Changing tastes and imperfect information

Nick Janetos*

April 14, 2016

PRELIMINARY AND INCOMPLETE

Abstract

I analyze a model of fads. A fad is when a choice with no intrinsic value becomes popular, then unpopular. For example, in the 1960s, tailfins on cars were popular, in the 1970s, they were not. In the model, fads are driven through the channel of imperfect information. Some players have better information about the past actions of other players, and are interested in communicating this through their own action choices. I show that in equilibrium, better informed players initially pool on a single action choice. Over time, the rest of the players learn which action the better informed players are pooling on, and start to mimic them. Once a 'tipping point' is reached, the better informed players switch to a different action, and the process repeats.

*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, Steve Matthews, Andy Postlewaite, Alvaro Sandroni, Kris Shaw, and last, and least, Jan Tilly, for comments and suggestions.

1 Introduction

Imagine two people attend a charity event. One inherited a fortune worth billions, the other is a janitor who recently won the lottery and is now worth millions. Assume both are interested in appearing wealthy. Whose clothing looks more expensive? A naive application of signaling theory might lead one to predict the billionaire. After all, he can easily afford a solid gold chain with his name engraved on the front. But the janitor wears the more expensive looking clothing. Why? The answer is that while the janitor can afford to buy the same clothing as the billionaire, he didn't – because he didn't know what clothing to buy. Instead, the janitor bought the sort of clothing he imagined a billionaire would wear, while the billionaire avoided buying the sort of expensive clothing that he thinks a janitor would think a billionaire would wear.

This story applies more broadly: to the choice of smartphone brands, or baby names, or model of cars. The common characteristics of these examples are that all alternatives are functionally equivalent, but that there is often a right choice and a wrong choice, and that the right choice shifts over time. Naming your daughter 'Cherri' is functionally equivalent to naming your daughter 'Elizabeth', but one name has negative connotations that the other lacks. Colloquially, the words 'taste', or 'fashion' are used to refer to this knowledge of the right choices to make – 'good taste' versus 'bad taste'. Economists use the word 'taste' as a synonym for preferences, and assume tastes are fixed over time, but a distinguishing characteristic of taste is that it shifts over time and is correlated across individuals – good taste in 1960 was bad taste in 1970.

In this paper, I ask: Why do tastes shift? Recall the janitor and the billionaire: Perhaps there are two relevant choices for clothing, red hats versus blue hats, but only the billionaire knows the blue hats are in good taste. Over time, the janitor might learn blue hats are in good taste, and start to wear them. But soon, it will be too late. The billionaire has no interest in wearing a style of hats which everyone knows is in good taste. Red hats become good taste, and the cycle repeats. Importantly, this is not a story of differentially costly actions, as in Pesendorfer (1995), or different preferences or action choices available to players, as in Karni and Schmeidler (1990), rather, it is a story of different *information* available to players.

It is not, however, clear that this story can be formally modeled. After all, even if the janitor ignorantly believes that red hats were in good taste, surely he is at least aware of his ignorance, and so could make the correct decision in good taste by simply doing the opposite of whatever he was inclined to do in the first place! I show that this story can be formally modeled in a dynamic continuous time setting. In the model, there is an infinite sequence of two sorts of short-lived

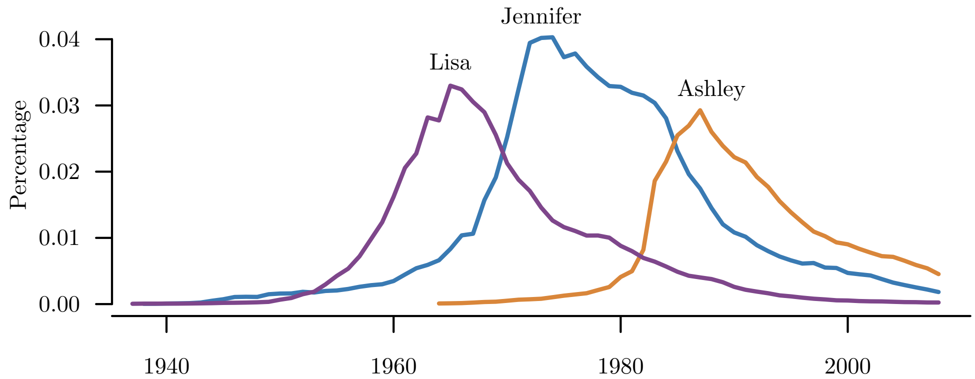

Figure 1: Baby names are an example of changing tastes. *The popularity of the names Lisa, Jennifer, and Ashley, 1940 – 2010, measured as a percentage of total baby names.*

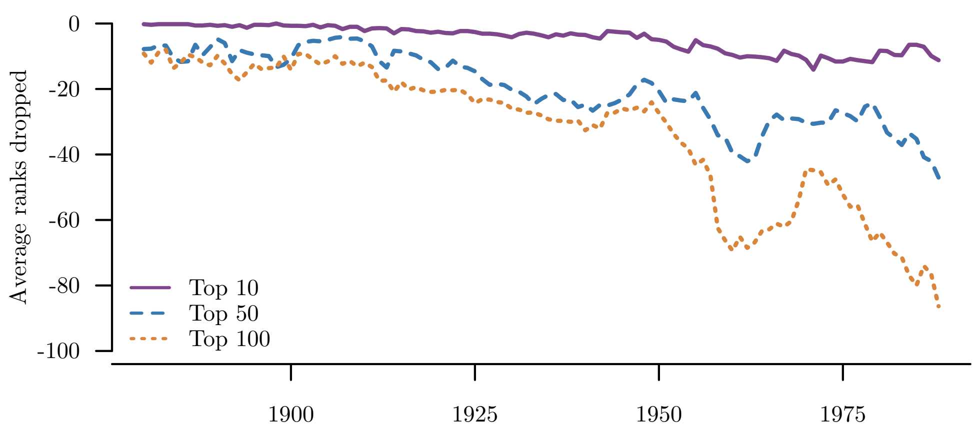

Figure 2: The length of time for which a name is popular has fallen over the past century. *The average number of rankings a name in the top 10, 50, and 100 names could expect to fall in the next decade, by year, 1880 – 1985.*

players. Following Karni and Schmeidler (1990), I refer to them as αs and βs. I assume that all players choose between one of two actions, and that all players share the same, fixed preferences – all want to make the same choice as αs, and a different choice than βs. I show how these preferences can capture, in a reduced form way, instrumental preferences induced by signaling or matching subgames, in which, for example, players match with players who choose the same action, interpreted as something like the choice of restaurant to attend, and want to attend the same restaurant as α players.

In each instant, these players play a stage game in which they have access to some private information. Importantly, the α players have access to better information than the β players. In section 2, I begin by analyzing the stage game independently, and I take this private information as exogenously given. I show that the α players can use their better private information to coordinate on an action, while the βs mix between actions. I then consider the full model in continuous time with an infinite horizon. Here, I allow the information available to players to depend endogenously on actions taken in the past. As in the literature on fictitious play, I assume that players can see some average of actions taken in the past, so that βs can learn over time what actions the αs are coordinating on. Unlike games with fictitious play, however, players are rational Bayesians, and αs have access to an average which places more weight on more recent actions, giving them more up-to-date information on the actions of other players.

I focus on stationary equilibria, in which player beliefs are independent of the time they enter the game. In section 3, I characterize all stationary equilibria using techniques from the literature on dynamical systems. I show that behavior in stationary equilibria takes one of two forms: either the αs and the βs 'pool' on the same action, or a cycle emerges in which the αs and βs cycle back and forth between the two actions. As in the intuitive story of the janitor and the billionaire, the αs 'lead' the βs, meaning that initially in the cycle, all the αs pool on one (arbitrary) action, (which is now the 'good taste' action), while the βs almost all poll on the other action (the 'bad taste' action). Over time, the βs learn which action the αs are pooling on, and gradually, more and more of them start to take that action. At a certain 'tipping point', enough βs are taking the good taste action, that the αs switch to the other action. At the tipping point, tastes switch. What was formally the bad taste action becomes the good taste action, and the cycle repeats.

In section 4, I consider some implications of the model. Although players do not have direct preferences for conformity (conformity is when players are more likely to take an action the more they see other players taking that action), or anti-conformity (which is when players are less likely to take an action the more they see other players taking that action), I show that β players in the model have an instrumental preference for conformity, while α players sometimes conform, and sometimes anti-conform. I consider an augmentation of the model with a matching subgame, in which actions are interpreted as a location choice, players match with other players at the same location they chose, and players derive some utility from matching with α type players. I explicitly compute the equilibrium in this case using the characterization of equilibria from section 3, and consider some comparative statics of the model.

I show that, contrary to intuition, improving the information available to β players does not improve their payoff, rather, the period of the cycles decreases, and switching occurs more often. This has implications for policies which involve education. For example, elite colleges value specific extra-curricular activities, such as playing the violin, when evaluating applications from potential students. If they value these activities because they signal that the candidate is aware of the sorts of things that successful candidates do to gain acceptance to an elite college, then encouraging high school students to play the violin because it improves their chances of acceptance erodes its signaling value, possibly causing colleges to switch to valuing a more obscure extra-curricular activity, like playing lacrosse. I show also that the rate at which switching occurs is driven not by the rate at which α players learn, but the rate at which β players learn, suggesting that an explanation for the stylized observation that in modern times fads are speeding up (see figure 1) is the spread of broadly available technologies such as broadcast television and Facebook.

1.1 Previous literature

The switching dynamics in this game are generated by the information structure of the game. This in in contrast to the previous literature on these sorts of dynamics (Karni and Schmeidler, 1990; Matsuyama, 1992; Bernheim, 1994; Frijters, 1998; Caulkins et al., 2007), 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 norms 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 possible in the static stage game. From this perspective, this paper is a method of refining correlated equilibria in the static game, in which the specific refinement depends on the modeler's specification of the information technology available to players in the dynamic game.

Finally, 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 this paper, I consider some non-stationary dynamics, specifically, I consider the impact of improvements in the information available to βs. I show that, contrary to naive intuition, improving this technology does not increase the payoffs of β players. Rather, the lifespan of fads in the model shortens, with players cycling faster through actions, fast enough so that the β players are left in the same situation they were before.

2 Model

This section is organized into two parts. In the first, I describe the static stage game. Here, some players have access to private information, which they use to coordinate on actions. In the static stage game, I initially take private information as exogenously given. In the dynamic game, the private information available to players will depend on past actions.

2.1 Stage game

A large number of players, some called αs and some called βs, play the following simultaneous-move stage game. For concreteness, assume there is a unit mass of both αs and of βs. First, nature draws the state, $(s_\alpha, s_\beta) \in [0,1]^2$, according to some probability measure μ on the unit square. I denote the realized state by (s_α, s_β), and the random variable by $(\tilde{s}_\alpha, \tilde{s}_\beta)$. Then, players observe some private information: The α players observe (s_α, s_β) directly, while the β players observe only s_β. Each player then chooses a binary action $a \in \{0, 1\}$, and receives some payoff, U, which depends on the complete action profile. I focus on symmetric stationary strategies, denoted $A_\alpha(s_\alpha, s_\beta) \in [0,1]$ and $A_\beta(s_\beta) \in [0,1]$, mapping private information into the probability of choosing $a = 1$. After actions are chosen, each player who chose action a receives a payoff $U(a, a_\alpha, a_\beta)$, where a_α, a_β are the fractions of α and β players choosing 1 in state (s_α, s_β), which, since there are many players, are simply

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

I make the following assumptions on payoffs:

Assumptions

A1 $U(1, a_\alpha, a_\beta)$ is increasing in a_α, and decreasing in a_β, and $U(0, a_\alpha, a_\beta)$ is decreasing in a_α, and increasing in a_β.

A2 $U(1, a, a) = U(0, a, a)\, \forall a \in [0, 1]$.

Intuitively, assumption A1 requires that all players prefer to take actions taken by α players, and all players prefer not to take actions taken by β players. Assumption A2 requires that players have no intrinsic preferences over actions, beyond the preferences induced by the actions of other players, because if the same fraction of αs and βs are choosing both actions, then all players are indifferent between actions.

The following utility specification satisfies these conditions:

Example (Linear utility)**.** *The payoff to a player is the fraction of α players choosing his action, minus the fraction of β players choosing his action, i.e.*

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

Model discussion

A natural question is why the player's preferences are such that they prefer to take the same actions as players with better information about the actions taken by others. One explanation is that those with better information also happen to be the ones who are well-connected, wealthy, and likely to be more successful. Are these the only cases where we should imagine that better information about other's actions is valued? For an alternate explanation, consider the literature on folk theorems in repeated games.[1] Broadly speaking, one conclusion of this literature is that, when players are better informed about the actions of other players, cooperation is easier to sustain. Empirically, researchers have documented case studies showing the importance of such self-supporting behavior, and its dependence upon knowledge of the actions of others (Ostrom, 1990). In fact, one theory for the evolution of human intelligence is that it arose from the necessity of tracking ever-larger amounts of information about the actions of other humans in one's tribe (Flinn et al., 2005).

From this perspective, it is not hard to imagine that being well-informed about the actions of others is fundamentally important for success in any sort of social interaction. When two strangers arrange on Craigslist to exchange a sofa for cash, they are playing an important repeated game, known throughout human history as the 'Will This Stranger Hit Me On The Head And Take My Sofa?' game. When a sofa manufacturer meets with a fabric supplier, and they make a verbal agreement for the fabric supplier to supply high quality fabric to the sofa manufacturer, they are playing another classic repeated game, called the 'Will My Suppliers Rip Me Off With Low Quality Fabric' game.[2,3] At the meeting, both the sofa manufacturer and the fabric supplier have a vested interest in communicating to the other that they are well-informed about the actions people around them are taking, since this in turn communicates that they will be quick to punish deviations, which in turn makes their partnership more valuable.

[1] See Mailath and Samuelson (2006) for a survey.

[2] They are also playing the 'Will This Stranger Hit Me On The Head And Take All My Sofas' game.

[3] The literature on the 'Will My Suppliers Rip Me Off With Low Quality Fabric' game is more commonly known as relational contracting.

2.2 Equilibria of the static stage game

To fix ideas, I derive some Bayesian equilibria of the static stage game. There are many such equilibria, and they depend on our specification of μ, the probability measure on the state space. Later, when I introduce the dynamic model, we will discipline the set of possible probability measures on the state space by allowing μ to evolve endogenously, depending on actions taken by players. For now, let μ be some arbitrary probability measure on the state space, for which conditional probabilities are well-defined.

First, it is easy to verify that the strategy profile which specifies that both sorts of players mix independently 50/50 between 0 and 1, disregarding their private information, is an equilibrium of the game. To see this, recall assumption A2, which states that players receive a payoff of zero regardless of their action choice, as long as equal proportions of the other players are choosing the same action. In this case, exactly 50% of αs and 50% of βs choose the same actions, so that all players are indifferent between all actions. This equilibrium is analogous to a pooling equilibrium in signaling games, in that all player types choose the same action, and no inferences may be drawn about the player type from the action.

On the other hand, it is not an equilibrium for α players to coordinate on action 1, while the βs mix, that is, $A_\alpha(s_\alpha, s_\beta) \equiv 1, A_\beta(s_\beta) \equiv \frac{1}{2}$, since by assumption A1, every β player then strictly prefers to choose $a = 1$:

$$U(1, 1, 0.5) > U(0, 1, 0.5).$$

Are there equilibria in which the α players condition on their additional private information in a non-trivial way? Re-phrased, 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 make the following simple observation:

Claim 1. *Let $\langle A_\alpha, A_\beta \rangle$ be an equilibrium strategy profile. If $A_\alpha(s_\alpha, s_\beta) > A_\beta(s_\beta)$ for some (s_α, s_β), then $A_\alpha(s_\alpha, s_\beta) = 1$, and similarly, if $A_\alpha(s_\alpha, s_\beta) < A_\beta(s_\beta)$, then $A_\alpha(s_\alpha, s_\beta) = 0$.*

Proof. By assumptions A1 and A2, if $A_\alpha(s_\alpha, s_\beta) > A_\beta(s_\beta)$, then

$$\begin{aligned} U(1, A_\alpha(s_\alpha, s_\beta), A_\beta(s_\beta)) &> U(1, A_\alpha(s_\alpha, s_\beta), A_\alpha(s_\alpha, s_\beta)) \\ &= U(0, A_\alpha(s_\alpha, s_\beta), A_\alpha(s_\alpha, s_\beta)) \\ &> U(0, A_\alpha(s_\alpha, s_\beta), A_\beta(s_\beta)). \end{aligned}$$

Since α players see the state, they strictly prefer to choose $a = 1$ when $U(1, A_\alpha(s_\alpha, s_\beta), A_\beta(s_\beta)) > U(0, A_\alpha(s_\alpha, s_\beta), A_\beta(s_\beta))$, and so $A_\alpha(s_\alpha, s_\beta) = 1$ is the only strategy consistent with optimal behavior. Similar reasoning establishes the result for $a = 0$. □

Intuitively, if the α players condition on their private information in a non-trivial way, so that they play a different action profile than βs, then they must be pooling on a single action. A direct corollary of this reasoning is that the β players must always be indifferent between actions:

Claim 2. *Let $\langle A_\alpha, A_\beta \rangle$ be an equilibrium strategy profile. Then the following indifference condition is satisfied everywhere:*

$$\frac{P(A_\alpha(\tilde{s}_\alpha, \tilde{s}_\beta) = 0 \mid s_\beta)}{P(A_\alpha(\tilde{s}_\alpha, \tilde{s}_\beta) = 1 \mid s_\beta)} = \frac{U(1, 1, A_\beta(s_\beta)) - U(0, 1, A_\beta(s_\beta))}{U(0, 0, A_\beta(s_\beta)) - U(0, 1, A_\beta(s_\beta))}, \tag{2}$$

Proof. Fix some s_β. If $A_\beta(s_\beta) \in (0, 1)$, then (2) is simply the necessary indifference condition for mixing to be a best reply. To see this, note that the payoff-relevant uncertainty for a β player arises from his lack of knowledge about the actions of α players, more specifically, his lack of knowledge about whether α players are pooling on $a = 1$, or pooling on $a = 0$, or mixing in such a way so that $A_\alpha(s_\alpha, s_\beta) = A_\beta(s_\beta)$. His expected payoff from choosing $a = 1$ is

$$\begin{aligned}
&\mathrm{E}[U(1, A_\alpha(\tilde{s}_\alpha, \tilde{s}_\beta), A_\beta(\tilde{s}_\beta)) \mid s_\beta] \\
&= \mathrm{P}[A_\alpha(\tilde{s}_\alpha, \tilde{s}_\beta) = 1 \mid s_\beta] \times U(1, 1, A_\beta(s_\beta)) \\
&+ \mathrm{P}[A_\alpha(\tilde{s}_\alpha, \tilde{s}_\beta) = 0 \mid s_\beta] \times U(1, 0, A_\beta(s_\beta)) \\
&+ \mathrm{P}[A_\alpha(\tilde{s}_\alpha, \tilde{s}_\beta) = A_\beta(\tilde{s}_\beta) \mid s_\beta] \times 0,
\end{aligned}$$

and a similar expression holds for $a = 0$. Equating the expected payoff from $a = 1$ to $a = 0$ and re-writing yields (2).

If, on the other hand, $A_\beta(s_\beta) = 1$, so that β players pool on $a = 1$, then it cannot be that with positive probability, α players choose $a = 0$, that is, $\mathrm{P}(A_\alpha(\tilde{s}_\alpha, \tilde{s}_\beta) = 0 \mid \tilde{s}_\beta = s_\beta) = 0$. This follows because, by claim 1, α players are either pooling on $a = 1$ or $a = 0$ in the event that β players are pooling on $a = 1$. In the event that α players are pooling on $a = 0$, then by assumption A2, β players are indifferent between actions, and in the event that α players are pooling on $a = 1$, then by assumption A1, β players strictly prefer to choose $a = 1$. Hence, if there is a nonzero probability that α players are choosing $a = 1$, then β players strictly prefer $a = 1$, a contradiction.

Similar reasoning establishes that if $A_\beta(s_\beta) = 0$ for some s_β, it must be that $A_\alpha(s_\alpha, s_\beta) = 0$, for all s_α, and therefore, by assumption A2, β players are indifferent between actions when pooling on 0 or 1. □

The above reasoning establishes the following: For some fixed μ, if we can partition the state space into a region on which the αs choose the action $a = 1$ and a region on which they choose $a = 0$, and find a function $A_\beta(s_\beta)$ which solves (2) for all values s_β realized under μ, then we have an equilibrium in which the α players can coordinate in a non-trivial way using their private information. For example:

Example. *Say* $\mu = \frac{1}{2} \circ (0,0) + \frac{1}{2} \circ (1,0)$, *and say* $A_\alpha(s_\alpha, s_\beta)$ *is*

$$A(s_\alpha, s_\beta) = \begin{cases} 1 & s_\alpha \geq \frac{1}{2} \\ 0 & s_\alpha < \frac{1}{2}. \end{cases}$$

Equation (2) *then implies, in the linear utility case, that* $A_\beta(0) = \frac{1}{2}$.[4]

This equilibrium is one in which the αs have access to some hidden 'sunspot', which allows them to coordinate on an action. For example, imagine that s_β represents cable television, while s_α represents broadcast television, and people who can afford to watch cable television would like to choose the same action 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. But in reality, information disseminates. Eventually, the βs may learn which action the αs are taking. In the next section, I augment the model so that the set of possible μ is endogenously determined, in such a way so that over time β players can eventually learn which action α players are coordinating on. Can αs still use their private information to coordinate? Perhaps surprisingly, they can, and furthermore, any payment achievable in the static game for any information structure is achievable in equilibrium the dynamic game.

2.3 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 a square, $[0,1]^2$, but elements of the state space are now denoted $(s_\alpha(t), s_\beta(t))$, to represent the dependence on time. I maintain the assumption that both αs and βs observe the value of $s_\beta(t)$, but that only αs observe $s_\alpha(t)$. Again, imagine that s_β represents a source of information available to everyone in the game, such broadcast television, while s_α represents a source of information available only to a subset of the population, such as cable television.

[4]Values of $s_\beta \neq 0$ do not occur with positive probability, the specification of the β players strategy for $s_\beta \neq 0$ depends on the specification of off-equilibrium path beliefs, which I omit here for simplicity.

A strategy profile is $\langle A_\alpha, A_\beta \rangle$, functions

$$A_\alpha : \mathbb{R}_+ \times [0,1] \times [0,1] \to [0,1], \text{ denoted } A_\alpha(t, s_\alpha, s_\beta) \text{ and}$$
$$A_\beta : \mathbb{R}_+ \times [0,1] \to [0,1], \text{ denoted } A_\beta(t, s_\beta),$$

mapping private information into the probability of taking action $a = 1$. I focus on stationary strategy profiles, independent of time, and I denote them by $A_\alpha(s_\alpha, s_\beta), A_\beta(s_\beta)$.

Fix a strategy profile. The game proceeds as follows. First, an initial condition is drawn from $\mu \in \Delta[0,1]^2$. I denote the random variable by $(\tilde{s}^0_\alpha, \tilde{s}^0_\beta)$, and (s^0_α, s^0_β) is the realized initial condition. The game *outcome path* is action paths $a_\alpha(t), a_\beta(t)$ and signal paths $s_\alpha(t), s_\beta(t)$, for $t \in [0, \infty)$, such that, $\forall t$, except perhaps on a (Lebesque) measure zero set,

$$a_\alpha(t) = A_\alpha(s_\alpha(t), s_\beta(t)) \tag{3}$$
$$a_\beta(t) = A_\beta(s_\beta(t)) \tag{4}$$
$$\dot{s}_\alpha(t) = f_\alpha(s_\alpha(t), a_\alpha(t), a_\beta(t)) \tag{5}$$
$$\dot{s}_\beta(t) = f_\beta(s_\beta(t), a_\alpha(t), a_\beta(t)), \tag{6}$$

where f_α, f_β are continuous functions representing the laws of motion for s_α and s_β. I make the following assumptions on f_α, and analogous assumptions for f_β:

Assumption

A3 Let $a_\alpha > a'_\alpha$ and $a_\beta > a'_\beta$. Then

$$\big(f_\alpha(s_\alpha, a'_\alpha, a_\beta) - f_\alpha(s_\alpha, a_\alpha, a_\beta)\big) \times (a_\alpha - s_\alpha) > 0$$
$$\big(f_\alpha(s_\alpha, a_\alpha, a_\beta) - f_\alpha(s_\alpha, a'_\alpha, a_\beta)\big) \times (a'_\alpha - s_\alpha) > 0$$
$$\big(f_\alpha(s_\alpha, a_\alpha, a'_\beta) - f_\alpha(s_\alpha, a_\alpha, a_\beta)\big) \times (a_\beta - s_\alpha) > 0$$
$$\big(f_\alpha(s_\alpha, a_\alpha, a_\beta) - f_\alpha(s_\alpha, a_\alpha, a'_\beta)\big) \times (a'_\beta - s_\alpha) > 0.$$

Assumption A3 requires that the evolution of the state variable is linked to player's action's choice in the straightforward way. If more α players choose action $a = 1$, then the state adjusts toward 1, if more β players choose $a = 1$, then the state adjusts toward $a = 1$. For example, it may be that

$$\dot{s}_\alpha(t) = (r_\alpha a_\alpha(t) + r_\beta a_\beta(t) - s_\alpha(t)) m_\alpha \tag{7}$$
$$\dot{s}_\beta(t) = (r_\alpha a_\alpha(t) + r_\beta a_\beta(t) - s_\beta(t)) m_\beta, \tag{8}$$

where $r_\alpha, r_\beta, m_\alpha, m_\beta$ are positive parameters of the model satisfying $r_\alpha + r_\beta =$

$1, m_\alpha > m_\beta$.

Outcome paths are neither guaranteed to exist, nor to be unique. In the case where an outcome path does not exist, payoffs are assumed to be $-\infty$, in the case where the outcome path is not unique, I select one consistent with a discrete time approximation to the model. The details are unimportant for analysis of the model and relegated to appendix B. For the remainder of the body of this paper, all strategies are implicitly taken to be measurable functions for which the outcome path exists and is unique.

Equations (3) and (4) state that the action paths a_α and a_β, which represent the fractions of αs and βs who took action 1 at time t, are consistent with the strategies of players. Equations (5) and (5) govern the evolution of s_α and s_β. We need to specify some rule to govern their evolution; ideally, such a rule would allow us to model the concept that the state adjusts to reflect actions taken in the past, but imperfectly. The specific choice of a rule may reflect the interpretation of $s_\alpha(t), s_\beta(t)$.

For example, one such interpretation is that s_α, s_β are the fractions of players observed to be choosing action $a = 1$ on either of two social media platforms, one called α and one called β. The way these platforms work is the following: Every time a player chooses an action, his action choice is photographed with some probability and posted on both social media platforms. It remains visible on the platform for some time, but photographs are removed from the platforms exogenously with flow rates m_α, m_β respectively. Formally, this story is represented by (**??**). The solution[5] satisfies

$$s_\alpha(t) = e^{-m_\alpha t} s_\alpha(0) + m_\alpha \int_0^t e^{-m_\alpha(\tau - t)}(r_\alpha a_\alpha(\tau) + r_\beta a_\beta(\tau))\, d\tau$$
$$s_\beta(t) = e^{-m_\beta t} s_\beta(0) + m_\beta \int_0^t e^{-m_\beta(\tau - t)}(r_\alpha a_\alpha(\tau) + r_\beta a_\beta(\tau))\, d\tau,$$

that is, $s_\alpha(t), s_\beta(t)$ are exponentially weighted moving averages of actions taken by players in the past, up to the initial conditions, s_α^0, s_β^0.[6] The evolution of the state variable is governed by the parameters $r_\alpha, r_\beta, m_\alpha$, and m_β. The parameters r_α, r_β are *intratemporal* weights, determining the relative chance each player is photographed and appears on the social network, while m_α, m_β are *intertemporal* weights, parameterizing how current the information on each social network is about the actions players are taking.

This specification of s_α and s_β captures the idea that players have some in-

[5] In the sense of Carathéodory.

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

formation about the actions of other players, but that this information is delayed, and does not immediately reflect changes in action choices. The assumption that $m_\alpha > m_\beta$ captures the idea that the αs have access to more up-to-date information than βs, since a higher value for m_α places more weight on more recent actions. An exponentially weighted moving average is, of course, simply one out of many which we could have chosen. We might instead consider a moving window, or, a number which updates only intermittently according to some Poisson process. In appendix C, I show by example that similar results may be obtained under these two alternatives, but in a less tractable way.

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 $U(a, a_\alpha(t), a_\beta(t))$, satisfying assumptions A1, A2. 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, 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. Note that α players see the realization of $\tilde{s}_\alpha(t), \tilde{s}_\beta(t)$, hence, their belief updating process over the value of the state is trivial, whereas β players only observe the value of $\tilde{s}_\beta(t)$, and hence their beliefs are determined in equilibrium by the marginal distribution of $\tilde{s}_\alpha(t)$.

I say a strategy profile A_α, A_β is an *equilibrium* if (note that I omit the expectation operator for α players, since their belief updating is trivial)

$$\text{Supp}(A_\alpha(s_\alpha, s_\beta)) \subset \underset{a\in\{0,1\}}{\arg\max}\ U(a, \tilde{a}_\alpha(t), \tilde{a}_\beta(t)) \tag{9}$$

$$\text{Supp}(A_\beta(s_\beta)) \subset \underset{a\in\{0,1\}}{\arg\max}\ \text{E}_\mu[U(a, \tilde{a}_\alpha(t), \tilde{a}_\beta(t)) \mid \tilde{s}_\beta(t) = s_\beta], \tag{10}$$

and these posterior beliefs are computed using Bayes' rule (where possible).

Fix a strategy profile, A_α, A_β, and choose some initial state (s^0_α, s^0_β). This induces a state path, the value of which at time t I denote $R_{s^0_\alpha, s^0_\beta}(t) = (s_\alpha(t), s_\beta(t))$. The *orbit* of (s^0_α, s^0_β) is $\{R_{s^0_\alpha, s^0_\beta}(t) | \forall t \in \mathbb{R}_+\}$, the set of values of the state variable traced out over time. The state (s^0_α, s^0_β) is called a *fixed point* of the system if its orbit is a single point. It is called a *periodic point* if there exists $P > 0$ such that $R_{s^0_\alpha, s^0_\beta}(t) = R_{s^0_\alpha, s^0_\beta}(t + P)\, \forall t \in \mathbb{R}_+$, and P is called a *period* of the orbit.[7]

These definitions are pointwise definitions, in the sense that they are independent of the distribution over initial conditions. The following definitions are with respect to the distribution over initial conditions, μ. I say a strategy profile $\langle A_\alpha, A_\beta\rangle$, is *stationary with respect to* μ if μ is an invariant probability measure under the dynamic system defined by (5), (6), that is, if $S \in [0, 1]^2$ is some measurable set

[7]The preceding definitions are standard in the literature on dynamical systems, see citation.

under μ, then

$$\mu(S) = \mu\left(R_S(t)\right) \forall t, \tag{11}$$

where $R_S(t)$ should be interpreted as the set of states at time t induced by beginning somewhere in state S, i.e, $R_S(t) = \{R_{s^0}(t) | s^0 \in S\}$. Condition (11) formalizes the intuition that a player entering a stationary game at time $t = 0$ should, before he observes any private information, have the same beliefs over the state variable as a player entering the game at any other time. Note, however, that we do not require that the state variable itself be stationary, only beliefs over the location of the state variable. Condition (11) could be interpreted in many ways. It could be interpreted as a rigorous alternative to assuming that players do not know the time at which they enter the game, and have some sort of improper uniform prior over the time they enter, so that their beliefs do not condition directly on t. Alternatively, we could interpret it as formalizing the intuition that the game begins at time $t = -\infty$. and has been running for such a long time that a player's knowledge of the period in which he enters is irrelevant to the formation of his beliefs. Viewed from this perspective, the assumption that the prior distribution over initial conditions is be invariant with respect to the dynamical system defined by A_α, A_β is not so much a restriction on the model, but rather a regularity condition that allows us to interpret the initial conditions in a meaningful way, within the context of a game that has both a well-defined probability space (ruling out an improper uniform prior) and an initial node (ruling out time beginning at $-\infty$)

Finally, I say that A_α, A_β, μ is a *stationary equilibrium* if it is an equilibrium and A_α, A_β are stationary with respect to μ. An application of a generalization of the Poincaré-Bendixson theorem[8] gives the following result:

Proposition 1. *Fix A_α, A_β, μ, a stationary equilibrium. Then every point in the support of μ is either a fixed point or a periodic point.*

Proposition 1 implies that the support of μ can be partitioned into a set of periodic orbits and fixed orbits. We do not lose much in the analysis of the model if we assume that the support of μ consists of a single periodic orbit, or a single fixed orbit, since it happens that more general distributions can always be 'stitched together' by combining various periodic or fixed orbits. Therefore, for simplicity, results in the remainder of the paper will be stated in terms of the case in which the support of μ consists of a single orbit.

[8]See Hájek (1968) for the theorem and appendix A for details.

3 Equilibria

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

Proposition 2. *Let A_α, A_β be a stationary equilibrium with respect to μ, and say the support of μ is a single orbit. Then either*

1. *The support of μ is a fixed orbit, (s_α^0, s_β^0), and*

$$A_\alpha(s_\alpha^0, s_\beta^0) = A_\beta(s_\beta^0), \tag{12}$$

or,

2. *with probability 1, $A_\beta(\tilde{s}_\beta^0)$ solves*

$$\left|\frac{f_\beta(s_\beta, 1, A_\beta(s_\beta))}{f_\beta(s_\beta, 0, A_\beta(s_\beta))}\right| = \frac{U(1,1,A_\beta(s_\beta)) - U(0,1,A_\beta(s_\beta))}{U(0,0,A_\beta(s_\beta)) - U(1,0,A_\beta(s_\beta))}, \tag{13}$$

and $A_\alpha(\tilde{s}_\alpha, \tilde{s}_\beta) \in \{0,1\}$.

Condition (12), $A_\alpha(s_\alpha^0, s_\beta^0) = A_\beta(s_\beta^0)$, corresponds to a fixed orbit at (s_α^0, s_β^0) (see point A in figure). Note that when $A_\alpha(s_\alpha, s_\beta) = A_\beta(s_\beta)$, by assumption A3, the state variable is fixed at $(s_\alpha, s_\beta) = (A_\alpha(s_\alpha, s_\beta), A_\beta(s_\beta))$. Condition (13) corresponds to a periodic orbit. It is an indifference condition, and is the dynamic equivalent of the indifference condition in equilibria of the stage game, (2). The difference is, while before the left hand side was the ratio of two conditional probabilities derived from an arbitrarily chosen probability measure, now, they are explicitly determined jointly by the specification of the information technology, (5), (6), and strategic behavior of the players in the game, A_α, A_β.

Conditions (12) and (13) are also sufficient for equilibria – in the following sense:

Proposition 3. *Let $\langle A_\alpha, A_\beta \rangle$ be a strategy profile. If for all $s_\alpha, s_\beta \in (0,1)^2$, it is the case that either*

1. *$A_\alpha(s_\alpha, s_\beta) = A_\beta(s_\beta)$, or,*

2. *A_β solves (13) and $A_\alpha(s_\alpha, s_\beta) \in \{0,1\}$, then*

there exists an invariant probability measure μ on $(0,1)^2$ such that A_α, A_β is a stationary equilibrium with respect to μ.

To see the intuition behind condition (13), consider the case in which f_α, f_β are given by (7), (8), and utility is linear, (1). In this case, (13) becomes (note that m_β

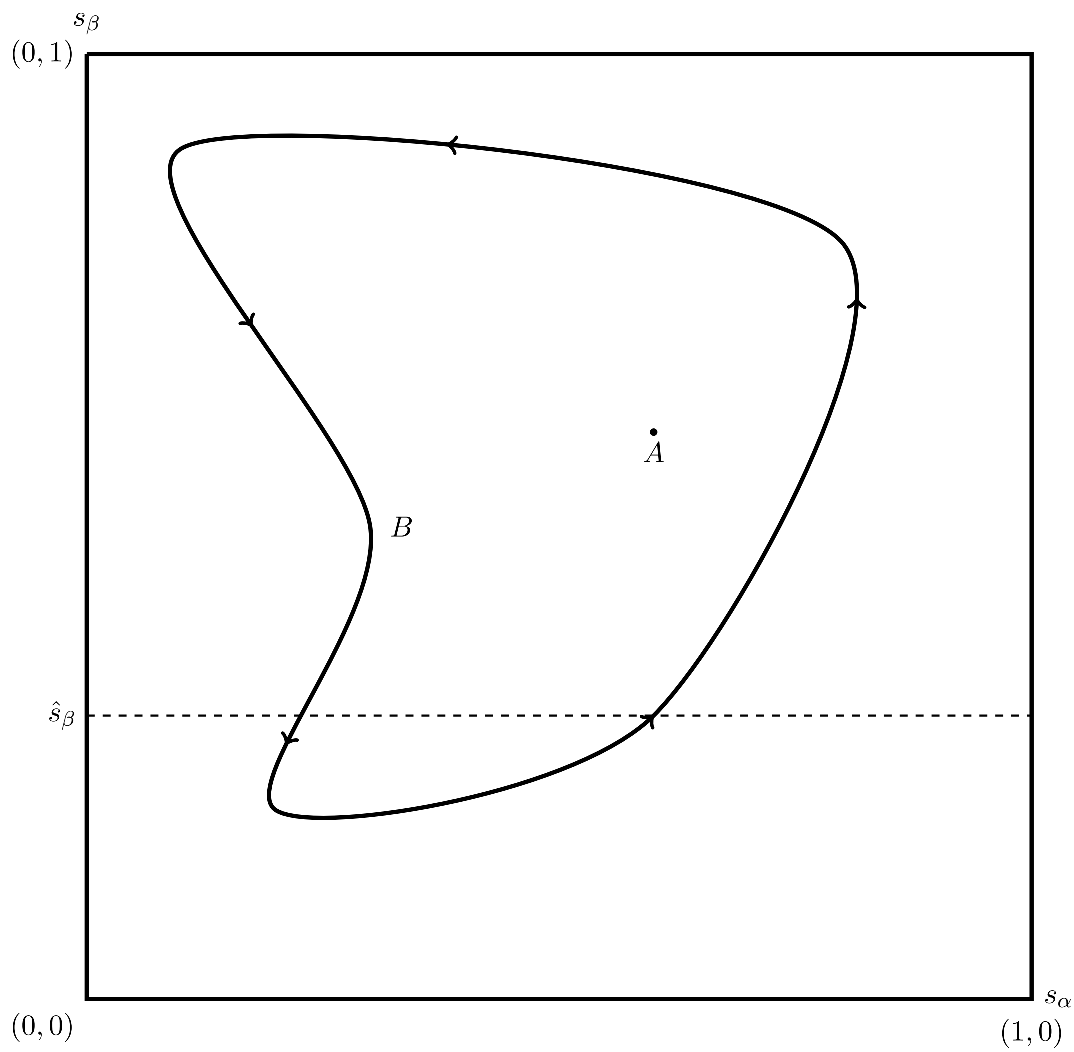

Figure 3: The two sorts of orbits possible in a stationary equilibrium, A, a fixed orbit, and B, a periodic orbit.

drops out):

$$\frac{r_\alpha + r_\beta A_\beta(s_\beta) - s_\beta}{r_\beta A_\beta(s_\beta) - s_\beta} = -\frac{1 - A_\beta(s_\beta)}{A_\beta(s_\beta)}, \tag{14}$$

and equation (14) has a unique solution,

$$A_\beta(s_\beta) = s_\beta. \tag{15}$$

3.1 Induced preferences for conformity

An interpretation of equation (15) 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 s_β, and mimicking it. In fact, this induced preference for conformity is a feature of the general model:

Proposition 4 (Instrumental preferences for conformity)**.** *Say A_α, A_β is a stationary equilibrium with respect to μ. Then with probability 1, if s_β, s'_β are two draws from $\tilde{s}_\beta$ and $s'_\beta \geq (>) s_\beta$, $A_\beta(s'_\beta) \geq (>) A_\beta(s_\beta)$.*

Proof. From proposition 2, we have that either the αs play the same action as the βs, or mix according to (13).

Consider the case in which (13) holds at both s_β and s'_β, and say $s'_\beta > s_\beta$ but $A_\beta(s'_\beta) < A_\beta(s_\beta)$. Then, by assumptions A1 and A2, we have

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

and so

$$\frac{U(1,1,A_\beta(s_\beta)) - U(0,1,A_\beta(s_\beta))}{U(0,0,A_\beta(s_\beta)) - U(1,0,A_\beta(s_\beta))} < \frac{U(1,1,A_\beta(s'_\beta)) - U(0,1,A_\beta(s'_\beta))}{U(0,0,A_\beta(s'_\beta)) - U(1,0,A_\beta(s'_\beta))}. \tag{16}$$

On the other side, from assumptions A3, we have

$$\begin{aligned}
\left|\frac{f_\beta(s_\beta,1,A_\beta(s_\beta))}{f_\beta(s_\beta,0,A_\beta(s_\beta))}\right| &> \left|\frac{f_\beta(s'_\beta,1,A_\beta(s_\beta))}{f_\beta(s'_\beta,0,A_\beta(s_\beta))}\right|\\
&> \left|\frac{f_\beta(s'_\beta,1,A_\beta(s'_\beta))}{f_\beta(s'_\beta,0,A_\beta(s'_\beta))}\right|
\end{aligned}$$

which together with (16) contradicts (13).

This establishes the result in the case that the state path is a periodic orbit, the case in which the state path is a fixed point is trivial and omitted. □

That is, with two independent draws from $\tilde{s}_\beta$, it is almost certain that the β players will be more likely to choose $a = 1$ when the draw is higher. Proposition 4 is interesting 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.

One criticism of this interpretation is that players in the game only 'weakly' have a preference for conformity, in the sense that they are playing mixed strategies and so are indifferent between all strategies. For example, in the game where f_α, f_β are given by (7), (8), and utility is linear, (1), we proved that the equilibrium strategy for β players was $A_\beta(s_\beta) = s_\beta$, but each individual β player is indifferent between choosing this and the 'anti-conformist' strategy $A_\beta(s_\beta) = 1 - s_\beta$, in which they choose the opposite of what they see others choosing.

While it is true that β players are indifferent between all actions, consider the thought experiment of eliciting preferences from players by asking them which action they prefer to choose, and how strong their preference is. Regardless of the fact that players are indifferent between actions, β players will at least strictly prefer to *report* that they strongly prefer the action they chose, since to report a weak preference marks them as a β player. That is, if all players strongly prefer to appear to be α players, then all players have an incentive to pretend to have strong preferences for the action they in fact chose.

To illustrate the connection between conformity and proposition 2, I sketch an intuitive proof of proposition 2. The full details are in the appendix.

In a world where $A_\beta(0.6) = 1$, what inferences should a β player make about whether αs are choosing $a = 1$? It may have been true at some point that αs were all taking the action $a = 1$ – but not for very long, since in a world where βs are conformist, all the βs start choosing $a = 1$ as well. If the β players are 'too' conformist, then each individual β player rationally infers that, if he sees a majority of players taking an action, he should choose the opposite action if he wants to mimic the αs.

What if, instead, the βs are anti-conformist,[9] in the sense that $A_\beta(0.6) = 0$? Via the same reasoning, in such a world each individual β recognizes that, conditional on seeing the majority of players taking a certain action, it is very likely that all the other βs are choosing a different action, and so choosing to conform to the

[9]An anti-conformist is distinct from a non-conformist. A non-conformist ignores the actions of others. An anti-conformist pays careful attention to the actions of others, then does the opposite.

majority action is the best reply. The degree to which βs are conformists is therefore determined by the intersection of these two forces. The full proof of proposition 2 is in the appendix.

4 Applications

In this section, I investigate applications of proposition 2.

4.1 Matching game

Imagine that player's action choices are interpreted as a choice between locations (for example, $a = 0$ is bars on the east side of town, and $a = 1$ is bars 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 -1. The payoff of a player who chooses $a = 1$ is therefore derived using Bayes' rule as

$$\begin{aligned} U(a, a_\alpha, a_\beta) &= \text{P}(\text{Meeting an } \alpha \mid \text{Choosing } a = 1) \\ &= \frac{W_\alpha a_\alpha}{W_\alpha a_\alpha + W_\beta a_\beta}, \end{aligned}$$

where here, W_α, W_β are weights, satisfying $W_\alpha, W_\beta > 0$ and $W_\alpha + W_\beta = 1$, which parameterize the relative population size of αs and βs. We will now proceed to apply proposition 2 to characterize all symmetric (meaning $A_\alpha(s_\alpha, s_\beta) = 1 - A_\alpha(1 - s_\alpha, 1 - s_\beta)$ and $A_\beta(s_\beta) = 1 - A_\beta(1 - s_\beta)$) stationary equilibria.

First, we characterize all stationary equilibria with fixed orbits, corresponding to condition (12) of proposition 2. This is a straightforward task. Let (s_α^0, s_β^0) denote the fixed orbit. Since $A_\alpha(s_\alpha^0, s_\beta^0) = A_\beta(s_\beta^0) := a^*$, and since we are in a fixed orbit, we have $\dot{s}_\alpha = \dot{s}_\beta = 0$, and the laws of motion become

$$\begin{aligned} 0 &= m_\alpha(r_\alpha A_\alpha(s_\alpha^0, s_\beta^0) - r_\beta A_\beta(s_\beta^0) - s_\alpha^0) \\ 0 &= m_\beta(r_\alpha A_\alpha(s_\alpha^0, s_\beta^0) - r_\beta A_\beta(s_\beta^0) - s_\beta^0), \end{aligned}$$

which implies that $s_\alpha^0 = s_\beta^0 = r_\alpha A_\alpha(s_\alpha, s_\beta) + r_\beta A_\beta(s_\beta) = a^*$. It is then not hard to see, by assumption A2 on the utility U, that any $a^* \in [0, 1]$ then characterizes the full set of stationary equilibria with respect to a μ which puts probability 1 on $(s_\alpha^0, s_\beta^0) = (a^*, a^*)$.

Now, we characterize all stationary equilibria with periodic orbits, corresponding to condition (13) in proposition 2. To do so, we begin by solving condition (13), which here is

$$\frac{r_\alpha + r_\beta A_\beta(s_\beta) - s_\beta}{r_\beta A_\beta(s_\beta) - s_\beta} = -\frac{W_\alpha/(W_\alpha + W_\beta A_\beta(s_\beta))}{W_\alpha/(W_\alpha + W_\beta(1-a_\beta))},$$

to obtain the functional form $A_\beta(s_\beta)$ must satisfy at all possible s_β in the support of μ, yielding

$$A_\beta(s_\beta) = \frac{s_\beta(1+W_\alpha) - r_\alpha W_\alpha}{1 + W_\alpha(r_\beta - r_\alpha)}. \tag{17}$$

Proposition 2 states that, if A_β satisfies (17), and $A_\alpha \in \{0,1\}$, then there is an invariant probability measure μ so that A_α, A_β is a stationary equilibrium with respect to μ. Constructing this equilibrium is therefore a matter of choosing regions in (s_α, s_β) space on which $A_\alpha = 0$ and $A_\alpha = 1$ appropriately, which we now proceed to do.

Let us look for an α strategy which splits (s_α, s_β) space into two regions, with a boundary defined by a function $\varphi(s_\beta)$. (See figure 4.1 for an illustration.)

We search for a symmetric strategy, meaning $\varphi(s_\beta) = 1 - \varphi(1-s_\beta)$. Construction of $\varphi(s_\beta)$ proceeds in the following way:

1. Pick a starting value, $s_\beta^0 < \frac{1}{2}$.
2. The corresponding point on the boundary is $(\varphi(s_\beta^0), s_\beta^0)$.
3. Together with this initial condition, 17, and $A_\alpha = 1$, a differential equation is induced by (5), (6). We solve this differential equation to find the path of $s_\alpha(t), s_\beta(t)$, along the region where $A_\alpha = 1$.
4. Once $s_\alpha(t), s_\beta(t)$ 'hits' the point $(1 - s_\beta^0, 1 - \varphi(1 - s_\beta^0))$, at $t = P/2$, the game enters into the phase where αs coordinate on $a = 0$. Repeating the process, we solve for the outcome path along the region where $A_\alpha = 0$, until $(s_\alpha(P), s_\beta(P)) = (\varphi(s_\beta^0), s_\beta^0)$.

Given s_β^0, we have two unknowns: $\varphi(s_\beta^0)$, and P, the period of the orbit. The solution to the differential equation (5), (6) yields two equations, $(s_\alpha(P/2), s_\beta(P/2)) = (1 - \varphi(1-s_\beta^0), 1-s_\beta^0)$ and $(s_\alpha(P), s_\beta(P)) = (\varphi(s_\beta^0), s_\beta^0)$, which we can solve to find $\varphi(s_\beta^0)$. Let's proceed to do that. (In what follows I omit some algebra.)

First, given our initial condition $(\varphi(s_\beta^0), s_\beta^0)$, (17), and $A_\alpha = 1$, the solution to the differential equation (5) is given by

$$s_\beta(t) = 1 - r_\alpha W_\alpha + e^{-\kappa t}(s_\beta^0 - (1 - r_\alpha W_\alpha)), \tag{18}$$

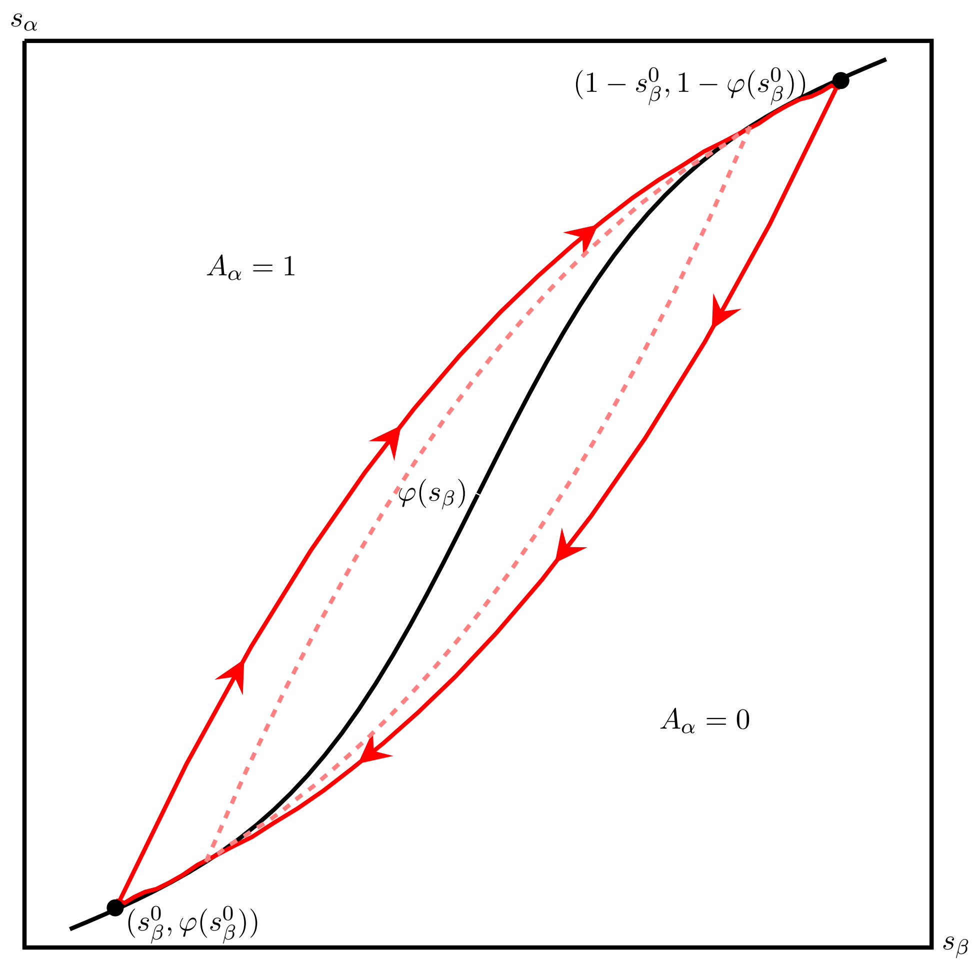

Figure 4: Two possible outcome paths (solid and dashed), under the matching specification of utility.

where

$$\kappa := \frac{m_\beta r_\alpha W_\beta}{2r_\beta W_\alpha + W_\beta}$$

is a growth term. The solution to (6) is considerably more cumbersome, but, e.g., in the simple case when $W_\alpha = W_\beta = r_\alpha = r_\beta$, it reduces to

$$s_\alpha(t) = \frac{12m_\alpha s_\beta^0 e^{-\frac{m_\beta}{4}t} + e^{-m_\alpha t} m_\alpha s_\beta^0 (18\varphi(s_\beta^0) - 21) + 9m_\alpha m_\beta}{4m_\alpha - m_\beta}. \tag{19}$$

To be clear – (18) and (19) are only expressions for the evolution of the state along the region where $A_\alpha = 1$, that is, on the intervals $[0, P/2] \cup [P, 3P/2] \cup \ldots$. We search for values $P, \varphi(s_\beta^0)$ which solve the system of equations

$$s_\beta(P/2) = 1 - s_\beta^0$$
$$s_\alpha(P/2) = 1 - \varphi(s_\beta^0),$$

that is, values so that $s_\alpha(t), s_\beta(t)$ hit the boundary φ at the right time and in the right place. Once we solve for $P, \varphi(s_\beta^0)$, repeating the same process (but with initial values $1 - \varphi(s_\beta^0), 1 - s_\beta^0$) and solving for the path of $s_\alpha(t), s_\beta(t)$ on the region where $A_\alpha = 0$, suffices to find the full equilibrium outcome path. (See figure (4.1).)

The expression for $\varphi(s_\beta^0)$ is cumbersome and unintuitive.

The period length in equilibrium is given by

$$P = \frac{(1 - W_\alpha(r_\alpha - r_\beta)) \log\left(\frac{1-(s_\beta^0 W_\beta + r_\alpha W_\alpha)}{s_\beta^0 W_\beta + r_\beta W_\alpha}\right)}{m_\beta r_\alpha W_\beta}. \tag{20}$$

Period length and player payoffs

How is the period of cycles in the model affected by the parameters of the model?

Proposition 5. *The period length is*

1. *Decreasing in m_β,*
2. *unaffected by m_α,*
3. *decreasing in r_α, and*
4. *increasing in r_β.*

The proof follows from a straightforward comparative statics exercise and is in the appendix. Intuitively, 1 happens because as the information of the β players improves, they take less time to learn which state the α players are coordinating, and so reach the tipping point, $\varphi(s_\beta^0)$, sooner. 3 and 4 occur for similar reasons – as r_α increases, the actions of α players count more prominently in the averages observed by β players, and so are more informative about the action which α players are pooling on.

As m_β increases, the information available to β players becomes more current – so, one might expect the average payoffs received by β players to improve. In fact, this is not the case, which we see by computing expected payoffs to both player types:

Proposition 6. *The average expected payoffs to a β-type player are*

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

and to an α type player are

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

Note that both V_α and V_β are independent of m_α, m_β. The proof is in the appendix, but intuitively, a player's payoff in this parameterization of the game is the probability that the player is an α player, conditional only on the action they are observed to be taking, and the complete history of the game. A β type player always has the option of randomizing 50/50 between actions. If he were to do so, on average, half the time he would pick the same action as α players, and half the time he would pick the wrong action and receive a payoff of zero, meaning that in equilibrium, his expected payoff should satisfy

$$V_\beta \geq \tfrac{1}{2} \times V_\alpha + \tfrac{1}{2} \times 0. \tag{21}$$

In fact, since by construction a player is indifferent between actions at all information sets, he is also indifferent between playing the equilibrium strategy and randomizing 50/50, so that (21) is an equality, $V_\beta = \frac{1}{2}V_\alpha$. The last step follows from the martingale property of Bayesian updating: Since a player's payoff is the probability that the player is an α, conditional only on the action that player took,

5 Conclusion

This model highlights one way in which information can drive the non-trivial dynamics of fashion fads and cycles, which are usually explained via appeals either to irrational behavior, or differences in technology or preferences. I close this paper with a discussion of what I think are some interesting predictions of the model.

In the section on non-stationary dynamics, I considered the impact of a sudden, exogenous improvement in the information available to the β type players, and I showed via example that somewhat paradoxically, better information for β-type players did not improve their payoffs. This formalizes the intuition stated in an axiomatic definition of 'cool' – the 'three rules of cool',[10] which I paraphrase as:

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 is what causes cool to move on.

The model predicts that attempts to identify the next stage in a trend from publicly available data will fail, for the same intuitive reason everyday investors are cautioned not to try to 'play the market' based on publicly available information – if you did manage to identify the next cool thing, it would no longer be the next cool thing.

The model also provides an explanation for the perception that the period in cultural cycles is perceived to be decreasing, if the adoption of technologies such as Facebook is perceived as an improvement in information available to all players. Online platforms such as Twitter are notorious for news cycles which last a few days and then are forgotten. If the purpose of these cycles is not to deliver informational content about fundamentals to users of the platform, but rather give users an opportunity to signal to each other that they're better informed about the world, then shorter periods are a consequence of improving technology.[11]

Improvements in information might also be interpreted within the context of an environment closer to that of Spence (1973)'s canonical model of job market signaling. Every year, high school students engage in extra-curricular activities, in hopes of improving their chances of acceptance to a high-ranking college. Imagine, for example, that students have a choice of taking Japanese as a foreign language, which improves their chances of acceptance to a high-ranking college. A theory of human capital implies that their chances of acceptance are higher because knowing Japanese is an intrinsically valuable skill. The Spence signaling model implies that their chances of acceptance are higher because better qualified students find it easier to learn a challenging language like Japanese, thereby separating themselves from the less qualified students. Under either theory, educating students about the fact that their chances of acceptance are higher if they learn Japanese can be welfare improving: Under the human capital theory, more students learn a valuable skill; under the Spence signaling model, fewer qualified students accidentally fail to take

[10] Malcolm Gladwell, 'The Coolhunters', *The New Yorker*, May 17th, 1997: 'The act of discovering whats cool is what causes cool to move on, which explains the triumphant circularity of coolhunting: because we have coolhunters like DeeDee and Baysie, cool changes more quickly, and because cool changes more quickly, we need coolhunters like DeeDee and Baysie.'

[11] find some article

Japanese because they were perhaps unaware that choice of foreign language is the signaling technology currently in use.

This model provides a third rationale for extra-curricular activities. According to this explanation, returns to extra-curricular activities are driven not by the fact that they improve the human capital of participants, or that they signal the quality of applicants because of heterogeneity in ability, but rather because they signal that the individual participating in the extra-curricular activities is aware of the sorts of things a successful applicant does in order to receive acceptance at a high-ranking college, and this awareness is correlated with other qualities of potential students that colleges value. It may be correlated with wealth, or intelligence. Or it may be correlated with the race of the applicant, as documented by sociologists (Karabel, 2006), who argue that the adoption of 'holistic' acceptance criteria by elite colleges at the beginning of the 20th century was a method of excluding first Jews, and later Asians, from the student body. By using the vague concept of 'character', which often meant participation in activities mostly practiced by white Protestants, elite colleges could plausibly deny admission to otherwise qualified candidates. Over time, the model predicts, minority students learn which activities the colleges favor, at which point, the standards shift. For example,[12] Asian students are underrepresented at elite colleges relative to their academic success. At one time, it was the case that colleges looked for participation in activities such as music, debating, or charity work, in which Asians were underrepresented. Eventually, Asian students began to diversity into these activities; today, elite colleges are relatively more interested in extra-curricular activities such as lacrosse, squash, or rowing, in which Asians are today underrepresented.

References

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.

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.

[12]'The model minority is losing patience', *The Economist*, October 3rd, 2015.

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

Flinn, M. V., D. C. Geary, and C. V. Ward (2005): "Ecological dominance, social competition, and coalitionary arms races: Why humans evolved extraordinary intelligence," *Evolution and Human Behavior*, 26, 10–46.

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.

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

Karabel, J. (2006): *The Chosen*, Mariner Books, 4th ed.

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*.

Mailath, G. and L. Samuelson (2006): *Repeated games and reputations: Long-run relationships*, Oxford University Press.

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.

Ostrom, E. (1990): *Governing the commons: The evolution of institutions for collective action*, Cambridge university press.

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

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

Spence, A. M. (1973): "Job market signaling," *The Quarterly Journal of Economics*, 87, 355–74.

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.