoff of zero (since all the high type players are choosing $1 - a_{N-1}$), instead of 1. With probability $1 - \frac{2}{N+1}$, at least one of the past N actions differs, and by mimicking a_{t-1}, a low type player receives a payoff of $\frac{1}{2}$ (since all players are mimicking the past action) instead of 0. The low type player's best reply condition is satisfied, then, if and only if

$$\frac{2}{N+1} \times 0 + \left(1 - \frac{2}{N+1}\right) \times \frac{1}{2} \geq \frac{2}{N+1} \times 1 + \left(1 - \frac{2}{N+1}\right) \times 0.$$

This condition is satisfied if and only if $N \geq 5$.

Similar reasoning establishes that players are best replying in successive periods $t = N + 1, N + 2, \ldots$, if beliefs are stationary. In fact, they are. To see this, fix t, and let P_n

denote the probability that $a_{t-n} = a_{t-n+1} = a_{t-1}$. Then by construction,

$$
\begin{aligned}
P_N &= \overbrace{P_{N-1}}^{\text{Probability that last } N-1 \text{ were the same}} + \overbrace{\frac{1}{2}P_N}^{\text{Probability that last } N \text{ were the same but low type was drawn}} \\
P_n &= P_{n-1} \, \forall 1 < n < N \\
P_1 &= \frac{1}{2}P_N \\
\sum_n P_n &= 1.
\end{aligned}
$$

The unique solution for these beliefs is

$$
\begin{aligned}
P_N &= \frac{2}{N+1} \\
P_n &= \frac{1}{N+1} \, \forall n < N,
\end{aligned}
$$

so that beliefs are stationary, and so the result follows for $N \geq 5$. $\square$

This simple model illustrates the main channel driving periodic dynamics in the main model of the paper. Low type player mimic past actions, absent a better idea of whether other players are conforming to the previous action or not. High type players occasionally conform, and occasionally anti-conform when they percieve that sufficiently many other players are choosing the same action. In the remainder of the paper, I analyze a richer model, capable of providing comparative statics results. The main channel, however, is similar.

3 Static stage game

In the previous section, we considered a simple discrete time game, intended to illustrate the intuition driving the results of this paper. The presentation of the continuous time game is split into two parts. In the first, this section, I describe the static stage game, and prove some illustrative results about equilibria in this environment. Here, the information structure is taken to be exogenous. In the second, Section 4, I allow the private information of players to vary over time, and investigate how this leads to cyclical behavior.

A continuum of players, some called αs and some called βs, play the following simultaneous-move stage game. First, nature draws a state, $\vec{\rho} \in [0,1]^2$, from a distribution $\mu \in \Delta([0,1]^2)$. Denote the realized state by $(\rho_\alpha, \rho_\beta)$, and the random variable by $(\tilde{\rho}_\alpha, \tilde{\rho}_\beta)$. Players then observe private information: The α players observe $(\rho_\alpha, \rho_\beta)$, while the β players observe only ρ_β. Each player simultaneously chooses an action $a \in \{0, 1\}$. I focus on symmetric mixed strategies, denoted $A_\alpha(\rho_\alpha, \rho_\beta) \in [0,1]$ and $A_\beta(\rho_\beta) \in [0,1]$, mapping private information into the probability of choosing $a = 1$.[2]

Once players have chosen actions, each player receives a payoff, which is a function of his action, and the fraction of players of each type choosing his action. Formally, each player's payoff function is $U(a, a_\alpha, a_\beta)$, where a is his action, and a_α and a_β are the fractions of α and β players choosing action 1. Note that α and β players have the same payoff function. Since there are many players, by the law of large numbers we have, in state $(\rho_\alpha, \rho_\beta)$,

$$\begin{aligned} a_\alpha(\rho_\alpha, \rho_\beta) &= A_\alpha(\rho_\alpha, \rho_\beta), \\ a_\beta(\rho_\beta) &= A_\beta(\rho_\beta). \end{aligned}$$

I impose the following assumption on payoffs:

[2] The analysis in the paper goes through if, instead of focusing on symmetric mixed strategies, we were to focus on pure strategies in which players condition on their index in such a way so that the same fractions of player types choose the same actions. By the law of large numbers, this can always be done. E.g., we could have instead of the symmetric mixed strategy $A_\beta(\rho_\beta)$, we could have

$$A_{\beta i}(\rho_\beta) = \begin{cases} 1 & i \leq A_\beta(\rho_\beta) \\ 0 & i > A_\beta(\rho_\beta), \end{cases}$$

where i is the index of the player. Under the appropriate assumptions on i (namely that it be uniformly distributed on $[0,1]$) the two approaches are equivalent.

Assumption 1. *$U(1, a_\alpha, a_\beta) > U(0, a_\alpha, a_\beta)$ if and only if $a_\alpha > a_\beta$.*

Assumption 1 formalizes the intuition that all players strictly prefer an action if more α than β players are choosing that action. This may be because players have reputational concerns for appearing to be α players, conditional on the action they take. Or, players may have intrinsic preferences for choosing the same actions as αs (Akerlof and Kranton, 2000), because they identify with αs but not with βs. Or, after the stage game, players may go on to play a subgame, the payoffs to which depend on the action chosen in the stage game, as illustrated in the following example.

Example 3 (Matching utility)**.** *Interpret the action as a location choice, $a = 0$ or $a = 1$, and introduce a matching subgame in which players are randomly matched to another player at the location they choose. They then receive a payoff of 1 if their partner is an α type, and a payoff of 0 otherwise. The expected payoff of a player at location a the probability that a randomly drawn player is an α, conditional on the action taken,*

$$U(a, a_\alpha, a_\beta) = \begin{cases} \frac{a_\alpha}{a_\alpha + a_\beta} & a = 1 \\ \frac{1 - a_\alpha}{1 - a_\alpha + 1 - a_\beta} & a = 0, \end{cases} \tag{2}$$

On the boundaries, when $a_\alpha = a_\beta = 0$ or $a_\alpha = a_\beta = 1$, assume $U(a, a_\alpha, a_\beta) = 0 \, \forall a$.[3]

The payoff specification in Example 3 will be maintained as an example for the remainder of the paper.

3.1 Equilibria of the stage game

To fix ideas, I derive Bayesian equilibria of the static stage game. There are many such equilibria.

[3] It is only important that $U(1, a_\alpha, a_\beta) = U(0, a_\alpha, a_\beta)$. On the boundaries, these payoffs technically violate Assumption 1, since U is constant in a_β when $a_\alpha = 0$.

The game has trivial equilibria, indexed by $p \in [0,1]$, in which both sorts of players disregard their private information and mix independently between 0 and 1, each choosing $a = 1$ with probability p. This follows from Assumption 1, which implies $U(1,p,p) = U(0,p,p)\,\forall p \in [0,1]$, and so mixing is trivially a best response when $A_\alpha(\rho_\alpha, \rho_\beta) = A_\beta(\rho_\beta) = p$.

On the other hand, in no equilibria do α players coordinate on action 1 while β's mix, i.e., $A_\alpha(\rho_\alpha, \rho_\beta) \equiv 1, A_\beta(\rho_\beta) \in (0,1)\,\forall \rho_\beta$. In this case, it is commonly known that $A_\alpha(\rho_\alpha, \rho_\beta) = 1$, and so $a_\alpha > a_\beta$, hence; by Assumption 1, $U(1,1,a_\beta) > U(0,1,a_\beta)$, and so β players strictly prefer to choose action 1.

Are there equilibria in which the α players condition on their additional private information in a non-trivial way? More precisely, are there equilibria in which the α players, with positive probability, choose a different action profile than the β players? The answer is yes. To see this, first, we note that in any equilibrium, α players are always either all playing 0, all playing 1, or mixing with the same frequency as β players:

Lemma 1. *Let $\langle A_\alpha, A_\beta \rangle$ be an equilibrium strategy profile. Then $A_\alpha(\rho_\alpha, \rho_\beta) \in \{0, A_\beta(\rho_\beta), 1\}$ for all $(\rho_\alpha, \rho_\beta)$.*

Proof. The result follows directly from Assumption 1. If $A_\alpha(\rho_\alpha, \rho_\beta) < A_\beta(\rho_\beta)$, then $U(1, a_\alpha, a_\beta) < U(0, a_\alpha, a_\beta)$, and so $A_\alpha(\rho_\alpha, \rho_\beta) = 0$ is the unique best reply. Similarly, if $A_\alpha(\rho_\alpha, \rho_\beta) > A_\beta(\rho_\beta)$, then $A_\alpha(\rho_\alpha, \rho_\beta) = 1$ is the unique best reply. Finally, if $A_\alpha(\rho_\alpha, \rho_\beta) = A_\beta(\rho_\beta)$ the result holds. □

Second, we note that in any equilibrium, β players are indifferent between actions:

Lemma 2. *Let $\langle A_\alpha, A_\beta \rangle$ be an equilibrium strategy profile. Then β players are always indifferent between actions 0 and 1. That is, the following indifference condition is satisfied for* *all* $\rho_\beta \in [0,1]$*:*

$$P(A_\alpha(\tilde{\rho}_\alpha, \tilde{\rho}_\beta) = 0 \mid \rho_\beta) \times \big(U(0,0,A_\beta(\rho_\beta)) - U(0,1,A_\beta(\rho_\beta))\big)$$
$$= P(A_\alpha(\tilde{\rho}_\alpha, \tilde{\rho}_\beta) = 1 \mid \rho_\beta) \times \big(U(1,1,A_\beta(\rho_\beta)) - U(0,1,A_\beta(\rho_\beta))\big) \quad (3)$$

Proof. Fix $\hat{\rho}_\beta \in [0,1]$. Either $A_\beta(\hat{\rho}_\beta) = 0$, or $A_\beta(\hat{\rho}_\beta) \in (0,1)$, or $A_\beta(\hat{\rho}_\beta) = 1$.

First, if $A_\beta(\hat{\rho}_\beta) \in (0,1)$, the β players are mixing and so indifferent between actions.

Second, if $A_\beta(\hat{\rho}_\beta) = 0$, then by Lemma 1, $A_\alpha(\rho_\alpha, \hat{\rho}_\beta) \in \{0,1\}$ for all $\{\rho_\alpha \mid (\rho_\alpha, \hat{\rho}_\beta) \in \text{Supp}(\mu)\}$. Define p^0, p^1 by

$$p^a := \mathrm{P}_\mu[A_\beta(\tilde{\rho}_\alpha, \tilde{\rho}_\beta) = a \mid \tilde{\rho}_\beta = \hat{\rho}_\beta],\ a \in \{0,1\}.$$

Then a β player's expected payoff from $a = 0$ is weakly less than his expected payoff from $a = 1$, since by Assumption 1

$$p^0 U(0,0,0) + p^1 U(0,1,0) = p^0 U(1,0,0) + p^1 U(0,1,0)$$
$$\leq p^0 U(1,0,0) + p^1 U(1,1,0),$$

with strict inequality if and only if $p^1 = 0$. The best response condition then implies that β's are indifferent between $a = 0$ and $a = 1$.

The case in which $A_\beta(\hat{\rho}_\beta) = 1$ is analogous. To write the indifference condition, a β player's expected payoffs from action a is

$$\mathrm{E}[U(a, A_\alpha(\tilde{\rho}_\alpha, \tilde{\rho}_\beta), A_\beta(\tilde{\rho}_\beta)) \mid \rho_\beta]$$
$$= p^1 U(a,1,A_\beta(\rho_\beta)) + p^0 U(a,0,A_\beta(\rho_\beta)) + (1 - p^0 - p^1) U(a, A_\beta(\rho_\beta), A_\beta(\rho_\beta)),$$

Equating the expected payoff from $a = 1$ to $a = 0$ and rewriting with Assumption 1 (which implies $U(1, A_\beta(\rho_\beta), A_\beta(\rho_\beta)) = U(0, A_\beta(\rho_\beta), A_\beta(\rho_\beta))$ always) yields (3). $\square$

Combining Lemmas 1 and 2 together yields the following characterization of equilibria. For ease of exposition, we now assume that μ has full support on $[0,1]^2$ and is absolutely continuous with respect to Lebesgue measure on the unit square.[4]

Proposition 3. *$\langle A_\alpha, A_\beta \rangle$ is an equilibrium of the static game if and only if, for all $(\rho_\alpha, \rho_\beta) \in [0,1]^2$,*

1. *$A_\alpha(\rho_\alpha, \rho_\beta) \in \{0, A_\beta(\rho_\beta), 1\}$ and*

2. *A_β solves equation (3).*

Proof. That these conditions are necessary follows directly from Lemmas 1 and 2 and the fact that μ has full support on $[0,1]^2$.

To see that these conditions are also sufficient, note first that if all α players are choosing action $a = 1$, then by Assumption 1, $U(1, 1, a_\beta) \geq U(0, 1, a_\beta)\, \forall a_\beta \in [0,1]$, so that choosing $a = 1$ is a best reply. Similar reasoning suffices to show that when all α players choose $a = 0$, it is a best reply for all α players to choose $a = 0$. Finally, when $A_\alpha(\rho_\alpha, \rho_\beta) = A_\beta(\rho_\beta)$, by Assumption 1 α players are indifferent between actions, and so are best responding. This establishes that $A_\alpha(\rho_\alpha, \rho_\beta) \in \{0, 1\}$ is a sufficient condition for α players to be best responding. For β players, equation (3) is the indifference condition, and so every action is a best reply. $\square$

For example, consider the special case in which utility is given by the matching subgame, example 3. Let $\rho_\beta \mapsto g(\rho_\beta)$ be any mapping into $[0,1]$ satisfying

$$F_{\rho_\beta}(g(\rho_\beta)) \in (\tfrac{1}{3}, \tfrac{2}{3})\, \forall \rho_\beta \in [0,1], \tag{4}$$

[4] This assumption is mostly for notational convenience, the following statements continue to hold for more general probability distributions, but only on the support of μ.

where F_{ρ_β} is the cumulative distribution function of the marginal distribution of $\tilde{\rho}_\alpha$ for some fixed value of ρ_β. Proposition 3 characterizes equilibria in which α players choose 1 if $\rho_\alpha \geq g(\rho_\beta)$, and 0 otherwise.

Proposition 4. *Say α's play according to*

$$A_\alpha(\rho_\alpha, \rho_\beta) = \begin{cases} 1 & \rho_\alpha \geq g(\rho_\beta) \\ 0 & \rho_\alpha < g(\rho_\beta), \end{cases} \tag{5}$$

and β's play according to

$$A_\beta(\rho_\beta) = 2 - 3F_{\rho_\beta}(g(\rho_\beta)). \tag{6}$$

Then under the matching specification of utility, (2)*, $\langle A_\alpha, A_\beta \rangle$ is an equilibrium.*

Proof. The proof proceeds by showing that (6) solves the indifference condition (3). It then follows from Proposition 3 that $\langle A_\alpha, A_\beta \rangle$ is an equilibrium, since by construction, A_α satisfies $A_\alpha(\rho_\alpha, \rho_\beta) \in \{0, A_\beta(\rho_\beta), 1\} \forall \rho_\alpha, \rho_\beta \in [0,1]^2$.

To that end, note that by construction of the α player's strategy, (5), we have

$$P_\mu(A_\alpha(\tilde{\rho}_\alpha, \tilde{\rho}_\beta) = 0 \mid \rho_\beta) = F_{\rho_\beta}(g(\rho_\beta))$$
$$P_\mu(A_\alpha(\tilde{\rho}_\alpha, \tilde{\rho}_\beta) = 1 \mid \rho_\beta) = 1 - F_{\rho_\beta}(g(\rho_\beta)).$$

The indifference condition then becomes

$$F_{\rho_\beta}(g(\rho_\beta)) \left(\frac{1}{2 - A_\beta(\rho_\beta)} \right) = (1 - F_{\rho_\beta}(g(\rho_\beta))) \left(\frac{1}{1 + A_\beta(\rho_\beta)} \right),$$

solving yields (6), which is a well-defined strategy only when condition 4 is satisfied, which, by assumption, it is. (When $F_{\rho_\beta}(g(\rho_\beta)) \notin (\frac{1}{3}, \frac{2}{3})$, β players strictly prefer one or the other

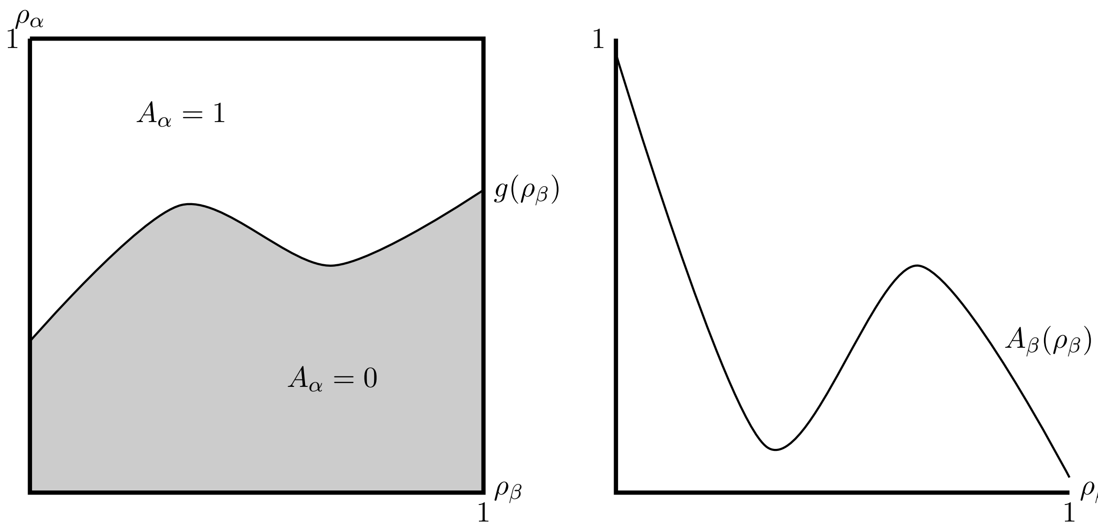

Figure 2: Graphical illustration of equilibria in the static stage game.

Note: On the left: α players coordinate using the (arbitrary) boundary $(\rho_\beta, g(\rho_\beta))$. On the right: In equilibrium, β players are more likely to choose action 1 when they believe α players are more likely to choose that action.

action, which, by Lemma 2, cannot be an equilibrium.) □

This equilibrium is one in which the αs have exclusive access to some hidden 'sunspot', which allows them to coordinate on an action, represented by the value of ρ_α. (Since ρ_β only ever takes one value, the β-type players have access to no private information.) For example, imagine that ρ_β represents broadcast television, while ρ_α represents cable television, and people who can afford to watch cable television would like to choose the same action as other people who can afford to watch cable television. The model suggests they can do so by coordinating on the information they see through cable television. In equilibrium, β players are aware that this coordination is taking place, and must draw inferences about which action α players are coordinating on.

The function $g(\rho_\beta)$ determines how likely it is that α players will choose action $a = 1$. The lower is $g(\rho_\beta)$, the greater the probability that, conditional on ρ_β, the α players are coordinating on action $a = 1$. The β players equilibrium strategy, (6), is also decreasing in $g(\rho_\beta)$. In this sense, β players are mimicking α players, but the degree to which this occurs is driven by β players beliefs about the extent to which other β players are choosing an action.

What are the expected payoffs to players under the strategy profile described in Proposition 4? Let V_α and V_β denote the *ex-ante* expected payoffs to each type of player, i.e.,

$$V_\alpha := \mathrm{E}_\mu[A_\alpha(\tilde{\rho}_\alpha, \tilde{\rho}_\beta)U(1, A_\alpha(\tilde{\rho}_\alpha, \tilde{\rho}_\beta), A_\beta(\tilde{\rho}_\beta)) + (1 - A_\alpha(\tilde{\rho}_\alpha, \tilde{\rho}_\beta))U(0, A_\alpha(\tilde{\rho}_\alpha, \tilde{\rho}_\beta), A_\beta(\tilde{\rho}_\beta))]$$

$$V_\beta := \mathrm{E}_\mu[A_\beta(\tilde{\rho}_\beta)U(1, A_\alpha(\tilde{\rho}_\alpha, \tilde{\rho}_\beta), A_\beta(\tilde{\rho}_\beta)) + (1 - A_\beta(\tilde{\rho}_\beta))U(0, A_\alpha(\tilde{\rho}_\alpha, \tilde{\rho}_\beta), A_\beta(\tilde{\rho}_\beta))].$$

Then α players can do well, as summarized in the following result:

Proposition 5. *Let $\langle A_\alpha, A_\beta \rangle$ be any equilibrium in which A_α is determined by (5) for some boundary g. Then*

$$V_\alpha = 2V_\beta,$$

and under the matching specification of utility, (2), $V_\alpha = \frac{2}{3}$ and $V_\beta = \frac{1}{3}$.

Proof. The full proof is in Appendix A; here I present an intuitive proof. By construction, β players are indifferent between strategies, by construction; hence, a β player who randomizes 50/50 between actions receives the same payoff *ex-ante* as in equilibrium. Half the time, such a player chooses the same action as an α player and receives V_α, half the time he does not and receives a payoff of zero, therefore,

$$V_\beta = \tfrac{1}{2}V_\alpha.$$

When payoffs are induced by the matching subgame, (2), the payoff to a player is the expected probability that the player is an α, conditional on his action, so, the *ex-ante* probability that a player is an α should equal the expected payoff of players, and since the probability that a player is an α player is $\frac{1}{2}$,

$$\frac{1}{2} = \frac{1}{2}V_\alpha + \frac{1}{2}V_\beta = V_\beta + \frac{1}{2}V_\beta = \frac{3}{2}V_\beta,$$

so $V_\beta = \frac{1}{3}$ and $V_\alpha = \frac{2}{3}$. $\square$

Surprisingly, the payoff to an α player in equilibrium is independent of the strategy A_α, as long as $A_\alpha(\rho_\alpha, \rho_\beta) \in \{0, 1\}$.

In the static game presented, the information structure is taken to be exogenous. In the next section, I augment the model so that state variable evolves and is allowed to depend on past actions, in such a way so that over time β players can eventually learn which action α players are coordinating on. Do the features of the static game carry over into the dynamic environment? In sections 4 and 5 I show that they do: α players can still condition on their private information, and can extract the same payoffs in the dynamic game that they do in the static game. I show that this occurs when behavior is cyclical — in order to extract the same payoffs, α players need to periodically switch between actions. Furthermore, I show how the dynamic structure of the game induces particular forms for $g(\rho_\beta), \mu$, and the β's strategy; I give intepretations of these forms and analyze comparative statics.

4 Dynamic game

Now, time is continuous, $t \in [0, \infty)$. In each instant t, a continuum of short-lived players play the static game. The state space is still the unit square, $[0, 1]^2$, but elements of the state space are now denoted $\vec{\rho}^{\,t} = (\rho^t_\alpha, \rho^t_\beta)$, to represent the dependence on time. I maintain the assumption that both αs and βs observe the value of ρ^t_β, but that only αs observe ρ^t_α. Again, imagine that ρ^t_β represents a source of information available to everyone in the game, such broadcast television, while ρ^t_α represents a source of information available only to a subset of the population, such as cable television.

The game proceeds as follows. First, an initial condition is drawn from $\mu^0 \in \Delta[0,1]^2$. I denote the random variable by $\tilde{\rho}$, and ρ^0 is the realized initial condition. Given an initial condition ρ^0, the game *outcome path* is action paths $a_\alpha^t(\rho^0), a_\beta^t(\rho^0)$ and state paths $\rho_\alpha^t(\rho^0), \rho_\beta^t(\rho^0)$ for $t \in [0,\infty)$, where, analogously to the static game, a_α^t and a_β^t are the fractions of α and β players choosing action $a = 1$ at time t. For convenience, the dependence on the (stochastic) initial conditions will be omitted where it is clear, so that the (stochastic) value of the outcome path at time t is denoted $(\tilde{a}_\alpha^t, \tilde{a}_\beta^t, \tilde{\rho}_\alpha^t, \tilde{\rho}_\beta^t)$ to represent $(a_\alpha^t(\tilde{\rho}_\alpha^0, \tilde{\rho}_\beta^0), a_\beta^t(\tilde{\rho}_\alpha^0, \tilde{\rho}_\beta^0), \rho_\alpha^t(\tilde{\rho}_\alpha^0, \tilde{\rho}_\beta^0), \rho_\beta^t(\tilde{\rho}_\alpha^0, \tilde{\rho}_\beta^0))$.

It will be of interest to consider two special sorts of outcome paths. A point ρ^* is a *fixed point* under $\vec{\rho}^{\,t}$ if

$$\vec{\rho}^{\,t}(\rho^*) = \rho^* \, \forall t \in [0,\infty).$$

A point ρ^* is a *periodic point* if there exists $P > 0$ such that

$$\rho^P(\rho^*) = \rho^*,$$

and the minimum such P is called the *period.* The set of states traced out by the state path is called the *orbit.* An orbit is called *fixed* if it consists of a single fixed point, and *periodic* if it consists entirely of periodic points with some period $P > 0$.

Information structure

We would now like to impose a dynamic structure on players' private information to capture the idea that $\vec{\rho}^{\,t}$ is somehow representative of actions taken in the past. To do so, I now impose the following interpretation on the meaning of $\vec{\rho}^{\,t}$. When players make an action choice, I assume that the choice (but not the player's type, or the time at which they chose the action) is visible for some time afterward. Although a player's type is not known, I assume one type's action choices may be more visible than others. Imagine, for example, that once a player chooses which style of clothing to wear, they wear it for some exogenously stochastic amount of time before replacing it, at which point their choice is no longer visible.[5] The value of ρ_α^t represents the fraction of players observed to be 'wearing' $a = 1$ at some location where the replacement rate of old clothing is higher than at some other location, represented by ρ_β^t. I assume that old action choices disappear at some exogenous Poisson rate r_α and r_β, respectively. Formally, the change in the fraction of players observed to have been choosing $a = 1$, for small time increments, evolves approximately according to

$$\rho_\alpha^{t+\varepsilon} \approx (1 - \varepsilon r_\alpha)\rho_\alpha^t + \varepsilon r_\alpha(\lambda_\alpha a_\alpha^t + \lambda_\beta r_\beta^t)$$
$$\rho_\beta^{t+\varepsilon} \approx (1 - \varepsilon r_\beta)\rho_\beta^t + \varepsilon r_\beta(\lambda_\alpha a_\alpha^t + \lambda_\beta r_\beta^t).$$

Here, $\lambda_\alpha a_\alpha^t + \lambda_\beta a_\beta^t$ is some average of the actions being taken by each player type. λ_α and λ_β are *intratemporal* weights, parameterizing how visible a particular player type's action choices are. Accordingly, we take $\lambda_\alpha + \lambda_\beta = 1$ and $\lambda_\alpha, \lambda_\beta > 0$. When λ_α is large, ρ_α^t and ρ_β^t mostly reflect the actions of high types, conversely, when λ_β is large, ρ_α^t and ρ_β^t mostly reflect the actions of low types.

While λ_α and λ_β are interpreted as intratemporal weights, r_α and r_β are *intertemporal* weights, adjusting how rapidly it is that old actions disappear. Consistent with the interpretation above, I assume that $r_\alpha > r_\beta > 0$, so that ρ_α^t represents a more 'up-to-date' average of the actions being taken.

[5] When the player replaces his clothing, we could imagine that he returns to the game and makes a new purchase decision, but 'forgets' how long it has been since the prior purchase decision. This is consistent with a focus on stationary equilibria in which a player's beliefs are independent of calendar time.

In the limit as $\varepsilon \to 0$, we derive

$$\dot{\rho}^t_\alpha = r_\alpha(\lambda_\alpha a^t_\alpha + \lambda_\beta a^t_\beta - \rho^t_\alpha) \tag{7}$$

$$\dot{\rho}^t_\beta = r_\beta(\lambda_\alpha a^t_\alpha + \lambda_\beta a^t_\beta - \rho^t_\beta)., \tag{8}$$

the laws of motion of $\vec{\rho}^{\,t}$.[6] A solution to (7) and (8) is

$$\rho^t_\alpha = e^{-r_\alpha t}\rho^0_\alpha + r_\alpha \int_0^t e^{-r_\alpha(\tau - t)}(\lambda_\alpha a^\tau_\alpha + \lambda_\beta a^\tau_\beta)\, d\tau \tag{9}$$

$$\rho^t_\beta = e^{-r_\beta t}\rho^0_\beta + r_\beta \int_0^t e^{-r_\beta(\tau - t)}(\lambda_\alpha a^\tau_\alpha + \lambda_\beta a^\tau_\beta)\, d\tau, \tag{10}$$

which makes explicit the fact that we assume that $\vec{\rho}^{\,t}$ is an exponentially-weighted moving average of past actions, with weights r_α and r_β, up to ρ^0, the initial condition.[7][8]

This specification of $\vec{\rho}^{\,t}$ captures the idea that players have some information about the actions of other players, but that this information is delayed, and does not immediately reflect changes in action choices. The assumption that $r_\alpha > r_\beta$ captures the idea that the αs have access to more up-to-date information than βs, since a higher value for r_α places more

[6] More generally, we might specify that $(\rho^t_\alpha, \rho^t_\beta)$ evolve according to some law of motion which depends on the action profile,

$$\dot{s}^t_\alpha = f_\alpha(\rho^t_\alpha, a^t_\alpha, a^t_\beta)$$
$$\dot{s}^t_\beta = f_\beta(\rho^t_\beta, a^t_\alpha, a^t_\beta),$$

in which case $f_\alpha(\rho_\alpha, a_\alpha, a_\beta) = r_\alpha(\rho_\alpha - \lambda_\alpha a_\alpha - \lambda_\beta a_\beta)$, $f_\beta(\rho_\beta, a_\alpha, a_\beta) = r_\beta(\rho_\beta - \lambda_\alpha a_\alpha - \lambda_\beta a_\beta)$ corresponds to an exponentially weighted moving average; the assumption that ρ_α, ρ_β are exponentially weighted moving averages is tractable and has a simple interpretation compared to the general case.

[7] There is some ambiguity as to the meaning of a solution to a discontinuous differential equation, which (7), (8) may be. Here, I mean a Carathéodory solution, that is, the solution should satisfy

$$s^t = \int_0^t \dot{s}^\tau \, d\tau + s^0.$$

[8] An alternate way to have set up the model would been to have had time begin at $-\infty$, which motivates an interpretation of $\rho^0_\alpha, \rho^0_\beta$ as representing, in some reduced-form way, the state of the system at time 0.

weight on more recent actions.[9]

Strategies

A strategy profile is now functions mapping calendar time and the observed state variable into the probability of choosing action $a = 1$, denoted $A^t_\alpha(\rho_\alpha, \rho_\beta), A^t_\beta(\rho_\beta)$. A strategy is *stationary* means it is independent of calendar time, $A^t_\alpha = A_\alpha, A^t_\beta = A_\beta \forall t \in [0, \infty)$.

A strategy profile $\langle A^t_\alpha, A^t_\beta \rangle$, induces an action path through—analogously to the static setting—the conditions

$$a^t_\alpha = A^t_\alpha(\rho^t_\alpha, \rho^t_\beta) \tag{11}$$

$$a^t_\beta = A^t_\beta(\rho^t_\beta) \, \forall t \in [0, \infty) \tag{12}$$

from the law of large numbers; it induces a state path through the laws of motion (9), (10).[10]

Payoffs

For a fixed outcome path, the payoffs to a player in period t are the same as in the static game, that is, if a player chooses action a at time t, his payoff is $U(a, a^t_\alpha, a^t_\beta)$, satisfying Assumption 1. A player who sees private information s updates his beliefs over initial conditions, as previously noted, for every initial condition there is a unique outcome path,

[9] An exponentially weighted moving average is, of course, simply one out of many which we could have chosen.

[10] Outcome paths satisfying (11), (12), (9), and (10) are neither guaranteed to exist, nor to be unique. In the case where an outcome path does not exist, as may occur, for example, if A_α or A_β are not measurable functions, then payoffs may be assumed to be $-\infty$, but we will not consider equilibria with this property. In the case where the outcome path is not unique, one may be selected according to any arbitrary rule. For example, we might select one consistent with a discrete-time approximation to the model. Since any individual player's deviations do not affect the outcome path, the precise rule selected is unimportant. For the remainder of this paper, all strategies are implicitly taken to be measurable functions, so that (9), (10) are well-defined, and each strategy profile is considered to induce to a unique outcome path. Lemma 7 characterizes the way in which there may be multiple outcome paths, briefly, an outcome path's orbit may contain both a fixed point and a periodic point, however, behavior on the outcome path is still consistent with the equilibrium characterizations derived in this paper.

hence, beliefs over initial conditions induce beliefs over the value of the outcome path at every time t. I denote random variables with tildes and realizations without tildes.

Equilibrium

An equilibrium consists of a strategy profile and beliefs over the state variable at each time t, denoted $\langle A^t_\alpha, A^t_\beta, \mu^t\rangle$, such that players are best responding to their beliefs, and beliefs are consistent with the outcome path induced by $\langle A^t_\alpha, A^t_\beta\rangle$.

Formally, players are best responding given beliefs when (note that the expectation operator is omitted for α players, since their belief updating process is trivial)

$$\operatorname{Supp}(A^t_\alpha(\rho^t_\alpha, \rho^t_\beta)) \subseteq \underset{a\in\{0,1\}}{\arg\max}\ U(a, a^t_\alpha, a^t_\beta) \tag{13}$$

$$\operatorname{Supp}(A^t_\beta(\rho^t_\beta)) \subseteq \underset{a\in\{0,1\}}{\arg\max}\ \mathrm{E}_{\mu^t}[U(a, \tilde{a}^t_\alpha, \tilde{a}^t_\beta) \mid \tilde{\rho}^t_\beta = \rho_\beta]\, \forall t \in [0, \infty) \tag{14}$$

and beliefs are consistent with equilibrium behavior when, for all measurable subsets $S \subset [0,1]^2$,

$$\mu^0(S) = \mu^t(\rho^t_\alpha(S), \rho^t_\beta(S))\, \forall t \in [0, \infty). \tag{15}$$

An equilibrium is *stationary* means the strategy profile and beliefs are stationary, so that $A^t_\alpha = A_\alpha, A^t_\beta = A_\beta$, and $\mu^t = \mu\, \forall t \in [0, \infty)$. Stationarity does not imply stationarity of the state path, just that players beliefs should be independent of the time they entered the game.

4.1 Equilibria of the dynamic game

I formally state a characterization of stationary equilibria in this game and discuss its implications. The full proof is in the appendix.

Proposition 6. *Let $\langle A^t_\alpha, A^t_\beta, \mu^t \rangle$, be an equilibrium. Then, for all t,*

1. *A^t_α satisfies*

$$A^t_\alpha(\tilde{\rho}_\alpha, \tilde{\rho}_\beta) \in \{0, A^t_\beta(\tilde{\rho}_\beta), 1\} \text{ and} \tag{16}$$

2. *$A^t_\beta(\rho_\beta)$ solves with probability 1*

$$\begin{aligned} &P_{\mu^t}(A_\alpha(\tilde{\rho}^t_\alpha, \tilde{\rho}^t_\beta) = 0 \mid \tilde{\rho}^t_\beta = \rho_\beta) \times (U(0, 0, A^t_\beta(\rho_\beta)) - U(1, 0, A^t_\beta(\rho_\beta))) \\ &\qquad = P_{\mu^t}(A_\alpha(\tilde{\rho}^t_\alpha, \tilde{\rho}^t_\beta) = 1 \mid \tilde{\rho}^t_\beta = \rho_\beta) \times (U(1, 1, A^t_\beta(\rho_\beta)) - U(0, 1, A^t_\beta(\rho_\beta))) \end{aligned} \tag{17}$$

Furthermore, these conditions are sufficient for equilibria, in the following sense: Let $\langle A^t_\alpha, A^t_\beta \rangle$ be a strategy profile, and say μ^t are probability distributions over $[0,1]^2$ consistent with the outcome path induced by $\langle A^t_\alpha, A^t_\beta \rangle$. If for all $t \in [0, \infty)$, $\rho_\alpha, \rho_\beta \in Supp(\mu^t)$, it is the case that $A_\alpha(\rho_\alpha, \rho_\beta)$ satisfies (16) and $A_\beta(\rho_\beta)$ satisfies (17), then $\langle A^t_\alpha, A^t_\beta, \mu^t \rangle$, is an equilibrium.

Proposition 6 is simply a re-statement of Lemmas 1 and 2, from the static environment in the dynamic environment. If we require equilibria to be stationary, we can derive stronger results than Proposition 6:

Proposition 7. *Let $\langle A_\alpha, A_\beta \rangle$, μ be a stationary equilibrium. Then, for all $\vec{\rho} \in Supp(\mu)$, A_α satisfies (16), A_β solves*

$$\begin{aligned} &|\lambda_\alpha + \lambda_\beta A_\beta(\rho_\beta) - \rho_\beta| \times \big(U(0, 0, A_\beta(\rho_\beta)) - U(1, 0, A_\beta(\rho_\beta))\big) \\ &\qquad = |\lambda_\beta A_\beta(\rho_\beta) - \rho_\beta| \times \big(U(1, 1, A_\beta(\rho_\beta)) - U(0, 1, A_\beta(\rho_\beta))\big) \end{aligned} \tag{18}$$

almost surely, and $\lambda_\alpha + \lambda_\beta A_\beta(\rho_\beta) - \rho_\beta \geq 0$ and $\lambda_\beta A_\beta(\rho_\beta) - \rho_\beta \leq 0$.

Furthermore, these conditions are sufficient for equilibria, in the following sense: Let

$\langle A_\alpha, A_\beta\rangle$ be a stationary strategy profile. If for all $\vec{\rho} \in [0,1]^2$, it is the case that $A_\alpha(\rho_\alpha, \rho_\beta)$ satisfies (16) and $A_\beta(\rho_\beta)$ satisfies (18), then there exists a probability measure μ on $[0,1]^2$ such that $\langle A_\alpha, A_\beta, \mu\rangle$, is a stationary equilibrium.

Propositions 6 and 7 look similar, and so it is worthwhile to consider their differences. Proposition 6 is the dynamic version of Proposition 3, and the proof is similar. In Proposition 3, μ was taken to be exogenous. Proposition 6 has nothing further to say about μ^t beyond the equilibrium requirement that it be consistent with player's behavior. For stationary equilibria, however, it is possible to say more about μ. Specifically, the marginal distributions $\mathrm{P}_{\mu^t}(A_\alpha(\tilde{\rho}^t_\alpha, \tilde{\rho}^t_\beta) = 0 \mid \tilde{\rho}^t_\beta = \rho_\beta)$ and $\mathrm{P}_{\mu^t}(A_\alpha(\tilde{\rho}^t_\alpha, \tilde{\rho}^t_\beta) = 1 \mid \tilde{\rho}^t_\beta = \rho_\beta)$ may be characterized, which yields (18), and furthermore, given strategy profiles $\langle A_\alpha, A_\beta\rangle$ satisfying (16) and (18), Proposition 7 states that consistent equilibrium beliefs exist, while Proposition 6 has nothing to say about the existence of consistent beliefs for a given strategy profile.

A focus on stationary equilibria is often justified through an argument that they represent, in some way, the long-run of a non-stationary equilibrium of the game. Since I mainly focus on stationary equilibria for the rest of the paper, in Appendix B I show via numerical simulation that it is not unusual for non-stationary equilibrium behavior to result in convergence to a stationary equilibrium.

4.2 Instrumental preferences for conformity

Under the matching specification of utility, it is possible to explicitly characterize equilibrium strategy profiles in stationary equilibria for β players by applying (18) from Proposition 7:

Example 4 (Matching utility I cont.)**.** *Here,*

$$U(a, a_\alpha, a_\beta) = \begin{cases} \frac{a_\alpha}{a_\alpha + a_\beta} & a = 1 \\ \frac{1-a_\alpha}{1-a_\alpha+1-a_\beta} & a = 0, \end{cases}$$

when $a_\alpha \neq 0$ or $a_\beta \neq 0$, and 0 if $a_\alpha = a_\beta = 0$. Then equation (18) becomes

$$r_\beta(\rho_\beta - \lambda_\beta A_\beta(\rho_\beta)) \times \frac{1}{1+(1-A_\beta(\rho_\beta))} = r_\beta(\rho_\beta - \lambda_\alpha - \lambda_\beta A_\beta(\rho_\beta)) \times \frac{1}{1+A_\beta(\rho_\beta)}, \quad (19)$$

Solving yields

$$A_\beta(\rho_\beta) = \rho_\beta. \quad (20)$$

That is, under the matching specification of utility, β players mimic the actions they see have been taken in the past. It is not clear ex-ante *that β players should display this behavior, after all, they do not have direct preferences for conformity, in the sense that their payoffs are not necessarily increasing in the number of other players choosing the same action. Is this a general feature of the model, or is it specific to the matching specification of utility? In this section, I argue that it is a general feature of the model, in the sense that β players are more likely to take an action the more they see that other players have chosen that action in the past, in every stationary equilibrium.*

An interpretation of equation (20) is that the β players are 'endogenously' conformist. Their strategy could be interpreted as the players sampling an action from the social network represented by ρ_β, and mimicking it. In fact, this induced preference for conformity is a feature of the general model:

Proposition 8 (Instrumental preferences for conformity)**.** *Say $\langle A_\alpha, A_\beta, \mu\rangle$ is a stationary equilibrium. If ρ_β, ρ'_β are two draws from $\tilde{\rho}_\beta$ and $\rho'_\beta > \rho_\beta$, then with probability 1, $A_\beta(\rho'_\beta) \geq A_\beta(\rho_\beta)$.*

Proof. Pick $\rho_\beta, \rho'_\beta \in \text{proj}_\beta(\text{Supp}(\mu))$. If $\rho'_\beta = \rho_\beta$, then $A_\beta(\rho_\beta) = A_\beta(\rho'_\beta)$. So say $\rho'_\beta > \rho_\beta$, but $A_\beta(\rho'_\beta) < A_\beta(\rho_\beta)$. Then, by Assumption 1, we have

$$
\begin{aligned}
&U(1,1,A_\beta(\rho_\beta)) < U(1,1,A_\beta(\rho'_\beta)) \\
&U(0,1,A_\beta(\rho_\beta)) > U(0,1,A_\beta(\rho'_\beta)) \\
&U(0,0,A_\beta(\rho_\beta)) > U(0,0,A_\beta(\rho'_\beta)) \\
&U(1,0,A_\beta(\rho_\beta)) < U(1,0,A_\beta(\rho'_\beta)),
\end{aligned}
$$

and so

$$U(0,0,A_\beta(\rho_\beta)) - U(1,0,A_\beta(\rho_\beta)) > U(0,0,A_\beta(\rho'_\beta)) - U(1,0,A_\beta(\rho'_\beta)) \tag{21}$$

$$U(1,1,A_\beta(\rho_\beta)) - U(0,1,A_\beta(\rho_\beta)) < U(1,1,A_\beta(\rho'_\beta)) - U(0,1,A_\beta(\rho'_\beta)). \tag{22}$$

On the other hand, (recall by Proposition 7 that $\lambda_\beta A_\beta(\rho_\beta) - \rho_\beta \leq 0$, $\lambda_\beta A_\beta(\rho'_\beta) - \rho_\beta \leq 0$):

$$|\lambda_\alpha + \lambda_\beta A_\beta(\rho_\beta) - \rho_\beta| \leq \left|\lambda_\alpha + \lambda_\beta A_\beta(\rho'_\beta) - \rho_\beta\right| \tag{23}$$

$$|\lambda_\beta A_\beta(\rho_\beta) - \rho_\beta| \geq \left|\lambda_\beta A_\beta(\rho'_\beta) - \rho_\beta)\right|. \tag{24}$$

Together, (21), (22), (23), and (24) contradict (18), which holds with probability 1, and so it must be with probability 1 that $A_\beta(\rho'_\beta) \geq A_\beta(\rho_\beta)$, the desired result. □

That is, with two independent draws from $\tilde{\rho}_\beta$, it is almost certain that the β players will be more likely to choose $a = 1$ when the draw is higher. Proposition 8 is of interest because, as in Bernheim (1994), observed conformist behavior does not arise from a direct preference for conformity, rather, it arises from strategic incentives on the part of players to appear to have better information about the actions of other players. It provides a rational for why we might observe aesthetic preferences for conformity, and furthermore, why preferences for conformity might be viewed negatively by others, or associated with lower-class tastes

(Bourdieu, 1984).

5 Matching game application

In this section, I apply Proposition 7 to explicitly compute equilibria in the case where payoffs are determined by the matching subgame, as in example 3. Imagine that player's action choices are interpreted as a choice between locations (for example, $a = 0$ represents a bar on the east side of town, and $a = 1$ represents a bar on the west side of town). Once players have made the action choice, they travel to the location, and look for someone to match with. Matching with an α results in a payoff of 1, and matching with a β results in a payoff of 0. Say in addition that instead of unit masses of both sorts of players, there is a mass M_α of α players and M_β of β players. The payoff of a player who chooses action a is therefore derived using Bayes' rule as

$$\begin{aligned} U(a, a_\alpha, a_\beta) &= \mathrm{P}(\text{Meeting an } \alpha \mid \text{Choosing } a) \\ &= \begin{cases} \frac{M_\alpha a_\alpha}{M_\alpha a_\alpha + M_\beta a_\beta} & a = 1 \\ \frac{M_\alpha (1-a_\alpha)}{M_\alpha(1-a_\alpha) + M_\beta(1-a_\beta)} & a = 0. \end{cases} \end{aligned} \tag{25}$$

When $a_\alpha = a_\beta = 0$, or $a_\alpha = a_\beta = 1$, (25) is not well-defined when $a = 1$ or $a = 0$, say in this case that $U(1, 0, 0) = U(0, 1, 1) = \frac{M_\alpha}{M_\alpha + M_\beta}$, so that U satisfies Assumption 1.

There are many stationary equilibria of this game. By Lemma 7, in Appendix A, it is sufficient to focus on stationary equilibria with beliefs whose support consists of a single fixed point or a single periodic orbit. Furthermore, we seek equilibria which are *geometrically* *symmetric*, meaning

$$A_\alpha(\rho_\alpha, \rho_\beta) = 1 - A_\alpha(1 - \rho_\alpha, \rho_\beta)$$
$$A_\beta(\rho_\beta) = 1 - A_\beta(\rho_\beta),$$

(or geometrically, that strategy profiles should be invariant to 180° rotations). The restriction to geometrically symmetric equilibria is justified for two reasons: First, it is aesthetically pleasing; analysis of geometrically *a*symmetric equilibria may be done and produces similar results. Second, among strategy profiles in which α's always coordinate on 0 or 1, there is a unique geometrically symmetric strategy profile which produces the highest possible payoff for α players. (See Proposition 12 for a characterization of payoffs in this equilibrium.)

Fixed points

First, we characterize all stationary equilibria with beliefs whose support contains a fixed point. These are all equilibria $\langle A_\alpha, A_\beta, \delta_{(\rho^*_\alpha, \rho^*_\beta)}\rangle$, where $\delta_{(\rho^*_\alpha, \rho^*_\beta)}$ is a fixed point and δ is the Dirac measure which puts probability 1 on $(\rho^*_\alpha, \rho^*_\beta)$. The result is summarized in the following proposition:

Proposition 9. *$\langle A_\alpha, A_\beta, \delta_{(\rho^*_\alpha, \rho^*_\beta)}\rangle$ is a stationary equilibrium if and only if*

$$A_\alpha(\rho^*_\alpha, \rho^*_\beta) = A_\beta(\rho^*_\beta) = \rho^*_\alpha = \rho^*_\beta.$$

Proof. To show necessity, pick $s^* \in [0, 1]$, and say $\langle A_\alpha, A_\beta\rangle$ is a strategy profile satisfying

$$A_\alpha(s^*, s^*) = A_\beta(s^*) = s^*.$$

I claim that $\langle A_\alpha, A_\beta, \delta_{(s^*,s^*)}\rangle$ is a stationary equilibrium. Under this strategy profile, the state path under the sole possible realization of the state variable is

$$(\rho_\alpha^t, \rho_\beta^t) = (s^*, s^*)\,\forall t \geq 0,$$

and by Proposition 7, $\langle A_\alpha, A_\beta, \mu\rangle$ is a stationary equilibrium, since (16) is satisfied; when $s^* \in (0,1)$ equation (18) is satisfied since, e.g., $\mathrm{P}_\mu(A_\alpha(\tilde{\rho}_\alpha, \tilde{\rho}_\beta) = 1 \mid \tilde{\rho}_\beta = s^*) = 0$, and when $s^* \in \{0,1\}$, (18) is satisfied since, e.g., $U(1,1,A_\beta(s^*)) - U(0,1,A_\beta(s^*)) = 0$, by Assumption 1, so for all $s^* \in [0,1]$, both sides of the indifference condition (18) are zero. Hence $\langle A_\alpha, A_\beta, \delta_{(s^*,s^*)}\rangle$ is a stationary equilibrium.

To show sufficiency, say $\langle A_\alpha, A_\beta, \delta_{(\rho_\alpha^*,\rho_\beta^*)}\rangle$ is some stationary equilibrium. Note that in equilibrium the state variable is commonly known, hence, for β players to be indifferent we must have $A_\alpha(\rho_\alpha^*, \rho_\beta^*) = A_\beta(\rho_\beta^*)$. And by the law of motion (9), for $(\rho_\alpha^*, \rho_\beta^*)$ to be a fixed point, we must have

$$\begin{aligned}
0 &= r_\alpha(\lambda_\alpha A_\alpha(\rho_\alpha^*, \rho_\beta^*) + \lambda_\beta A_\beta(\rho_\beta^*) - \rho_\beta^*) \\
&= r_\alpha(A_\beta(\rho_\beta^*) - \rho_\beta^*) \\
&\implies A_\beta(\rho_\beta^*) = \rho_\beta^*.
\end{aligned}$$

Applying (10) and similar reasoning yields $A_\beta(\rho_\beta^*) = \rho_\alpha^*$, so $A_\alpha(\rho_\alpha^*, \rho_\beta^*) = A_\beta(\rho_\beta^*) = \rho_\alpha^* = \rho_\beta^*$, which is the desired result. $\square$

Periodic orbits

Now, in the more interesting case, we characterize all stationary equilibria with periodic orbits. To do so, we begin by solving condition (18), which when $A_\alpha(\rho_\alpha, \rho_\beta) \in \{0, 1\}$ becomes

$$(\lambda_\alpha + \lambda_\beta A_\beta(\rho_\beta) - \rho_\beta) \times M_\alpha/(M_\alpha + M_\beta(1 - a_\beta))$$
$$= -(\lambda_\beta A_\beta(\rho_\beta) - \rho_\beta) \times M_\alpha/(M_\alpha + M_\beta A_\beta(\rho_\beta)),$$

to obtain the functional form $A_\beta(\rho_\beta)$ must satisfy at all possible ρ_β in the support of μ. The result is summarized in Lemma 3:

Lemma 3. *Under the matching specification of utility,* (25)*, if* $A_\alpha(\rho_\alpha, \rho_\beta) \in \{0, 1\}$*, and*

$$A_\beta(\rho_\beta) = \frac{(2\frac{M_\alpha}{M_\beta} + 1)\rho_\beta - \frac{M_\alpha}{M_\beta}\lambda_\alpha}{2\frac{M_\alpha}{M_\beta}\lambda_\beta + 1}. \tag{26}$$

for all $(\rho_\alpha, \rho_\beta)$*, then there exists a probability measure* μ *such that* $\langle A_\alpha, A_\beta, \mu \rangle$ *is a stationary equilibrium.*

(26) may not result in a well-defined strategy profile on the entire state space (specifically, it may specify that $A_\beta(\rho_\beta) \notin [0, 1]$ at some ρ_β), hence, by Proposition 7, wherever this occurs we must have $A_\alpha(\rho_\alpha, \rho_\beta) = A_\beta(\rho_\beta)$. Finding stationary equilibria with periodic orbits is then a matter of choosing regions on which $A_\alpha = 1$ and $A_\alpha = 0$, solving (26) on the regions where it implies $A_\beta(\rho_\beta) \in [0, 1]$, and choosing $A_\alpha(\rho_\alpha, \rho_\beta) = A_\beta(\rho_\beta)$ on regions where (26) requires $A_\beta(\rho_\beta) \notin [0, 1]$. Together with (26) and the laws of motion (9), (10), this induces a dynamical system on the state space. Solving for a time average then yields μ. More specifically,

1. Guess an initial starting point, $(\rho_\alpha^0, \rho_\beta^0)$, and begin with $A_\alpha = 1$.

2. Solve forward the differential equation resulting from (26), the laws of motion (9), (10),

and the assumption that $A_\alpha(\rho_\alpha, \rho_\beta) = 1$,

$$\dot{s}^t_\alpha = r_\alpha \left(\lambda_\alpha + \lambda_\beta \left(\frac{(2\frac{M_\alpha}{M_\beta} + 1)\rho^t_\beta - \frac{M_\alpha}{M_\beta}\lambda_\alpha}{2\frac{M_\alpha}{M_\beta}\lambda_\beta + 1} \right) - \rho^t_\alpha \right) \tag{27}$$

$$\dot{s}^t_\beta = r_\beta \left(\lambda_\alpha + \lambda_\beta \left(\frac{(2\frac{M_\alpha}{M_\beta} + 1)\rho^t_\beta - \frac{M_\alpha}{M_\beta}\lambda_\alpha}{2\frac{M_\alpha}{M_\beta}\lambda_\beta + 1} \right) - \rho^t_\beta \right) \tag{28}$$

on some time interval, $[0, P^0)$, to yield $(\rho^t_\alpha, \rho^t_\beta)$ on $t \in [0, P^0)$.

3. At $t = P^0$, take the point $(\rho^{P^0}_\alpha, \rho^{P^0}_\beta)$ as an initial condition, and repeat the process, but now with $A_\alpha = 0$, to yield $(\rho^t_\alpha, \rho^t_\beta)$ on $[P^0, P^1)$.

4. Check whether the stationarity condition, $(\rho^{P^1}_\alpha, \rho^{P^1}_\beta) = (\rho^0_\alpha, \rho^0_\beta)$ is satisfied, if it is, then a stationary equilibrium has been found, with period $P^0 + P^1$, in which α players play a strategy satisfying

$$A_\alpha(\rho_\alpha, \rho_\beta) = \begin{cases} 1 & (\rho_\alpha, \rho_\beta) \in \{(\rho^t_\alpha, \rho^t_\beta) \mid t \in [0, P^0)\} \\ 0 & (\rho_\alpha, \rho_\beta) \in \{(\rho^t_\alpha, \rho^t_\beta) \mid t \in [P^0, P^1)\}, \end{cases}$$

and β players play according to (26).

For example, in the special case in which $M_\alpha = M_\beta = \lambda_\alpha = \lambda_\beta = \frac{1}{2}$, we might look for a geometrically symmetric equilibrium by fixing an initial starting value $\rho^0_\beta < \frac{1}{2}$, and guessing ρ^0_α as the corresponding value in the α dimension. The solution to the differential equation (28) may be obtained using standard methods as

$$\rho^t_\beta = e^{-\frac{\lambda_\alpha r_\beta}{2\lambda_\beta + 1}t}(\rho^0_\beta - (1 + \lambda_\beta)). \tag{29}$$

Substituting (29) into (27) yields an ordinary differential equation,

$$\dot{s}_\alpha^t = r_\alpha \left(\lambda_\alpha + \lambda_\beta \left(\frac{(2\frac{M_\alpha}{M_\beta}+1)e^{-\frac{\lambda_\alpha r_\beta}{2\lambda_\beta+1}t}(\rho_\beta^0 - (1+\lambda_\beta)) - \frac{M_\alpha}{M_\beta}\lambda_\alpha}{2\frac{M_\alpha}{M_\beta}\lambda_\beta + 1} \right) - \rho_\alpha^t \right) \tag{30}$$

which may also be solved using standard methods. We could proceed by taking $(\rho_\alpha^{P/2}, \rho_\beta^{P/2})$ as the initial conditions and repeating this process, but by our assumption of geometric symmetry, this is equivalent to $1-\rho_\beta^0 = \rho_\beta^{P/2}$, and $1-\rho_\alpha^0 = \rho_\alpha^{P/2}$, which yields two stationarity conditions from (29) and the solution to (30). Specifically, the stationarity condition resulting from (29) is

$$1 - \rho_\beta^0 = e^{-\frac{\lambda_\alpha r_\beta}{2\lambda_\beta+1}\frac{P}{2}}(\rho_\beta^0 - (1+\lambda_\beta)),$$

which, solving for P, implies the period length is

$$P = \frac{2(1+2\lambda_\beta)}{r_\beta \lambda_\alpha} \log\left(\frac{1-\rho_\beta^0+\lambda_\beta}{\lambda_\beta+\rho_\beta^0} \right).$$

Analogously, solving the stationarity condition $1-\rho_\alpha^0 = \rho_\alpha^{P/2}$ yields ρ_α^0 as a function of ρ_β^0, the point at which α players should switch from $a=0$ to $a=1$. The solution in this case is plotted numerically in Figure 3. The solid curve, $\rho_\alpha^{\text{switch}}(\rho_\beta)$, denotes the explicit function induced by the stationarity condition $1-\rho_\alpha^0 = \rho_\alpha^{P/2}$. It is the boundary at which stationarity requires α players to switch from one action to another. By Proposition 7, any strategy at which α players are choosing either 0 or 1 everywhere is a best reply for all α players. But the following strategy, in which α players switch at the switching boundary,

$$A_\alpha^{\text{switch}}(\rho_\alpha, \rho_\beta) = \begin{cases} 1 & \rho_\alpha \geq \rho_\alpha^{\text{switch}}(\rho_\beta) \\ 0 & \rho_\alpha < \rho_\alpha^{\text{switch}}(\rho_\beta), \end{cases} \tag{31}$$

is special, since by construction, it is the only (up to zero probability events) strategy consistent with all stationary beliefs. Combining this, together with Lemma 7 and Proposition 9, yields the following characterization of stationary, geometrically symmetric equilibria:

Proposition 10. *$\langle A_\alpha, A_\beta, \mu \rangle$ is a stationary, geometrically symmetric equilibrium if and only if*

$$A_\alpha(\rho_\alpha, \rho_\beta) \in \{A_\alpha^{switch}(\rho_\alpha, \rho_\beta), A_\beta(\rho_\beta)\}$$

almost surely, and A_β satisfies (26).

Armed with Proposition 10, we now analyze some comparative statics of this equilibrium. Of interest is periodic behavior in the game, and the main questions I ask are the following: What affects the period length? What affects the payoffs to players? And, what affects the switching strategy for α players, A_α^{switch}?

5.1 Comparative statics

5.1.1 Period length

As illustrated by Figure 3, there are many possible period lengths in equilibrium, depending on the initial choice of ρ_β^0. The following result characterizes the set of possible period lengths:

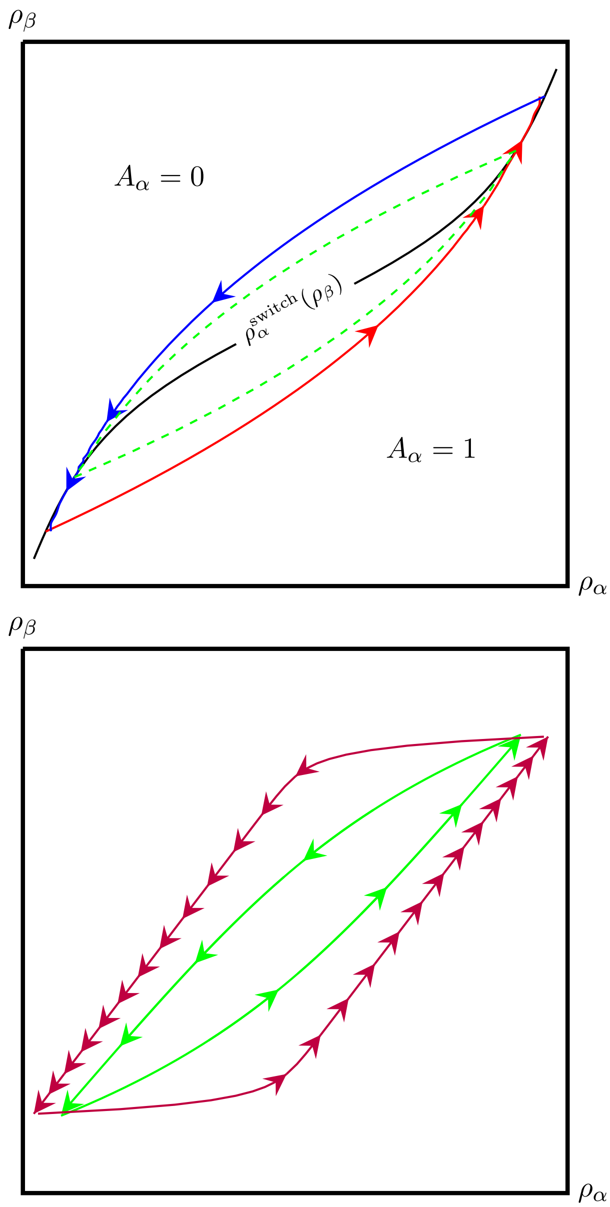

Figure 3: Left: Two possible outcome paths under the matching specification of utility. The solid outer outcome path has equitemporally spaced arrows to indicate speed. The inner dashed curve is another outcome path. When the invariant measure μ is ergodic if and only if it has support only on one outcome path, and the convex hull of all ergodic measures is the set of all nvariant measures. $\rho_\alpha^{\text{switch}}(\rho_\beta)$ represents the set of points at which α players switch in stationary equilibria, derived from the stationarity condition $1-\rho_\alpha^0 = \rho_\alpha^{P/2}$, hence, the indicated strategy at which α players choose 1 when $(\rho_\alpha, \rho_\beta) \geq (\rho_\alpha^{\text{switch}}(\rho_\beta), \rho_\beta)$ and 0 otherwise is the only geometrically symmetric α strategy profile consistent with all geometrically symmetric invariant measures (up to zero-probability events on the boundary of $\rho_\alpha^{\text{switch}}$.) Right: Two outcome paths, with $\lambda_\alpha = \lambda_\beta = M_\alpha = M_\beta = \frac{1}{2}$, but $r_\beta = 1$ in one (light, green) and $r_\beta = 1/4$ in the other (dark, purple). Arrows are equitemporally spaced to illustrate that when $r_\beta = 1/4$, the state path moves more slowly in both dimensions.

Proposition 11. *In every stationary equilibrium under the matching specification of utility, the period satisfies*

$$P \leq \left(\frac{2}{r_\beta \lambda_\alpha}\right)\left(1+\frac{M_\alpha}{M_\beta}(1+\lambda_\beta-\lambda_\alpha)\right)\log\left(1+\frac{M_\beta}{M_\alpha}\right). \tag{32}$$

Furthermore, for any P satisfying (32), there is a stationary equilibrium which induces an outcome path with that period length.

The proof follows from algebraic computation and is relegated to the appendix. Proposition 11 has the following implications for comparative statics of the period length:

Corollary 1. *The maximum possible period length in symmetric stationary equilibria, denoted P^*, satisfies*

1. $\dfrac{\partial P^*}{\partial r_\beta} < 0$, $\dfrac{\partial P^*}{\partial r_\alpha} = 0$,
2. $\dfrac{\partial P^*}{\partial \lambda_\alpha} < 0$, $\dfrac{\partial P^*}{\partial \lambda_\beta} > 0$,

Furthermore, there is a cutoff $\underline{r}_\beta$ such that

1. $\dfrac{\partial P^*}{\partial (M_\alpha/M_\beta)} < 0$ *when* $\lambda_\beta < \underline{r}_\beta$ *and*
2. $\dfrac{\partial P^*}{\partial (M_\alpha/M_\beta)} > 0$ *when* $\lambda_\beta > \underline{r}_\beta$.

Corollary 1 has the following implications about how long fads last. A broad stylized fact about the last century is that the pace of modern life is perceived to be faster today than it was in 1900. Today, the lifespan of fashions is measured in years, a century ago, they might be measured in decades. On Facebook or Twitter, topics trend and then are forgotten in the span of a week. The model suggests that what drives the speed of fads is the up-to-dateness of information available to β players, not that available to α players, as summarized in the following two corollaries to Proposition 11:

Corollary 1 is a statement about orderings of the set of possible period lengths. We could also consider the following: Fix some value of $\rho_\beta^0 < \frac{1}{2}$, and compare the periodic outcome paths, for different parameters, but for which ρ_β^0 is the smallest value achieved by ρ_β^t. (See Figure 3 for an illustration of two state paths derived in this way.) The same comparative statics hold, see Corollary 2 in Appendix A.

Corollary 1 implies that the channel which drives the lifespan of a fad is not the information available to the α players, but that available to the β players. An increase r_β means the average available to β players places more weight on more recent actions, and an increase in λ_α (which, since $\lambda_\alpha + \lambda_\beta = 1$, is a decrease in λ_β) means the average available to β players places more weight on actions taken by α players. Both reduce the lifespan of a fad, for the reason that the faster β players learn the action α players are coordinating on, the faster α players need to switch. This provides a rational for why fads are percieved to begin more often and end more rapidly today, if in 1900 it took days for upper-classes to learn about a new fashion, but months for lower-classes, the model predicts that fads should have lasted months; today, it may still take days for high-classes, but weeks for lower-classes, and accordingly fads last weeks.

5.1.2 Payoffs

The next result characterizes average payoffs in equilibrium. By average payoffs, I mean the payoff expected by a player *ex-ante*, before entering the game,

$$
\begin{aligned}
V_\alpha &:= \mathrm{E}_\mu[U(A_\alpha(\tilde{\rho}_\alpha, \tilde{\rho}_\beta)), A_\alpha(\tilde{\rho}_\alpha, \tilde{\rho}_\beta), A_\beta(\tilde{\rho}_\beta)] \\
V_\beta &:= \mathrm{E}_\mu[U(A_\beta(\tilde{\rho}_\beta)), A_\alpha(\tilde{\rho}_\alpha, \tilde{\rho}_\beta), A_\beta(\tilde{\rho}_\beta)].
\end{aligned}
$$

Proposition 12. *In every symmetric, stationary equilibrium in which* $A_\alpha(\rho_\alpha, \rho_\beta) \in \{0, 1\}$

and $A_\beta(\rho_\beta) \in (0,1)$ on the outcome path,

$$V_\alpha = \frac{M_\alpha}{M_\alpha + \frac{1}{2}M_\beta}$$
$$V_\beta = \frac{1}{2}\frac{M_\alpha}{M_\alpha + \frac{1}{2}M_\beta}.$$

Proof. β's are by construction indifferent between actions, and so, a β player who deviates to the strategy 'mix 50/50 between actions' receives the same payoff as one who mixes according to $A_\beta(\rho_\beta)$. Such a player, by chance, chooses the α-type player's actions half the time, and so

$$V_\beta = \frac{1}{2}V_\alpha. \tag{33}$$

When α players are coordinating on $a = 0$ or $a = 1$, their payoff is the probability that a randomly chosen player choosing that action is an α player in equilibrium. Since α players always choose either $a = 1$ or $a = 0$, consistency of beliefs requires that the average (unconditional) payoff of any player in the game be equal to the probability that an average player is an α, that is

$$\frac{M_\alpha}{M_\alpha + M_\beta}V_\alpha + \frac{M_\beta}{M_\alpha + M_\beta}V_\beta = \frac{M_\alpha}{M_\alpha + M_\beta},$$

which, re-arranging, implies

$$\frac{M_\beta}{M_\alpha}V_\beta = 1 - V_\alpha. \tag{34}$$

Combining (33) and (34) and solving for V_α, V_β yields the desired result. □

Proposition 12 implies that the payoffs to the different player types are independent of all model parameters except for the relative mass of α and β players in the population. In particular, it is independent of the timeliness of the β-type player's information, r_β. Intuitively, this occurs because improving the quality of β-type player's information has two effects: The first is that β players learn faster, and the second, from Corollary 1, is that α players switch more rapidly between actions. Here, the two effects exactly cancel. This result suggests that informing β players about what other players are doing may not have the intended effect. Imagine for example advising a high school student to go to college in order to maximize her earnings, and that going to college increases earnings only because it communicates that the recipient of a diploma is the sort of person who understands that the sort of thing one does to increase earnings is go to college.[11] Then, advising one or two high school students to go to college will benefit them, but making it commonly known that one should attend college to maximize earnings will result in employers seeking some other signal.

5.1.3 Strength of conformity

In Section 4.2 the concept of instrumental preferences for conformity in this model was introduced. We proved that β players were more likely to choose an action the more they'd seen that action chosen in the past. How much more likely they are to choose that action depends on model parameters, I interpret this as the degree of β player's conformity and call it κ, i.e.

$$\kappa := \frac{\partial}{\partial \rho_\beta} A_\beta(\rho_\beta).$$

The following result summarizes the comparative statics of the strength of conformity:

[11] This is valuable to employers if, for example, understanding that one should go to college is indicative of a high-class backround, or signals that an individual understands the other sorts of things that one should do to be successful in a career, or will be a good cultural fit in a particular job.

Proposition 13. *β players are more conformist when there are more α's, and when α's are more visible, i.e.,*

1. $\dfrac{\partial \kappa}{\partial (M_\alpha/M_\beta)} > 0$,
2. $\dfrac{\partial \kappa}{\partial \lambda_\alpha} > 0$.

Proof. The proof follows from straightforward calculation, since from (26),

$$\kappa = \frac{2(M_\alpha/M_\beta) + 1}{2(M_\alpha/M_\beta)\lambda_\beta + 1}.$$

□

Conditional on observing, e.g., a high value of ρ_β, a β player tends to think it is more likely that many β players are choosing $a = 1$, since the periods of time in which lots of both types of players are pooling on a single action last longer than the periods of time when αs are choosing a different action than everyone else (see Figure 3 for an illustration). An increase in the visibility of α players increases the degree to which β's are conformist because it reduces the length of the initial adoption period, in which α's are mainly choosing one action and β's mainly choosing another, hence, when λ_α is larger, conditional on observing a high value of ρ_β, a β type player believes it more likely that everyone, α's and β's alike, is pooling on $a = 1$, and so more β's must in equilibrium choose $a = 1$.

5.1.4 α strategies and anti-conformity

What determines the switching boundary, $\rho_\alpha^{\text{switch}}(\rho_\beta)$? In particular, we are interested in the average slope of the switching boundary for some point $(\rho_\beta, \rho_\alpha^{\text{switch}}(\rho_\beta))$ in the orbit of the state path, as well as an interpretation of its slope. Specifically, consider the following linear strategy α's might play for some fixed slope, $\gamma > 0$ (see Figure 4 for an illustration.)

$$A_\alpha(\rho_\alpha, \rho_\beta) = \begin{cases} 1 & \rho_\alpha - \frac{1}{2} \geq \gamma(\rho_\beta - \frac{1}{2}) \\ 0 & \rho_\alpha - \frac{1}{2} < \gamma(\rho_\beta - \frac{1}{2}), \end{cases} \tag{35}$$

If we fix some value of $\rho_\beta < 1/2$, then the linear strategy in which the switching boundary passes through the points $(\rho_\alpha^{\text{switch}}(\rho_\beta), \rho_\beta)$, $(1 - \rho_\alpha^{\text{switch}}(\rho_\beta), 1 - \rho_\beta)$ supports an equilibrium with a state path passing through those points, by Proposition 10, so that

$$\gamma = \frac{1 - 2\rho_\beta}{1 - 2\rho_\alpha^{\text{switch}}(\rho_\beta)}. \tag{36}$$

How to interpret γ? In Figure 5 is plotted $\rho_\alpha^t, \rho_\beta^t$, for a fixed starting value, as a function of time. The α's strategy while playing $a = 1$, as a function of time, can be divided into three periods. In the first, P^1, the α's coordinate on the action which appears to be a minority action, i.e., $\rho_\alpha^t < \frac{1}{2}$ and $\rho_\beta^t < \frac{1}{2}$. During this initial, very few β players are choosing action $a = 1$, and the payoff to the α players is high. In the second period, P_2, the α's coordinate on an action which appears to be the majority action, since $\rho_\alpha > \frac{1}{2}$, but which is still percieved as the minority action by ρ_β. Since $A_\beta(\frac{1}{2}) = \frac{1}{2}$ in every geometrically symmetric equilibrium, this is also a period of time during which the minority of β players are choosing the action. In the final period, P_3, majorities of both α and β players are coordinating on action $a = 1$. The switching point is determined as the point where the ratio of ρ_β to ρ_α is sufficiently high, that is, $a = 1$ is percieved by α players to have 'played out' as a trend. 'Sufficiently high' is measured by γ. So, γ may be interpreted as the degree to which α players are willing to wait before choosing the minority action. We interpreted the slope of the β's strategy, κ, as measuring the degree to which they mimicked the actions being taken by others, that is, the degree to which β players appeared to be conformist; analogously, we might interpret γ

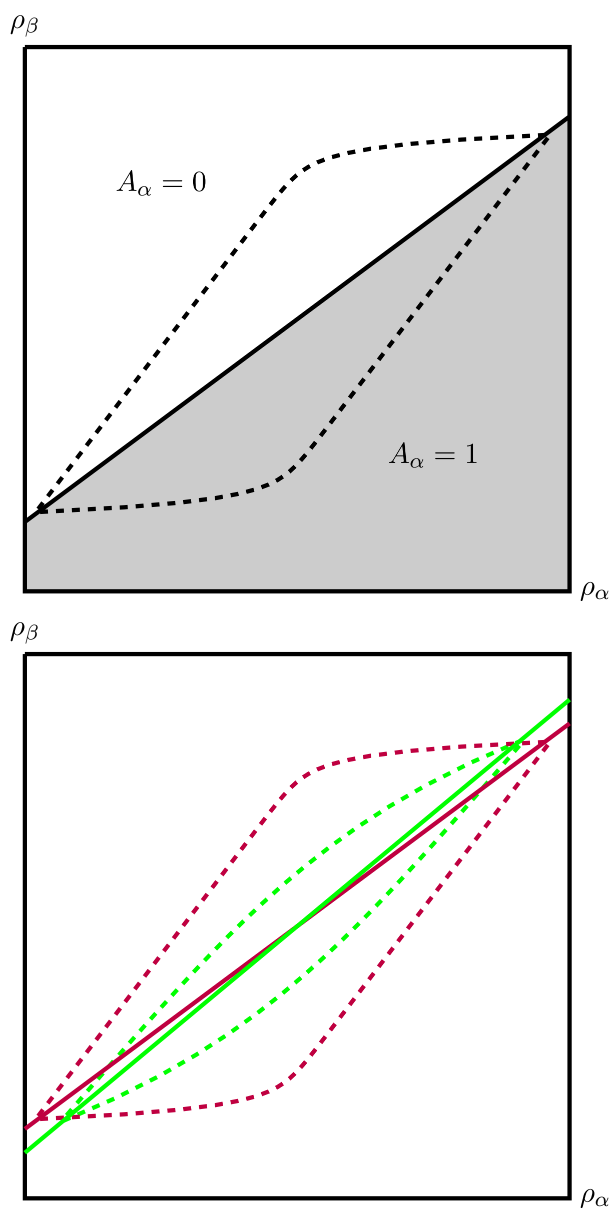

Figure 4: Left: Each outcome path in a geometrically symmetric equilibrium is supported by a corresponding linear strategy. Right: Moving from high r_β (light, green) to low r_β (dark, purple) decreases the slope of the linear switching boundary.

as measuring the degree to which α players anti-conform, a higher value of γ corresponds to α players appearing to be less willing to anti-conform. Under the interpretation of ρ_α, ρ_β as Facebook and Instagram, intuitively, an α player decides to anti-conform and choose $a = 0$ when it appears that the number of people choosing $a = 1$ on Facebook is large and sufficiently close to the number of people choosing $a = 0$ on Instagram, and γ measures the threshold at which the switch occurs.

Proposition 14. *In every equilibrium strategy of the form* (35), $\gamma < 1$. *Furthermore,* γ *satisfying* (36) *is increasing in* r_α, *and decreasing in* r_β.

Proof. The proof follows by substituting the analytical expression for $\rho_\alpha^{\text{switch}}(\rho_\beta)$ derived previously into the expression for γ, (36). It is then a straightforward comparative statics exercise to establish the result. □

Proposition 14 suggests that improving the information technology available to β players, or increasing r_β, increases the degree to which α players are anti-conformist. Intuitively, when β players learn what action players are coordinating on more quickly, α's not only switch more rapidly, as established in Proposition 1, but they also act as if more anti-conformist.

¿¿¿¿¿¿¿ fba381dca49f97033869a13fe7d5b25f8bf9118d

6 Extensions

So far, α players have differed from β players in three distinct ways: First, they observe strictly more information than β players. Second, they have access to a more up-to-date average of the actions of past players, since $r_\alpha > r_\beta$. And third, all players have preferences for matching the actions of α players, and mismatching the actions of β players. In this section, we consider the different cases, summarized in table **??**.

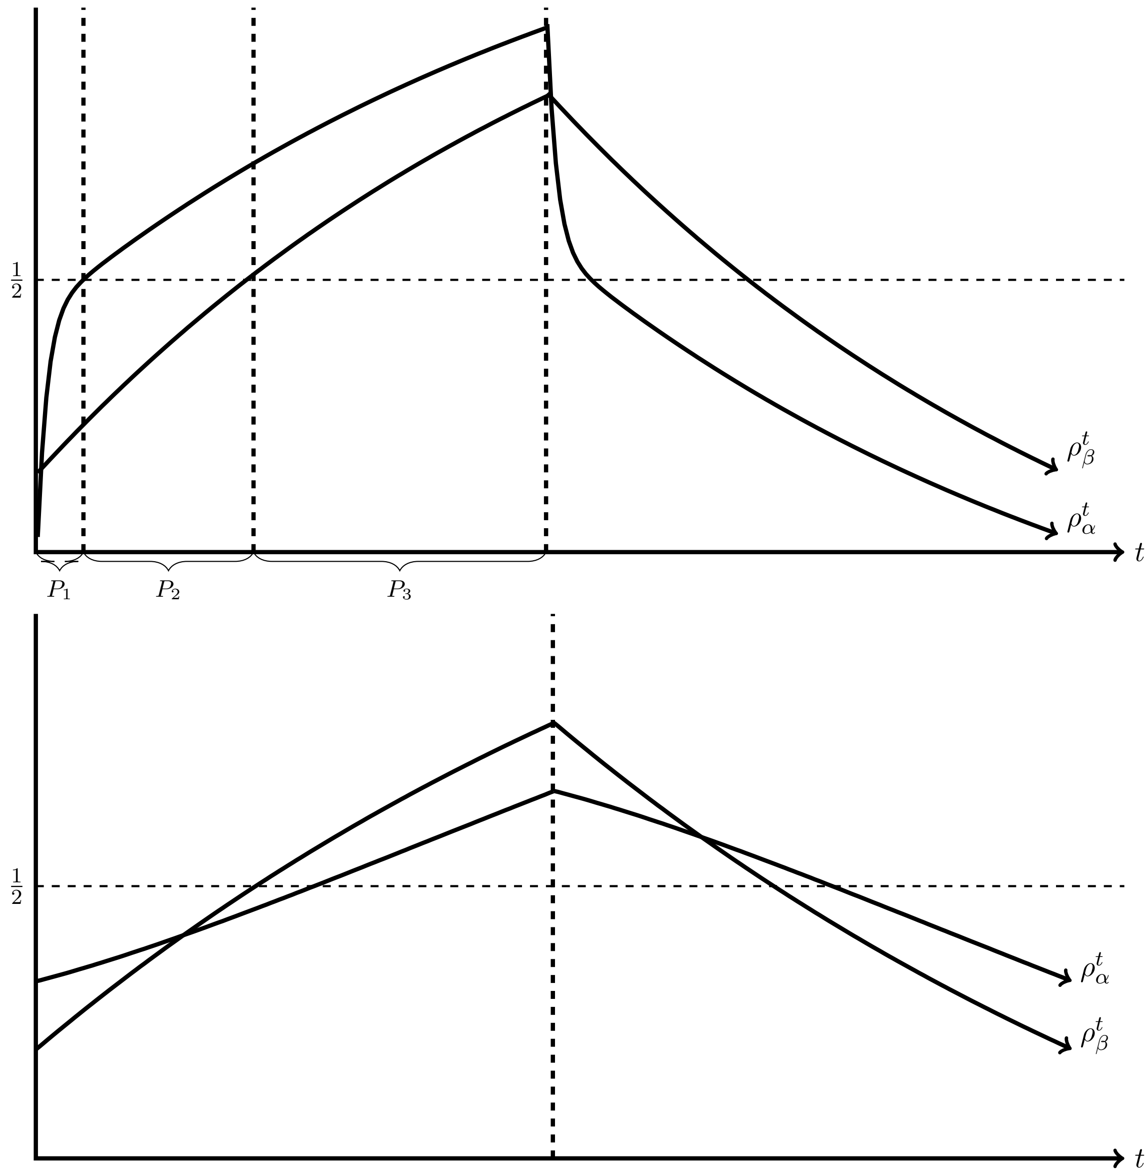

Figure 5: Top: One period of ρ_α, ρ_β as functions of time when $A_\alpha = A_\alpha^{\text{switch}}$. The α's strategy, as a function of time, can be divided into three periods. In the first, P^1, the α's coordinate on the action which appears to be a minority action, $a = 1$. During this initial, very few β players are choosing action $a = 1$, and the payoff to these α players is high. In the second period, P_2, the α's coordinate on an action which appears to be the majority action, according to ρ_α, but which is still percieved as the minority action by ρ_β. Since $A_\beta(\frac{1}{2}) = \frac{1}{2}$ in every geometrically symmetric equilibrium, this is also a period of time during which the minority of β players are choosing the action. In the final period, P_3, majorities of both α and β players are coordinating on action $a = 1$. The switching point is determined as the point where the ratio of ρ_β to ρ_α is sufficiently high, that is, $a = 1$ is percieved by α players to have 'played out' as a trend. The cycle then repeats. **Bottom**: The case in which $r_\alpha < r_\beta$ is a case in which α's anti-conform when they percieve everyone has been taking the *newer* action, and they coordinate on the action which appears relatively *older*.

	$r_\alpha > r_\beta$	$r_\alpha < r_\beta$
α's valued	Main paper	'Hipsters'
β's valued	Only fixed points	Only fixed points

Table 1: The three possible additional cases considered in Section 6.

In this section, I consider the following extensions to the model: First, I consider what happens if β players are the type who are valued, that is, I relax the assumptions on the utility function and allow $U(1, a_\alpha, a_\beta)$ to instead be increasing in a_β and decreasing in a_α. I show that only fixed-point equilibria are possible in this setting. Second, I consider what happens if $r_\alpha < r_\beta$, that is, α players have access to older information about the past actions of players. I show that periodic equilibria may still be supported, but that α player's strategies now may be interpreted as choosing actions which seem relatively older.

6.1 β players are valued

The case in which β players are the ones who are valued corresponds to the following modification to Assumption 1:

Assumption 2. *$U(1, a_\alpha, a_\beta) < U(0, a_\alpha, a_\beta)$ if and only if $a_\alpha > a_\beta$.*

Then the following result is straightforward:

Proposition 15. *Under Assumption 2, in every stationary equilibrium, $A_\alpha(\rho_\alpha, \rho_\beta) = A_\beta(\rho_\beta)$.*

Proof. Say there were some state $(\rho_\alpha, \rho_\beta) \in \text{Supp}(\mu)$ for which $A_\alpha(\rho_\alpha, \rho_\beta) > A_\beta(\rho_\beta)$. (The case in which the inequality is flipped is analogous.) Then by Assumption 1, $U(0, A_\alpha(\rho_\alpha, \rho_\beta), A_\beta(\rho_\beta)) > U(1, A_\alpha(\rho_\alpha, \rho_\beta), A_\beta(\rho_\beta))$, and so $A_\alpha(\rho_\alpha, \rho_\beta) = 0$. But since $A_\beta(\rho_\beta) \geq 0$, we have $A_\beta(\rho_\beta) \geq A_\alpha(\rho_\alpha, \rho_\beta)$, a contradiction. $\square$

Proposition 15 implies that the sorts of dynamics which occurred under Assumption 1 cannot occur if β-type players are the valued players. Players may all pool on an action

($A_\alpha = A_\beta = 0$), or they may randomize independently of what other players are doing. Intuitively, imagine an academic environment, in which appearing to be well-informed about what other academics are wearing is a signal, perhaps, that one pays insufficient attention to one's research. The academics who do pay attention to what others are wearing are always capable of mimicking the dress of those who do not, and so in equilibrium a wide range of dress styles are acceptable and no inferences are drawn from one's clothing. Or, imagine an office environment, where a similar dynamic may be at play, and where all workers coordinate on a uniform dress code, because standing out is frowned upon as a signal that one is insufficiently serious about the job.

6.2 $r_\alpha < r_\beta$

The case in which α players have access to older information than β players about what other players are choosing corresponds to the case in whcih $r_\alpha < r_\beta$. In this case, I assert without proof that Proposition 7 continues to hold,[12] and as before, we may construct the switching boundary, $\rho_\alpha^{\text{switch}}(\rho_\beta)$. Now, however, the α's switch the actions they play (compare to (31)):

$$A_\alpha^{\text{switch}}(\rho_\alpha, \rho_\beta) = \begin{cases} 0 & \rho_\alpha \geq \rho_\alpha^{\text{switch}}(\rho_\beta) \\ 1 & \rho_\alpha < \rho_\alpha^{\text{switch}}(\rho_\beta). \end{cases} \tag{37}$$

As before, we can define the average slope of the switching boundary, γ, as in (36), and the linear strategy, analogous to (35):

$$A_\alpha(\rho_\alpha, \rho_\beta) = \begin{cases} 1 & \rho_\alpha - \frac{1}{2} < \gamma(\rho_\beta - \frac{1}{2}) \\ 0 & \rho_\alpha - \frac{1}{2} \geq \gamma(\rho_\beta - \frac{1}{2}). \end{cases} \tag{38}$$

[12]The proof is the same, with minor modifications.

How to interpret γ now, when $r_\alpha < r_\beta$? When $r_\alpha > r_\beta$, the α players coordinated on actions which seemed recently more popular, until they saw that enough players had been coordinating on that action recently, in which case they switched to choose the action which appeared to be a minority action, but which also appeared to be relatively recently more popular. When $r_\alpha < r_\beta$, the α's have access to relatively older information than β players about the actions of other players, and so they switch to the minority action, they choose the action which appears to be relatively *older* in popularity (see Figure 5), and we have the following modification of Proposition 14 with an analogous proof:

Proposition 16. *In every equilibrium strategy of the form* (35), $\gamma > 1$. *Furthermore,* γ *satisfying* (36) *is decreasing in* r_α, *and increasing in* r_β.

Since in this case α players coordinate on the actions which seem relatively older and less popular, we might interpret this as the case of 'hipsters', while the case in which $r_\alpha > r_\beta$ was the case of 'fashion leaders'.

7 Conclusion

I conclude with a discussion of future avenues for research, as well as miscellaneous topics unsuitable for the main body of the paper.

In the United States, when high school students apply to prestigious colleges, admissions committees generally discriminate between students, who may almost universally have perfect test scores, on the basis of 'holistic' factors. In practice, they tend to look at whether a student engaged in particular extra-curricular activities. Several decades ago, those activities might have included playing the violin, or belonging to a chess club. Today, those activities might include belonging to a lacrosse team, or volunteering at a homeless shelter. Critics of holistic admissions factors charge that they provide a means for prestigious colleges to discriminate against Asian applicants.[13] The model suggests a channel through which this discrimination may take place. An admissions committee, disallowed from directly excluding candidates based on their last name, might instead favor applicants who engage in a traditionally non-Asian activity, such as playing the violin. Over time, Asian families learn about this preference, and soon, whether an applicant plays the violin is non-informative about his ethnicity, and the college admissions committee must switch to favoring playing lacrosse. Importantly, a candidate who plays lacrosse signals only that he understands that the sort of thing one does to get into a a prestigious college is play lacrosse. The model therefore suggests that programs which educate people about the sorts of things one does to get into a good college may be self-defeating. If a college cannot discriminate on the basis of extra-curricular activities, it may either move to explicit discrimination, or stop discriminating.

In 'The Coolhunt'[14], Malcolm Gladwell presents an axiomatic definition of the concept of 'cool', which I re-phrase here:

1. Cool cannot be manufactured, only observed.
2. Cool can only be observed by those who are themselves cool.
3. The act of observing cool causes cool to take flight.

The result of these three rules, writes Gladwell, is a 'closed loop, the hermeneutic circle of coolhunting', in which the cool are always chasing the next trend, and the adoption of that trend is the thing which causes it to be uncool. In this paper, I present a model which I argue formalizes these axioms.

[13]'For Asian Americans, a changing landscape on college admissions', *Los Angeles Times*, February 2015.
[14]'The Coolhunt', *The New Yorker*, March 1997.

Generally speaking, an economist's explanations for the existence of advertising tend to radically differ from everyone else's. One channel through which advertising is commonly perceived to work is by somehow manipulating the perceptions consumers draw of the product. Advertisers themselves see this as a crucial component of their craft.[15] In this story, someone who bought an Apple computer in the 1980s was not necessarily buying a better computer as much as he was buying an identity as a non-conformist artistic type, a perception that Apple encouraged with their famous '1984' and 'Think Different' campaigns. See Akerlof and Kranton (2000) for a formal economic model of identity. This paper suggests an alternate rationale for what it might mean to buy an identity, namely, that buying certain products might credibly signal to others something about oneself. Consumers of Apple products, for example, could credibly signal that they were the sort of people who understood that Apple products are cool. Apple's 'Think Different' campaign hardly mentioned the name of the company, instead featuring pictures of artists and inventors. In fact, a common feature of much modern advertising is that it is deliberately vague about the product being advertised. This is explained within the context of the model, if we imagine that an advertising firm is interested in designing an ad campaign for a product which people are buying to signal that they are part of a high-status group, then there is a tradeoff: Too little advertising, and the low-status group fails to buy the product, too much advertising, and the worth of the product to the high-status group is less. This model provides a framework, grounded in agents with standard preferences, in which to ask questions about the optimal release of information by an advertising firm which controls the parameters of the model.

[15] See, e.g., *Ogilvy on Advertising*, by David Ogilvy: 'These three brands have different images which appeal to different kinds of people. It isn't the whiskey they choose, it's the image. The brand image is 90% of what a distiller has to sell.'

References

Akerlof, G. A. and R. E. Kranton (2000): "Economics and identity," *Quarterly journal of Economics*, 715–753.

Aumann, R. J. (1987): "Correlated equilibrium as an expression of Bayesian rationality," *Econometrica: Journal of the Econometric Society*, 1–18.

Bernheim, B. D. (1994): "A theory of conformity," *Journal of Political Economy*, 102, pp. 841–877.

Bikhchandani, S., D. Hirshleifer, and I. Welch (1992): "A theory of fads, fashion, custom, and cultural change as informational cascades," *Journal of Political Economy*, 100, pp. 992–1026.

Birkhoff, G. D. (1931): "Proof of the ergodic theorem," *Proc. Natl. Acad. Sci. USA*, 17, 656–660.

Bourdieu, P. (1984): *Distinction: A social critique of the judgement of taste*, Harvard Univ Pr.

Caulkins, J. P., R. F. Hartl, P. M. Kort, and G. Feichtinger (2007): "Explaining fashion cycles: Imitators chasing innovators in product space," *Journal of Economic Dynamics and Control*, 31, 1535–1556.

Corneo, G. and O. Jeanne (1999): "Segmented communication and fashionable behavior," *Journal of Economic Behavior & Organization*, 39, 371–385.

Foley, C. A. (1893): "Fashion," *The Economic Journal*, 3, 458–474.

Foster, D. P. and R. V. Vohra (1997): "Calibrated learning and correlated equilibrium," *Games and Economic Behavior*, 21, 40–55.

Frijters, P. (1998): "A model of fashions and status," *Economic Modelling*, 15, 501 – 517.

Glendinning, P. (1994): *Stability, Instability, and Chaos: An Introduction to the Theory of Nonlinear Differential Equations*, Cambridge University Press.

Hájek, O. (1968): *Dynamical systems in the plane*, Academic Press.

Karni, E. and D. Schmeidler (1990): "Fixed preferences and changing tastes," *The American Economic Review*, 80, 262–267.

Kocer, Y. (2010): "Endogenous learning with bounded memory," *Economic Theory Center Working Paper*.

Matsuyama, K. (1992): "Custom versus fashion: Path-dependence and limit cycles in a random matching game," Tech. rep., Hoover Institution, Stanford University.

Milgrom, P. and J. Roberts (1991): "Adaptive and sophisticated learning in normal form games," *Games and economic Behavior*, 3, 82–100.

Pesendorfer, W. (1995): "Design innovation and fashion cycles," *American Economic Review*, 85, 771–92.

Plutarch (187): *Bíoi Parálléloi*, Self published (?).

Sunstein, C. (1995): "Social norms and social rules," The coase lecture, The Law School University of Chicago.

Yoganarasimhan, H. (2012a): "Cloak or flaunt? The fashion dilemma," *Marketing Science*, 31, 74–95.

——— (2012b): "Identifying the presence and cause of fashion cycles in data," Tech. rep., University of Washington Foster School of Business.

The online appendix may be found here.