

History, Expectations, and Leadership in the Evolution of Social Norms

DARON ACEMOGLU

Department of Economics, MIT and CIFAR

and

MATTHEW O. JACKSON

Department of Economics, Stanford University, Santa Fe Institute, and CIFAR

First version received January 2013; final version accepted September 2014 (Eds.)

We study the evolution of a social norm of “cooperation” in a dynamic environment. Each agent lives for two periods and interacts with agents from the previous and next generations via a coordination game. Social norms emerge as patterns of behaviour that are stable in part due to agents’ interpretations of private information about the past, influenced by occasional commonly observed past behaviours. For sufficiently backward-looking societies, history completely drives equilibrium play, leading to a social norm of high or low cooperation. In more forward-looking societies, there is a pattern of “reversion” whereby play starting with high (low) cooperation reverts towards lower (higher) cooperation. The impact of history can be countered by occasional “prominent” agents, whose actions are visible by all future agents and who can leverage their greater visibility to influence expectations of future agents and overturn social norms of low cooperation.

Key words: Cooperation, Coordination, Expectations, History, Leadership, Overlapping Generations, Repeated Games, Social Norms

JEL Codes: C72, C73, D7, P16, Z1

1. INTRODUCTION

By coordinating expectations and providing “frames of reference” for interpreting signals and information, *social norms* regulate many areas of human interactions and create self-reinforcing (stable) patterns of behaviour. These norms are not immutable, however. Rather, they emerge and change as a result of social and historical factors, and can also be influenced by ‘leadership’—in particular, the visible actions of prominent individuals, the subject of our analysis. For example, distrust between blacks and whites was a major aspect of social relations in South Africa, cemented by the harsh policies of the apartheid regime. One of the defining challenges for the new South African democracy, holding its first multi-racial elections in 1994, was to break this distrust. Many of the actions of South Africa’s new president and leader of the African National Congress, Nelson Mandela, can be interpreted as using his prominence to switch society to a more cooperating, trusting social norm. Mandela not only gave speeches advocating reconciliation, emphasizing the place of the white minority in the hoped-for “Rainbow Nation”, but even more prominently, he presented the 1995 Rugby World Cup trophy to the South African national team, the Springboks,

wearing their jersey, even though the team had long been a symbol of apartheid. Other symbolic gestures by prominent individuals have had equally long-lasting effects on social norms and expectations. Examples include George Washington's refusal to be considered for a third term in office, which changed the beliefs of many leading contemporaries who viewed the presidency as a form of monarchy (*e.g.* Wood, 2010) and enshrined the limited tenure of U.S. presidents, as well as Mahatma Gandhi's actions emphasizing non-violent resistance and religious tolerance against the background of mounting internal religious tensions and potential violent resistance to British colonialism.

Our main contribution is to develop a model in which the dynamics of behaviour emerge along a single equilibrium with evolving beliefs over time, which allows us to study the role of actions by prominent agents in driving and changing social norms. Our analysis clarifies when history-determined social norms are likely to emerge, and how they can change endogenously in response to prominent agents who coordinate the expectations of different players. We show that social norms can have a natural dynamic whereby greater cooperative behaviour or trust following certain salient events or actions will be endogenously eroded.

Our framework also formalizes the notion that social norms constitute distinct *frames of reference* that coordinate agents' expectations, and shape the interpretation of the information they receive and thus their behaviours. A particular social norm, for example, generalized trust, can persist because the expectation that others will be honest leads agents to interpret ambiguous signals as still being consistent with honest behaviour and thus overcoming occasional transgressions. In contrast, a social norm of distrust would lead to a very different interpretation of the *same* signals and a less trusting pattern of behaviour.

To communicate our main contribution in the clearest fashion, we focus on a coordination game with two actions: “*High*” and “*Low*”. *High* actions can be thought of as more “cooperative”. This base game has two pure-strategy Nash equilibria, and the one involving *High* actions by both players leads to higher payoffs for both players. We consider a society consisting of a sequence of players, each corresponding to a specific “generation”.¹ Each agent's payoff depends on his or her action (which is decided at the beginning of his or her life) and the actions of the previous and the next generation. Agents only observe a noisy signal of the action by the previous generation and so are unsure of the play in the previous period—and this uncertainty is maintained by the occasional presence of agents who find it dominant to play *High* or *Low* behaviour. In addition, a small fraction of agents are *prominent*. Prominent agents are distinguished by the fact that their actions are observed perfectly by all future generations. This formalizes the notion of shared (common) historical events and enables us to investigate conditions under which prominent agents can play a leadership role in changing social norms.

We show that a greatest equilibrium, which involves the highest likelihood of all agents choosing *High* behaviour, always exists as does a least equilibrium. The greatest (as well as the least) equilibrium path exhibits the behaviour at which we have already hinted. First, depending on the shared (common knowledge) history of play by prominent agents, a social norm involving most players choosing *High*, or a different social norm where most players choose *Low*, could emerge. These social norms shape behaviour because they set the *frame of reference*: what agents expect those in the past to have played, and those in the future to play, are governed by the prevailing social norm. So, the past history when coupled with equilibrium behaviour sets a prior belief about the past and future play of others. Because they only receive noisy information about past play, agents interpret the information they receive according to the prevailing social norm as

1. The assumption that there is a single player within each generation is adopted for simplicity and is relaxed in Section 5.

determined by the shared (common) history.² For example, even though the action profile (*High, High*) yields higher payoff, a *Low* social norm may be stable, when agents expect others in the past to have played *Low* (e.g. distrust between blacks and whites in South Africa, even though they would be better off with a more trusting approach to race relations). In particular, the first agent following a prominent *Low* play will know that at least one of the two agents he or she interacts with is playing *Low*, and this is often sufficient to induce him or her to play *Low*. The next player then knows that with high likelihood the previous player has played *Low* (unless he or she was exogenously committed to *High*), and so the social norm of *Low* becomes self-perpetuating. Moreover, highlighting the role of the interplay between *history* and *expectations* in the evolution of social norms, in such an equilibrium even if an agent plays *High*, a significant range of signals will be interpreted as coming from *Low* play by the next generation and will thus be followed by a *Low* response. This naturally discourages *High*, making it more likely for a *Low* social norm to persist. When prominent agents are rare, these social norms can last for a long time.

Secondly, except for the extreme settings we have just discussed in which all endogenous agents follow the action of the last prominent agent, behaviour fluctuates between *High* or *Low* as a function of the signals agents receive from the previous generation. In such situations, society tends to a steady-state distribution of actions. Convergence to this steady state exhibits a pattern that we refer to as *reversion*. Starting with a prominent agent who has chosen to play *High*, the likelihood of *High* play monotonically *decreases* as a function of the time elapsed since the last prominent agent (and likewise for *Low* play starting with a prominent agent who has chosen *Low*). The intuition for this result is as follows: an agent who immediately follows a prominent agent, let us say the period 1 agent, is sure that the previous agent played *High*, and so the period 1 (endogenous) agent will play *High*.³ The period 2 agent then has to sort through signals as it could be that the period 1 agent was exogenous and committed to *Low*. This makes the period 2 (endogenous) agent's decision sensitive to the signal that he or she sees. Then in period 3, an endogenous agent is even more reluctant to play *High*, as now he or she might be following an exogenous player who played *Low* or an endogenous agent who played *Low* because of a very negative signal. This continues to snowball as each subsequent player then becomes more pessimistic about the likelihood that the previous player has played *High* and so plays *High* with a lower probability. Thus, as the distance to the prominent agent grows, each agent is less confident that their previous neighbour has played *High*. Moreover, they also rationally expect that their next period neighbour will interpret the signals generated from their own action as more likely to have come from *Low* play, and this reinforces their incentives to play *Low*.

Thirdly and most importantly, this setting enables us to formally study *leadership-driven changes in social norms*. We show that prominent agents can counter the power of history by exploiting their visibility to change the prevailing social norm from *Low* to *High*. In particular, starting from a social norm involving *Low* play—as long as parameters are not so extreme that all *Low* is locked in—prominent agents can (and will) find it beneficial to switch to *High* and create a new social norm involving *High* play. The greater (in fact, in our baseline model, perfect) visibility of prominent agents means that: (i) they know that the next generation will be able to react to their change of action; and (ii) the prominent action is observed by all future agents who can then also adjust their expectations to the new norm, thus further incentivizing the next generation to play *High*. Both the understanding by all players that others will also have observed the action of the prominent agent (and the feedback effects that this creates) and the anticipation

2. History is summarized by the action of the last prominent agent. The analysis will make it clear that any other *shared understanding*, e.g. a common belief that a specific action was played at a certain point in time, could also play the same role.

3. This is true unless all endogenous non-prominent agents playing *Low* is the only equilibrium.

of the prominent agent that he or she can change the expectations of others are crucial for this type of leadership.

Finally, we also clarify that the leadership role of prominent agents is not reliant on perfect visibility to their partner. An agent who will be prominent to future generations but generates only imperfect signals to the next generation can also play a major role in changing social norms because the expectation that future generations may coordinate with his or her action makes his or her partner more willing to match his or her (noisily observed) behaviour, in turn making future generations even more likely to coordinate with him/her.

Social norms and conventions and their dynamics are the focus of several important literatures. First, the literature on dynamic and repeated games of incomplete information has studied how reputations affect behaviour, and the conditions for the emergence of more cooperative equilibria (see, *e.g.* Mailath and Samuelson (2006), for an excellent overview).⁴ For example, Tirole (1996) develops a model of “collective reputation” in which an individual’s reputation is tied to his or her group’s reputation because his or her past actions are only imperfectly observed.⁵ Tirole demonstrates the possibility of multiple steady states and shows that bad behaviour by a single cohort can have long-lasting effects.⁶ Secondly, the evolutionary game theory literature has studied the dynamics of social norms extensively (see Young, 2010, for a recent survey). In particular, Young (1993) and Kandori *et al.* (1993), and numerous papers building on their work, study evolutionary models where equilibrium behaviour in a coordination game played repeatedly by non-forward-looking agents follows a Markov chain and thus results in switches between patterns of play, which can be shaped by both history and random shocks.⁷ Our contributions relative to this extensive literature are as follows: (i) we consider an overlapping generations model of interactions rather than the random matching set-up on which most of this literature focuses; (ii) we assume that the structure of the game is common knowledge and study forward-looking Bayesian behaviour in which social norms affect the interpretation of signals of past behaviour, while much of the past literature on norms focuses on best response dynamics; (iii) we study the role of prominent agents in anchoring and sometimes changing social norms. In particular, our analysis clarifies how prominent agents (who are more visible to others) affect

4. Repeated games with overlapping generations of players or with asynchronous actions in which, as in our model, expectations of future play matter are also worth noting (*e.g.* Lagunoff and Matsui, 1997; Anderlini *et al.*, 2008).

5. Also related is Tabellini (2008) who, building on Bisin and Verdier (2001), endogenizes preferences in a prisoners’ dilemma game as choices of partially altruistic parents. The induced game that parents play has multiple equilibria, leading to very different stable patterns of behaviour in terms of cooperation supported by different preferences. See also Doepke and Zilibotti (2008) and Galor (2011) for other approaches to endogenous preferences.

See also Jackson and Peck (1991) who show the role of the interpretation of signals, history, and expectations as drivers of price dynamics in an overlapping generations model.

6. Other notable recent examples include Bidner and Francois (2013) who model the interactions between tolerance towards political transgressions, which is itself shaped by social norms, and choices by political leaders; and Belloc and Bowles (2013) who examine the interaction between conventions and institutions.

7. As examples, Azur (2004) models the dynamic process of tipping in an evolutionary context. Argenziano and Gilboa (2010) and Steiner and Stewart (2008) emphasize the role of history as a coordinating device in equilibrium selection, but using an approach in which expectations are formed on the basis of a similarity function applied to past history (thus more similar to the non-forward-looking behaviour in models of evolution and learning in games). Diamond and Fudenberg (1989), Matsuyama (1991), Krugman (1991), and Chamley (1999) discuss the roles of history and expectations in dynamic models with potential multiple steady states and multiple equilibria, and are thus also related.

A more distant cousin is the growing global games literature (*e.g.* Carlsson and van Damme, 1993; Morris and Shin, 1998; Frankel and Pauzner, 2000; and Burdzy *et al.* 2001). However, this literature is not concerned with why groups of individuals or societies in similar economic, social, and political environments end up with different patterns of behaviour and why there are sometimes switches from one pattern of behaviour to another.

not only those with whom they directly interact, but also future generations⁸; and (iv) we also show, under certain parameter values, the emergence of an interesting pattern of mean-reverting dynamics of expectations and behaviour, which also shapes social norms.

There is also a literature that models leadership, though mostly focusing on leadership in organizations (*e.g.* Myerson, 2011; Hermalin, 2012). The notion of leadership in our model, which builds on prominence and observability, is quite different from—and complementary to—the emphasis in that literature.⁹

The rest of the article is organized as follows. The next section discusses several applications that motivate our approach and its assumption. Section 3 presents the model. Section 4 contains our main results. Section 5 clarifies the role of prominence in coordinating expectations and considers several extensions. Section 6 concludes. Appendix A contains the main proofs, whereas Appendix B, which is available online as Supplementary Material, presents additional material, including some motivating examples, and proofs omitted from the article and Appendix A.

2. APPLICATIONS

We begin with a few examples of applications and settings that fit within our model and help motivate the questions and the analysis that follows.

2.1. Cooperation

The canonical application motivating many of our ideas is one of societal cooperation. Suppose that cooperation decisions are taken within the context of a partnership with payoffs:

	Cooperate	Not cooperate
Cooperate	β, β	$-\alpha, 0$
Not cooperate	$0, -\alpha$	$0, 0$

where $\alpha, \beta > 0$. This payoff matrix implies that it is a best response for an individual to cooperate when his or her partner is doing so and this yields the highest payoff to both players. However, not cooperating is also a Nash equilibrium. When cast in the context of a proper dynamic game, non-cooperation can thus emerge as a “social norm”.

The simplest way of placing these interactions into a dynamic setting is by using an overlapping-generations framework, where each individual plays this game with an agent from the previous generation and one from the next generation. We introduce “stickiness” in dynamic behaviour by assuming that each agent chooses a single action, which determines his or her payoffs in both of his or her interactions. We also introduce incomplete information, so that each individual will be acting on the basis of a noisy signal about the action of the previous generation

8. Ellison (1997) infuses a rational player into a society of fictitious players and shows that the rational agent has an incentive to be forward looking in sufficiently small societies. See also Blume (1995) and Matsui (1996) for other approaches to the evolution of play in the presence of rational players. However, we emphasize a calculus of prominent agents in shaping future social norms that is not present in the previous literature.

9. A recent empirical paper by Borowiecki (2012) investigates long-run persistence of preferences for cultural goods, such as classical music, which require coordinated demand/support for their production. Borowiecki finds that the *local* birth of prominent classical composers in the Renaissance is a significant predictor of current provision of cultural goods in Italian provinces today. He finds that a standard deviation (SD) increase of prominent composer births within a province in the renaissance correlates with a 0.4 SD increase in current cultural goods provision. Although one can imagine other explanations for such observations, it is an important finding with respect to the persistent impact of prominent behaviours.

(and realizes that the next generation will observe a noisy signal of his or her action). This allows for play to persist, but also to change over time. Prominent agents will be those whose actions can be seen (without noise or with less noise) by future generations.

2.2. Norms within a profession

A more substantive application is to norms of behaviour within a profession, which we illustrate in the context of a police force, but the same applies to any occupation, including governmental and private sector settings, in which there is a choice of corrupt versus ethical behaviour and these choices across individuals are complements.

Consider a police force where policemen work in teams of two consisting of one older and one younger policeman. Each generation consists of N policemen, and each older policeman is randomly matched with a younger policeman. Each policeman can demand and accept bribes during his activities, but the payoff to this depends on whether his partner is also corrupt. Some individuals are “exogenously committed” to one of the two actions (or in terms of payoffs it is dominant for them to be corrupt or not). For the remaining “endogenous” agents, monetary payoffs from corruption are given by the following matrix:

	Corrupt	Not Corrupt
Corrupt	κ, κ	$\phi\kappa, 0$
Not corrupt	$0, \phi\kappa$	$0, 0$

where $\kappa > 0$ and $0 < \phi < 1$. This captures the fact that a corrupt policeman teamed up with an honest policeman will be less effective in getting bribes. Without any penalties, corruption would be a dominant strategy for endogenous agents.

Police corruption, however, is routinely investigated. Old and currently corrupt police are caught with probability p and incur a penalty x . A policeman who is honest in old age but was corrupt in his youth is caught for past corruption with probability $p' \leq p$, and is subject to the same fine x if caught. The probability of getting caught when young is small.¹⁰

It is then clear that a policeman will never choose to be honest in his youth and corrupt in old age. Moreover, provided that p' is sufficiently close to p , they will also not choose corruption in their youth and honesty in old age, since this reduces the probability of getting caught by only a small amount. This implies that each policeman will choose the same behaviour both in the first and second periods of their lives, and corrupt policemen will be caught with probability p , implying a cost of corruption of px . Splitting the cost of corruption into a per period cost of $px/2$, the expected per-period payoff matrix of policemen can be written as

	Corrupt	Not corrupt
Corrupt	$\kappa - px/2, \kappa - px/2$	$\phi\kappa - px/2, 0$
Not corrupt	$0, \phi\kappa - px/2$	$0, 0$

When $px/(2\phi) > \kappa > px/2$, this payoff matrix is equivalent to the one introduced in the previous subsection with $\beta = \kappa - px/2 > 0$ and $\alpha = -(\phi\kappa - px/2)$, except that corruption in this game corresponds to the cooperative action in the previous subsection. The full information version of

10. We could also allow for an honest policeman whose partner is corrupt to be (falsely) convicted of corruption with some probability $p'' > 0$. This would increase the complementarities in payoffs, without changing the basic structure of payoffs. We analyse the simpler case where honest policemen earn 0 in each period.

this game has two pure-strategy equilibria, one in which both policemen are corrupt and one in which both are honest.

Now let us suppose that policemen decide whether to accept bribes under incomplete information. In particular, they observe some signal about the past behaviour of their first partner when they are young (but know little about their future partners). They also understand that a signal of their current actions will be observed by their future partners and may impact their behaviours. Our analysis will show that if policemen expect others to be corrupt, then they will interpret the signals they receive in this light, and will be more likely to choose corruption unless they receive very strong signals to the contrary.

Occasionally, some generations of policemen or some individuals have greater visibility, and their actions (heroic or villainous) are observed by many generations and would be part of the shared history of the police force. This can occur when a corruption scandal occurs and some corrupt behaviour is publicly reported; or when some particularly good, heroic, or honest behaviour is similarly publicly reported (*e.g.* in the press). Our framework will thus both highlight the importance of this sort of prominence on social norms and show how it can be studied.

2.3. *Conflict and trust between two groups*

A related model can be used to capture the dynamics of conflict and trust between two groups, such as black and white South Africans. We could think of players from each group taking turns (as in [Acemoglu and Wolitzky, 2014](#)). If payoffs are given by the cooperation game above, this can lock in the social norm of conflict and distrust between the two groups. In this light, our discussion of Nelson Mandela's leadership can be viewed as a switch by a prominent agent to cooperation, even when he believed that whites at the time were not cooperating and were expecting blacks not to cooperate.

2.4. *Collective action*

Consider the following model of collective political action. A state is ruled by an autocrat who can be forced to make concessions if citizens protest in an organized manner. Society has an overlapping generations structure as already described in the context of our previous applications, but with each generation populated by n agents. To maximize similarity with our other applications, suppose also that the country consists of n "neighbourhoods", and each agent is assigned to a neighbourhood (see Section 5 for essentially the same model without the neighbourhood structure). Suppose that the concessions the autocrat makes are neighbourhood-specific (*e.g.* reduce repression or make public good investments in that neighbourhood), so directly, each agent only cares about the his or her neighbourhood.

Each agent has a choice: protest or not. We assume that an agent who is the only one who protests within his or her neighbourhood—that is a young agent protesting when the old agent in the same neighbourhood does not, or an old agent doing so when the next generation does not—incurs an expected cost of $\alpha > 0$ (*e.g.* this could be the product of the probability of getting caught times the disutility from the punishment by the autocrat). If two agents protest in the neighbourhood, then they are able to force the hand of the autocrat, and receive a per person per period gain of $\beta > 0$. In addition, an agent who protests when young can be caught and punished when old even if he or she does not protest in the second period of his or her life. If this probability is sufficiently close to the probability of getting caught when protesting then an agent who protests in youth will also do so in old age.

This implies that the payoff matrix for the interaction between an old and young agent can then be represented as

	Protest	Not protest
Protest	β, β	$-\alpha, 0$
Not protest	$0, -\alpha$	$0, 0$

This is identical to the payoff matrices in our other applications.

In addition, individuals are uncertain, but observe informative signals about whether others are protesting: the young receive a signal of past play (either of a randomly chosen agent from the previous generation or from the behaviour of the agent in the same neighbourhood). For instance, the noise in the signal could come from autocrats' natural tendency to make it difficult to observe unrest. This results in a dynamic game of incomplete information as in our baseline model.

Prominence now has a clear counterpart as actions by agents who visibly protest and can serve to coordinate the expectations of future generations. In fact, in Section 5 we will return to a variation on this example (though without the neighbourhood structure) and interpret prominent agents as those that coordinate protests across a number of regions.

2.5. Other applications

The applications presented so far clarify that our framework is based on three features:

- complementarities in actions across individuals;
- some stickiness in an individual's behaviour across time due either to an up-front investment cost, switching costs, or liability for past actions; and
- incomplete information about the actions of others at the time of decision making.

Clearly, there are many additional applications that have these features including, among others: complementary investments (*e.g.* in partnerships or in matches); the leading of an army (with prominent leaders seen as leading an attack); tax avoidance or more general law breaking (as the probability of being caught may be lower when others also breaking the law); and the choice of general social values and morals that are taught to children (whether to be honest or not). Although our baseline model and the applications discussed here are in terms of the coordination game, in Section 5 we also show that much of the analysis and insights extend to applications where the stage game takes the form of a prisoner's dilemma.

3. THE MODEL

We consider an overlapping-generations model where agents live for two periods. We suppose for simplicity that there is a single agent born in each period (generation). This is extended to a setting with more individuals within each generation in Section 5. Each agent's payoffs are determined by his or her interaction with agents from the two neighboring generations (older and younger agents). Figure 1 shows the structure of interaction between agents of different generations.

3.1. Actions and payoffs

The action played by the agent born in period t is denoted $A_t \in \{High, Low\}$. An agent chooses an action only once. The stage payoff to an agent playing A when another agent plays A' is denoted $u(A, A')$. The total payoff to the agent born at time t is

$$(1 - \lambda)u(A_t, A_{t-1}) + \lambda u(A_t, A_{t+1}), \quad (1)$$




FIGURE 1
Demographics

where A_{t-1} designates the action of the agent in the previous generation and A_{t+1} is the action of the agent in the next generation. Therefore, $\lambda \in [0, 1]$ is a measure of how much an agent weighs the play with the next generation compared with the previous generation; when $\lambda = 1$ an agent cares only about the next generation’s behaviour, whereas when $\lambda = 0$ an agent cares only about the previous generation’s actions. The λ parameter thus captures discounting as well as other aspects of the agent’s life, such as what portion of each period the agent is active (*e.g.* agents may be relatively active in the latter part of their lives, in which case λ could be $> 1/2$).¹¹ In our baseline analysis, we take $u(A, A')$ to be given by the following matrix:

$$\begin{array}{cc}
 & \begin{array}{cc} \textit{High} & \textit{Low} \end{array} \\
 \begin{array}{c} \textit{High} \\ \textit{Low} \end{array} & \begin{array}{cc} \beta, \beta & -\alpha, 0 \\ 0, -\alpha & 0, 0 \end{array}
 \end{array}$$

where β and α are both positive. This payoff matrix captures the notion that, from the static point of view, both *(High, High)* and *(Low, Low)* are static equilibria given this payoff matrix—and so conceivably both *High* and *Low* play could arise as stable patterns of behaviour. *(High, High)* is clearly the payoff-dominant or Pareto optimal equilibrium.¹²

3.2. Agents, signals, and information

There are two characteristics of agents in this society.

First, agents are distinguished by whether they choose an action to maximize the utility function given in equation (1). We refer to those who do so as “*endogenous*” agents. There are also some committed or “*exogenous*” agents who will choose an exogenously given action. This might be because these “*exogenous*” agents have different preferences, or because of some irrationality or trembles. Any given agent is exogenous with probability 2π , and such an agent is

11. This parameter incorporates tastes such as discounting, which may in turn be influenced by savings and investment technologies and other social factors. It may also proxy for other dimensions of social organization. For example, if one society has a higher retirement age than another, this will change the relative time one spends with agents of various ages. These aspects are all rolled into a single parameter here for parsimony and expositional simplicity.

12. Depending on the values of β and α , this equilibrium is also risk dominant, but this feature does not play a major role in our analysis. We also note that the normalization of a payoff of 0 for *Low* is for convenience, and inconsequential. In terms of strategic interaction, it is the difference of payoffs between *High* and *Low* conditional on expectations of what others will do that matter, which is then captured by the parameters α and β .

exogenously committed to playing each of the two actions, *High* and *Low*, with equal probability (and all of this, independently of all past events). Throughout, we assume that $\pi \in (0, \frac{1}{2})$, and in fact, we think of π as small (though this does not play a role in our formal results). With the complementary probability, $1 - 2\pi > 0$, the agent is endogenous and chooses whether to play *High* or *Low* when young, and is stuck with the same decision when old.

Secondly, agents can be either “*prominent*” or “*non-prominent*” (as well as being either endogenous or exogenous). A noisy signal of an action taken by a non-prominent agent of generation t is observed by the agent in generation $t + 1$. No other agent receives any information about this action. In contrast, the actions taken by prominent agents are perfectly observed by all future generations. We assume that each agent is prominent with probability q (again independently of other events) and non-prominent with the complementary probability, $1 - q$. This implies that an agent is exogenous prominent with probability $2\pi q$ and endogenous prominent with probability $(1 - 2\pi)q$.

The different types of agents and their probabilities in our model are thus:

	Non-prominent	Prominent
Endogenous	$(1 - 2\pi)(1 - q)$	$(1 - 2\pi)q$
Exogenous	$2\pi(1 - q)$	$2\pi q$

Unless otherwise stated, we assume that $0 < q < 1$, so that both prominent and non-prominent agents are possible. We refer to agents who are endogenous and non-prominent as *regular* agents.

We now explain the information structure in more detail. Let h^{t-1} denote the public history at time t , which includes a list of past prominent agents and their actions up to and including time $t - 1$, and let h_{t-1} denote the last entry in that history. In particular, we can represent what was publicly observed in any period as an entry with value in $\{High, Low, N\}$, where *High* indicates that the agent was prominent and played *High*, *Low* indicates that the agent was prominent and played *Low*, and *N* indicates that the agent was not prominent. We denote the set of h^{t-1} histories by \mathcal{H}^{t-1} .¹³

In addition to observing $h^{t-1} \in \mathcal{H}^{t-1}$, an agent of generation t , when born, receives a signal $s_t \in [0, 1]$ about the behaviour of the agent of the previous generation, where the restriction to $[0, 1]$ is without loss of any generality (although the signal is irrelevant when the agent of the previous generation is prominent, it simplifies the specification of strategies if all agents see a signal even though this does not matter in the results). This signal has a continuous distribution described by a density function $f_H(s)$ if $A_{t-1} = High$ and $f_L(s)$ if $A_{t-1} = Low$. Without loss of generality, we order signals such that higher s has a higher likelihood ratio for *High*; that is so that $\frac{f_H(s)}{f_L(s)}$ is non-decreasing in s . To simplify the analysis, we maintain the assumption that $\frac{f_H(s)}{f_L(s)}$ is strictly increasing in s , so that the strict Monotone Likelihood Ratio Principle (MLRP) holds, and we take the densities to be continuous and positive.

Let $\Phi(s, x)$ denote the posterior probability that $A_{t-1} = High$ given $s_t = s$ under the belief that a (non-prominent) agent of generation $t - 1$ plays *High* with probability x . Then

$$\Phi(s, x) \equiv \frac{f_H(s)x}{f_H(s)x + f_L(s)(1-x)} = \frac{1}{1 + \frac{(1-x)f_L(s)}{x f_H(s)}}. \quad (2)$$

The game begins with a prominent agent at time $t = 0$ playing action $A_0 \in \{High, Low\}$.

13. As will become clear, it is irrelevant whether a prominent agent was exogenous or endogenous in the greatest or least equilibrium in our model, though such information could be used in other equilibria as a correlating device.

3.3. *Strategies, semi-Markovian strategies and equilibrium*

We can write the strategy of an endogenous agent of generation t as:

$$\sigma_t : \mathcal{H}^{t-1} \times [0, 1] \times \{P, N\} \rightarrow [0, 1],$$

written as $\sigma_t(h^{t-1}, s_t, T_t)$ where $h^{t-1} \in \mathcal{H}^{t-1}$ is the public history of play, $s_t \in [0, 1]$ is the signal observed by the agent of generation t regarding the previous generation's action, and $T_t \in \{P, N\}$ denotes whether or not the the agent of generation t is prominent. The number $\sigma_t(h^{t-1}, s_t, T_t)$ corresponds to the probability that the agent of generation t plays *High*. We denote the strategy profile of all agents by the sequence $\sigma = (\sigma_1, \sigma_2, \dots, \sigma_t, \dots)$.

We show below that the most relevant equilibria for our purposes involve agents ignoring histories that come before the most recent prominent agent. These histories are not payoff-relevant provided others are following similar strategies. We call these *semi-Markovian* strategies.

Semi-Markovian strategies are specified for endogenous agents as functions $\sigma_\tau^{SM} : \{High, Low\} \times [0, 1] \times \{P, N\} \rightarrow [0, 1]$, written as $\sigma_\tau^{SM}(a, s, T)$ where $\tau \in \{1, 2, \dots\}$ is the number of periods since the last prominent agent, $a \in \{High, Low\}$ is the action of the last prominent agent, $s \in [0, 1]$ is the signal of the previous generation's action, and again $T \in \{P, N\}$ is whether or not the current agent is prominent.

With some abuse of notation, we sometimes write $\sigma_t = High$ or Low to denote a strategy or semi-Markovian strategy that corresponds to playing *High (Low)* with probability one.

We analyse *Bayesian equilibria*, which we simply refer to as *equilibria*. More specifically, an equilibrium is a profile of endogenous players' strategies together with a specification of beliefs (conditional on each history and signal) such that: the endogenous players' strategies are best responses to the profile of strategies given their beliefs (conditional on each possible history and signal) and given their prominence; and beliefs are derived from the strategies and history according to Bayes' rule. Since $0 < q < 1$ and $\pi > 0$, all feasible histories and signal combinations are possible, and the sets of Bayesian equilibria, perfect Bayesian equilibria and sequential equilibria coincide.¹⁴

4. EQUILIBRIUM

We start with a few observations about best responses and then move to the characterization of the structure of equilibria.

4.1. *Best responses*

Given the utility function (1), an endogenous agent of generation t will have a best response of $A = High$ if and only if

$$(1 - \lambda)\phi_{t-1}^t + \lambda\phi_{t+1}^t \geq \frac{\alpha}{\beta + \alpha} \equiv \gamma, \tag{3}$$

where ϕ_{t-1}^t is the (equilibrium) probability that the agent of generation t assigns to the agent from generation $t - 1$ having chosen $A = High$. ϕ_{t+1}^t is defined similarly, except that it is also conditional

14. To be precise, any particular signal still has a 0 probability of being observed, but posterior beliefs are well-defined subject to the usual measurability constraints.

When $q = 0$ or $\pi = 0$ (contrary to our maintained assumptions), some feasible combinations of histories and signals have zero probability and then Bayesian and perfect Bayesian equilibria (appropriately defined for a continuum of signals) can differ. In that case, it is necessary to carefully specify which beliefs and behaviours off the equilibrium path are permitted as part of an equilibrium. For the sake of completeness, we provide a definition of equilibrium in Appendix A that allows for those corner cases.

on agent t playing *High*. Thus, it is the probability that the agent of generation t assigns to the next generation choosing *High* conditional on his or her own choice of *High*. Defining ϕ_{t+1}^t as this conditional probability is useful, since playing *Low* guarantees a payoff of 0, and the relevant calculation for agent t is the consequence of playing *High*, and will thus depend on ϕ_{t+1}^t .

The parameter γ encapsulates the payoff information of different actions in an economical way.¹⁵ In what follows, γ (rather than α and β separately) will be the main parameter affecting behaviour and the structure of equilibria.

4.2. Existence of equilibria

We say that a strategy σ is a *cut-off strategy* if for each t , h^{t-1} such that $h_{t-1} = N$ and $T_t \in \{P, N\}$, there exists $c_t(h^{t-1}, T_t)$ such that $\sigma_t(h^{t-1}, s, T_t) = 1$ if $s > c_t(h^{t-1}, T_t)$ and $\sigma_t(h^{t-1}, s, T_t) = 0$ if $s < c_t(h^{t-1}, T_t)$.¹⁶ Clearly, setting $\sigma_t(h^{t-1}, s, T) = 1$ (or 0) for all s is a special case of a cut-off strategy.¹⁷

We can represent a cut-off strategy profile by the sequence of cut-offs for each possible history:

$$c = (c_1^N(h^0), c_1^P(h^0), \dots, c_t^N(h^{t-1}), c_t^P(h^{t-1}), \dots),$$

where $c_t^T(h^{t-1})$ denotes the cut-off by agent of prominence type $T \in \{P, N\}$ at time t conditional on history h^{t-1} . In what follows, we define “greatest equilibria” using the natural Euclidean partial ordering in terms of the (infinite) vector of equilibrium cut-offs.

Proposition 1.

1. All equilibria are in cut-off strategies.
2. There exists an equilibrium in semi-Markovian cut-off strategies.
3. The set of equilibria and the set of semi-Markovian equilibria form complete lattices, and the greatest (and least) equilibria of the two lattices coincide.

The proof of this proposition relies on an extension of the well-known results for (Bayesian) games of strategic complements to a setting with an infinite number of players, presented in Appendix A. The proof of this proposition, like those of all remaining results in the article, is also provided in the Appendix.

Given the results in Proposition 1, we focus on extremal equilibria. Since the lattice of equilibria is complete, there is a unique maximal (and hence greatest or maximum) equilibrium and unique minimal (and hence least or minimum) equilibrium. Because, again from Proposition 1, these extremal equilibria are semi-Markovian, their analysis will be quite tractable. As is generally the case in this class of models, non-extremal equilibria can be much more complicated, and we will not focus on them.

We further simplify the exposition by focusing on the greatest equilibrium since each statement has an immediate analog for the least equilibrium, which we omit for brevity.

15. In particular, γ is a measure of how “risky” the *High* action is—in the sense that it corresponds to the probability that the other side should be playing *High* to make a player indifferent between *High* and *Low*. Put differently, it is the “size of the basin of attraction” of *Low* as an equilibrium.

16. Note that specification of any requirements on strategies when $s = c_t(h^{t-1}, T_t)$ is inconsequential as this is a zero probability event.

17. If $h_{t-1} = P$, the agent of generation t still receives a signal and the cutoff then simply determines a probability of playing *Low*. We can restrict the cutoffs to be 0 or 1 to obtain pure strategy equilibria (at least one of these is always a best response for agents following a prominent agent), or allow it to vary for mixed strategy equilibria. Extremal equilibria involve 0 or 1 cutoffs for agents following prominent agents and so our proofs work either way.

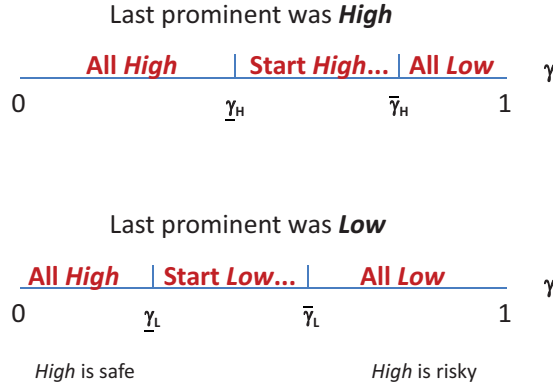


FIGURE 2

A depiction of the play of endogenous players in the greatest equilibrium, as a function of the underlying attractiveness of playing *Low* (γ), broken down as a function of the play of the last prominent player

4.3. A characterization of greatest equilibrium play

The structure of equilibria depends on the values of the parameters λ , γ , π , q , and the signal structure. To provide the sharpest intuitions, we focus on two key parameters: λ , which captures how forward/backward looking the society is, and $\gamma \equiv \alpha / (\beta + \alpha)$, which designates how risky playing *High* is (in the sense defined in the footnote 15).

Figure 2 summarizes the structure of equilibria simply as a function of γ (also fixing λ). This figure shows that the structure of equilibria will in general depend on the play of the last prominent agent. For example, when it is *High*, there is a key threshold, $\underline{\gamma}_H$, such that for $\gamma \leq \underline{\gamma}_H$, there is a *social norm* where all endogenous agents play *High* regardless of their signal. Above this threshold, a *High* social norm is no longer an equilibrium. As *High* becomes riskier (γ increases), generations immediately following a *High* prominent player will play *High*, but subsequent play deteriorates as agents become increasingly sceptical that the previous generation played *High*. Once *High* play is sufficiently risky ($\gamma > \bar{\gamma}_H$), then even a small chance of facing a future exogenous *Low* player is enough so that *High* is never played.

The picture following a prominent *Low* play is similar, but with different thresholds.

In this subsection, we focus on conditions for *High* and *Low* social norms, in which endogenous players follow the last prominent play regardless of their signals. The dynamics of play in the intermediate regions, where endogenous agents respond to their signals, is characterized in Section 4.4. Finally, in Section 4.5, we examine the role of endogenous prominent agents and their ability to lead a society away from a *Low* social norm.

Recall from (3) that an endogenous player is willing to play *High* if and only if $(1 - \lambda)\phi_{t-1}^t + \lambda\phi_{t+1}^t \geq \gamma$. Therefore, for *High* to be played by all endogenous players, it must be that (3) holds for all possible signals when all other endogenous agents (are expected to) play *High* and the last prominent agent played *High*. This is of course equivalent to (3) holding conditional on the lowest possible signal (when all endogenous agents play *High* and the last prominent agent played *High*) because this would ensure that *High* is a best response for any signal, effectively locking it in following the *High* play of the last prominent agent. The threshold for this to be the case is

$$\gamma \leq \underline{\gamma}_H \equiv (1 - \lambda)\Phi(0, 1 - \pi) + \lambda(1 - \pi). \tag{4}$$

The expression is intuitive noting that $\Phi(0, 1 - \pi)$ is the probability of last generation having played *High* conditional on the lowest possible signal, $s=0$, and $1 - \pi$ is the probability of

the next generation playing *High* (since only agents endogenously committed to *Low* will do otherwise). We will see that above this threshold, endogenous agents immediately following *High* prominent play will play *High*, because they are (fairly) confident that the last player played *High*, but this confidence will gradually erode over time, and the likelihood of *Low* play will increase (Section 4.4).

Equivalently, (in the case where the numerator and denominator below are positive) rearranging (4), *High* will lock in following a prominent *High* if and only if

$$\lambda \geq \frac{\gamma - \Phi(0, 1 - \pi)}{1 - \pi - \Phi(0, 1 - \pi)}.$$

This way of expressing the conditions for a *High* social norm is also intuitive. Recalling that λ is the weight that an agent places on the payoff from the interaction with the younger generation, it implies that if a society is sufficiently forward looking, then it is possible to sustain all *High* following a prominent *High*. Above this threshold, expectations that future generations will play *High* is sufficient to keep a player playing *High*, even under the lowest possible signal from the previous generation. In contrast, if agents are more backward looking than implied by this threshold, then their play is eventually molded by history (and the particular signals that they observe from the older generation).

Another key threshold is reached when *High* play becomes sufficiently risky that, regardless of signal, the possibility of facing future *Low* is so overwhelmingly costly that all *Low* becomes the only possible play. That threshold is described by $\gamma = \bar{\gamma}_H$.¹⁸ The characterization of greatest equilibrium play following a prominent *High* is then pictured in Figure 3. The shape in this figure follows since when $\lambda = 1$, agents only care about the next generation and so either players can sustain all *High* or all *Low*—making history completely irrelevant. When $\lambda = 0$, agents are entirely backward looking and so after a prominent play at least the first endogenous agent will play *High*, so that $\bar{\gamma}_H = 1$.

There are analogous thresholds that apply after a prominent *Low*. Following a similar logic to that above, all endogenous players indefinitely playing *High* following a prominent *Low* is possible if and only if

$$\gamma \leq \underline{\gamma}_L \equiv \lambda(1 - \pi), \quad (5)$$

since this requires that they prefer to play *High* even following a *Low* prominently, and it is entirely leveraged by their expectations of future play. This threshold can be expressed again in terms of λ , $\lambda \geq \gamma / (1 - \pi)$, and requires agents to be sufficiently forward looking—caring more about their match with the younger generation. It can also be noted that when (5) is satisfied then so is (4), and a *High* social norm can be sustained regardless of the actions of prominent agents.

The more interesting threshold is $\gamma = \bar{\gamma}_L$. Above this threshold, a *Low* prominent play is followed by all endogenous agents choosing *Low* regardless of their signal. Although, as is also the case for $\bar{\gamma}_H$, there is no general closed-form solution for $\bar{\gamma}_L$, there is one when $\underline{\gamma}_H \geq \bar{\gamma}_L$. In this case (Appendix A),

$$\bar{\gamma}_L = \gamma_L^* \equiv (1 - \lambda) \Phi(1, \pi) + \lambda(1 - \pi), \quad (6)$$

which is conveniently symmetric to $\underline{\gamma}_H$: *Low* will play prominent *Low* regardless of signal if γ is greater than a λ -weighted average of $1 - \pi$ (probability of *Low* play for the next generation)

18. This threshold does not have a closed-form solution, though it is a well-defined monotone function of parameters, including λ (Appendix A). An upper bound on this threshold is provided by the most optimistic belief an agent following a *High* prominent play could have: $\bar{\gamma}_H \leq (1 - \lambda) + \lambda(1 - \pi)$, and this is the exact threshold in the extreme cases of $\lambda = 0, 1$, but not necessarily in between. For intermediate values of λ , some future agents might not play *High* even with a very high signal from the first period agent who plays *High* (depending on his or her expectations about the future).

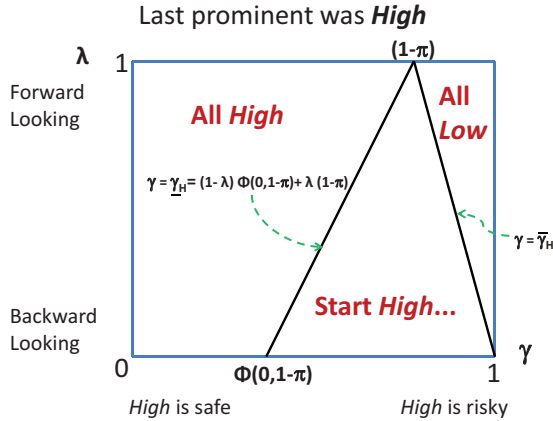


FIGURE 3

Play of endogenous players in the greatest equilibrium, as a function of how forward-looking they are (λ) and how risky High play is (γ), given that the last prominent player played High

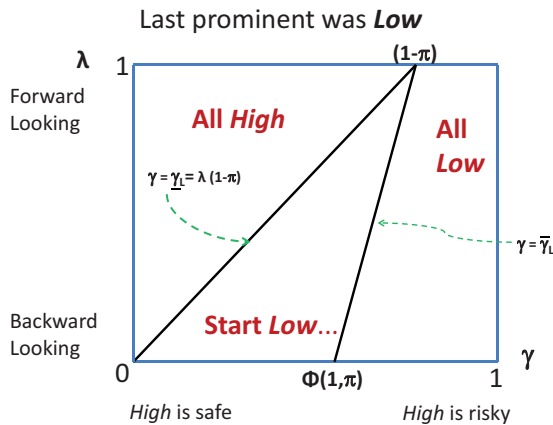


FIGURE 4

Play of endogenous players in the greatest equilibrium, as a function of how forward looking they are (λ) and how risky High play is (γ), given that the last prominent player played Low.

and $\Phi(1, \pi)$ (the expectation of Low play in the previous generation given the highest possible signal, $s = 1$).

This structure of equilibrium following Low prominent play is pictured in Figure 4, and the full characterization of the greatest equilibrium is provided in Proposition 2.

Proposition 2. *In the greatest equilibrium:*

- *If the last prominent play was High then:*
 - *If $\gamma \leq \underline{\gamma}_H$, then all endogenous agents play High;*
 - *if $\underline{\gamma}_H < \gamma < \bar{\gamma}_H$, then endogenous agents start playing High following prominent High, but then play Low for some signals; and*
 - *if $\bar{\gamma}_H < \gamma$, then all endogenous agents play Low.*

- *If the last prominent play was Low then:*
 - *If $\gamma \leq \underline{\gamma}_L$, then all endogenous agents play High,*
 - *if $\underline{\gamma}_L < \gamma < \bar{\gamma}_L$, then endogenous agents start playing Low following prominent Low, but then play High for some signals; and*
 - *if $\bar{\gamma}_L < \gamma$, then all endogenous agents play Low.*

Proposition 2 makes the role of history clear: for parameter values such that all *High* is not an equilibrium, the social norm is determined by history for at least some time. In particular, if prominent agents are rare, then society follows a social norm established by the last prominent agent for an extended period of time.

Nevertheless, our model also implies that social norms are not everlasting: switches in social norms take place following the arrival of exogenous prominent agents (committed to the opposite action). Thus when q is small, a particular social norm, determined by the play of the last prominent agent, emerges and persists for a long time, disturbed only by the emergence of another (exogenous) prominent agent who chooses the opposite action and initiates a different social norm.

Some key comparative statics are also clear from this proposition (more detailed comparative statics are given in Appendix B). First, Figures 3 and 4 make it clear that as λ increases—so that individuals care more about the future—all *High* following both *High* and *Low* prominent play is an equilibrium for a larger set of values of other parameters (*e.g.* more values of γ). This is intuitive: when agents are more forward looking, coordinating on *High* becomes easier (for the same reason that all *High* became the unique pure-strategy equilibrium and the complete information version of the model). The impact of an increase in the probability of exogenous agents, π , is to reduce $\bar{\gamma}_H$ since it increases the probability of exogenous *Low* players and thus the likelihood that an agent may want to switch to *Low* following a low signal. However, the corresponding impact on $\underline{\gamma}_L$ is ambiguous, since it increases the presence of both *High* and *Low* exogenous players.¹⁹ Finally, as information becomes more precise so that $\Phi(0, 1 - \pi)$ decreases (*i.e.* the 0 signal indicates *Low* play with a higher probability), all *High* is more difficult to sustain. This is because as agents receive more accurate information about the play of the previous generations, it becomes harder to convince them to play *High* following a signal indicating an exogenous *Low* play.

It is also useful to compare behaviour across the two possibilities for the last prominent play. First, clearly $\underline{\gamma}_H \geq \underline{\gamma}_L$, so that it is easier to sustain all *High* play following a prominent *High*, and similarly $\bar{\gamma}_H \geq \bar{\gamma}_L$, so that there are more situations where all *Low* is the only possible continuation following a prominent *Low* than a prominent *High*. More interestingly, we may also have $\bar{\gamma}_L < \underline{\gamma}_H$, so that the prominent agents lock in subsequent behaviour—a strong form of history dependence. Provided that $\lambda < 1$, the condition that $\underline{\gamma}_H > \bar{\gamma}_L$ (which is equivalent to $\underline{\gamma}_L < \underline{\gamma}_H$ with $\underline{\gamma}_L^*$ given by (6)) can be simply written as $\Phi(0, 1 - \pi) > \Phi(1, \pi)$.²⁰ Defining the least and greatest likelihood ratios as

$$m \equiv \frac{f_H(0)}{f_L(0)} < 1 \text{ and } M \equiv \frac{f_H(1)}{f_L(1)} > 1,$$

19. In particular, this depends how accurate signals about past play are because an equilibrium with all *Low* crucially depends on the inference about past play.

20. Recall that if $\bar{\gamma}_L \leq \underline{\gamma}_H$, then $\bar{\gamma}_L = \underline{\gamma}_L^*$ as defined in (6). Therefore, $\bar{\gamma}_L < \underline{\gamma}_H$ is equivalent to $\underline{\gamma}_L^* < \underline{\gamma}_H$.

the (necessary and sufficient) condition for $\bar{\gamma}_L < \underline{\gamma}_H$ is $\lambda < 1$ and

$$\frac{(1-\pi)^2}{\pi^2} > \frac{M}{m}. \tag{7}$$

This requires that m is not too small relative to M , so that signals are sufficiently noisy. Intuitively, when the greatest equilibrium involves all endogenous agents playing *Low*, this must be the unique “continuation equilibrium” (given the play of the last prominent agent). Thus, the condition that $\gamma > \bar{\gamma}_L$ ensures the uniqueness of the continuation equilibrium following a prominent agent playing *Low*—otherwise all *Low* could not be the greatest equilibrium. In this light, it is intuitive that this condition should require signals to be sufficiently noisy. Otherwise, players would react strongly to signals from the previous generation and could change to *High* behaviour when they receive a strong signal indicating *High* play in the previous generation and also expecting the next generation to receive accurate information regarding their own behaviour. Noisy signals ensure that each agent has a limited ability to influence the future path of actions and thus prevent multiple equilibria supported by coordinating on relatively precise signals of past actions.

4.4. The reversion of play over time

We now complete the characterization of the greatest equilibrium for the cases where $\underline{\gamma}_H < \gamma \leq \bar{\gamma}_H$ and $\underline{\gamma}_L < \gamma \leq \bar{\gamma}_L$, which involve the *reversion* of the play of regular players.

For example, when all *High* is not an equilibrium, then *High* play deteriorates following a prominent play of *High*. This is a consequence of a more general monotonicity result that shows that cut-offs always move in the same direction: that is, either thresholds are monotonically non-increasing or monotonically non-decreasing over time, so that *High* play either becomes monotonically more likely (if the last prominent play was *Low*) or monotonically less likely (if the last prominent play was *High*). So, for instance, when greatest equilibrium behaviour is not completely *High*, then *High* play deteriorates over time, meaning that as the distance from the last prominent *High* agent increases, the likelihood of *High* behaviour decreases and corresponding cut-offs increase.

Since we are focusing on semi-Markovian equilibria, we denote, with a slight abuse of notation, the cut-offs used by prominent and non-prominent agents τ periods after the last prominent agent by c_τ^P and c_τ^N , respectively. We say that *High* play is *non-increasing* over time if $(c_\tau^P, c_\tau^N) \leq (c_{\tau+1}^P, c_{\tau+1}^N)$ for each τ . We say that *High* play is *decreasing* over time, if, in addition, whenever $(c_\tau^P, c_\tau^N) \neq (0, 0)$ and $(c_\tau^P, c_\tau^N) \neq (1, 1)$, it follows that $(c_\tau^P, c_\tau^N) \neq (c_{\tau+1}^P, c_{\tau+1}^N)$. The concepts of *Low* play being non-decreasing and increasing over time are defined analogously.

The definition of decreasing or increasing play implies that when the cut-offs for endogenous agents are non-degenerate, they must strictly increase over time. This implies that, unless *High* play completely dominates, *High* play strictly decreases over time.

Proposition 3.

1. In the greatest equilibrium, cut-off sequences (c_τ^P, c_τ^N) are monotone: following a prominent agent choosing *High*, (c_τ^P, c_τ^N) are non-decreasing and following a prominent agent choosing *Low*, they are non-increasing.
2. If $\underline{\gamma}_H < \gamma < \bar{\gamma}_H$, then in the greatest equilibrium, *High* play is decreasing over time following *High* play by a prominent agent.
3. If $\underline{\gamma}_L < \gamma < \bar{\gamma}_L$, then in the greatest equilibrium, *High* play is increasing over time following *Low* play by a prominent agent.




FIGURE 5

Reversion of play from *High* to the greatest steady state

Figure 5 illustrates the behaviour of the cut-offs and the corresponding probabilities of *High* play for regular agents following a *High* prominent play. For the reasons explained in the paragraph preceding Proposition 3, prominent endogenous agents will have lower cut-offs and higher probabilities of *High* play than regular agents. Depending on the specific levels of parameters, it could be that prominent endogenous agents all play *High* for all signals and times, or it could be that their play reverts too.

The intuition for Proposition 3 is interesting. Immediately following a *High* prominent action, an agent knows for sure that he or she is facing *High* in the previous generation. Two periods after a *High* prominent action, he or she is playing against an agent from the older generation who knew for sure that they themselves were facing *High* in the previous generation. Thus, his or her opponent was likely to have chosen *High* himself. Nevertheless, since $\gamma > \underline{\gamma}_H$, there are some signals for which he or she will be sufficiently confident that the previous generation was of exogenous type and chose *Low* instead. Now consider an agent three periods after a *High* prominent action. For this agent, not only is there the possibility that one of the two previous agents was exogenous and committed to *Low* play, but also the possibility that his or her immediate predecessor received an adverse signal and decided to play *Low* instead. Thus, he or she is even more likely to interpret adverse signals as coming from *Low* play than was his or her predecessor. This reasoning highlights the tendency towards higher cut-offs and less *High* play over time. In fact, there is another more subtle force pushing in the same direction. Since $\gamma > \underline{\gamma}_H$, each agent also realizes that even when he or she chooses *High*, the agent in the next generation may receive an adverse signal, and the farther this agent is from the initial prominent agent, the more likely are the signals resulting from his or her choice of *High* to be interpreted as coming from a *Low* agent. This anticipation of how his or her signal will be interpreted—and thus become more likely to be countered by a play of *Low*—as the distance to the prominent agent increases creates an additional force towards reversion.

The converse of this intuition explains why there is improvement of *High* play over time starting with a prominent agent choosing *Low*. The likelihood of a given individual encountering *High* play in the previous generation increases as the distance to prominent agent increases, as Figure 5 shows.

Proposition 3 also implies that behaviour converges to a limiting (steady-state) distribution along sample paths where there are no prominent agents. Two important caveats need to be noted, however. First, this limiting distribution depends on the starting point. Moreover, the

limiting distribution following a prominent agent playing *Low* may be different from the limiting distribution following a prominent agent playing *High*. This can be seen by considering the case where $\bar{\gamma}_L < \gamma < \underline{\gamma}_H$, already discussed above: here the (trivial) limiting distribution is a function of the action of the last prominent agent that completely locks in play until the next prominent agent. Secondly, although there is convergence to a limiting distribution along sample paths without prominent agents, there is in general no convergence to a stationary distribution because of the arrival of exogenous prominent agents. In particular, provided that $q > 0$ (and since $\pi > 0$), the society will necessarily fluctuate between different patterns of behaviour. For example, when $\bar{\gamma}_L < \gamma < \underline{\gamma}_H$, as already pointed out following Proposition 2, the society will fluctuate between social norms of *High* and *Low* play as exogenous prominent agents arrive and choose different actions (even if this happens quite rarely).

Note also that there is an interesting difference between the ways in which reversion occurs in Proposition 3 starting from *Low* versus *High* play. Endogenous prominent agents are always at least weakly as willing to play *High* as are regular agents, since they will be observed and are thus more likely to have their *High* play reciprocated by the next agent. Thus, their cut-offs are always weakly lower and their corresponding probability of playing *High* is higher. Hence, if play starts at *High*, then it is the regular agents who are reverting more, that is, playing *Low* with a greater probability. In contrast, if play starts at *Low*, then it is the prominent agents who revert more, that is, playing *High* with a greater probability (and eventually leading to a new prominent history beginning with a *High* play). It is possible, for some parameter values, that one type of endogenous player sticks with the play of the last prominent agent (prominent endogenous when starting with *High*, and non-prominent endogenous when starting with *Low*), whereas the other type of endogenous player strictly reverts in play.²¹

4.5. Breaking the *Low* social norm

In this subsection, we illustrate how prominent agents can exploit their greater (and common knowledge) visibility to future generations to play a leadership role and break the *Low* social norm to induce a switch to *High* play. Consider a *Low* social norm where all regular agents play *Low*.²² Suppose that at generation t there is an endogenous prominent agent. The key question analysed in the next proposition is when an endogenous prominent agent would like to switch to *High* play to change the existing social norm.

Let $\tilde{\gamma}_L$ denote the threshold such that above this level, in the greatest equilibrium, all *regular* (endogenous but non-prominent) players choose *Low* following a prominent *Low*. It is straightforward to see that $0 < \tilde{\gamma}_L < \bar{\gamma}_L$ (provided $\lambda < 1$), and so this is below the threshold where all endogenous players choose *Low* (because, as we explained above, prominent endogenous agents are more willing to switch to *High* than regular agents).

Proposition 4. *Consider the greatest equilibrium:*

1. *Suppose that $\tilde{\gamma}_L \leq \gamma < \min\{\gamma_L^*, \underline{\gamma}_H\}$ and that the last prominent agent has played *Low*. Then there exists a cut-off $\tilde{c} < 1$ such that an endogenous prominent agent playing at least two periods after the last prominent agent and receiving a signal $s > \tilde{c}$ will choose *High* and break the *Low* social norm.*

21. The asymmetry between reversion starting from *Low* versus *High* play we are emphasizing here is distinct from the asymmetry that results from our focus on the greatest equilibrium. In particular, this asymmetry is present even if we focus on the least equilibrium.

22. The social norm in question here involves one in which all regular agents, but not necessarily endogenous prominent agents, play *Low*.

2. Suppose that $\gamma < \min\{\tilde{\gamma}_L, \underline{\gamma}_H\}$ and that the last prominent agent played *Low*. Then there exists a sequence of decreasing cut-offs $\{\bar{c}_\tau\}_{\tau=2}^\infty < 1$ such that an endogenous prominent agent playing $\tau \geq 2$ periods after the last prominent agent and receiving a signal $s > \bar{c}_\tau$ will choose *High* and switch play from the path of convergence to steady state to a *High* social norm.

The proposition and subsequent results are proven in [online Appendix B](#), which is available as Supplementary Material.

The results in this proposition are important and intuitive. Their importance stems from the fact that they show how prominent agents can play a crucial leadership role in society. In particular, the first part shows that starting with the *Low* social norm, a prominent agent who receives a signal from the last generation that is not too adverse (so that there is some positive probability that he or she is playing an exogenous type committed to *High* play) will find it profitable to choose *High*, and this will switch the entire future path of play, creating a *High* social norm instead. The second part shows that prominent agents can also play a similar role starting from a situation that does not involve a strict *Low* social norm—instead, starting with *Low* and reverting to a steady-state distribution. In this case, the threshold for instigating such a switch depends on how far they are from the last prominent agent who has chosen *Low*.

The intuition for these results clarifies how history and expectations shape the evolution of cooperation. Prominent agents can play a leadership role because they can exploit their impact on future expectations and their visibility by future generations to change a *Low* social norm into a *High* one. In particular, when the society is stuck in a *Low* social norm, regular agents do not wish to deviate from this, because they know that the previous generation has likely chosen *Low* and also that even if they were to choose *High*, the signal generated by their action would likely be interpreted by the next generation as coming from a *Low* action. For a prominent agent, the latter is not a concern, since his or her action is perfectly observed by the next generation. Moreover and perhaps more importantly from an economic point of view, his or her deviation from the *Low* social norm can influence the expectations of all future generations, reinforcing the incentives of the next generation to also switch their action to *High*.

5. DISCUSSION AND EXTENSIONS

In this section, we first clarify the role of prominence in coordinating expectations and enabling endogenous leadership. We then outline how our results extend to the case in which there are $n > 1$ agents within each generation. We also show how our analysis applies with different structures of payoffs, how imperfect prominence affects our results, and how the framework can be generalized to endogenize prominence.

5.1. Prominence, expectations and leadership

In this subsection we highlight the role of prominence in our model, emphasizing that prominence is different from (stronger than) simply being observed by the next generation with certainty. In particular, the fact that prominence involves being observed by all subsequent generations with certainty plays a central role in our results. To clarify this, we consider four scenarios.

In each scenario, for simplicity, we assume that there is a starting non-prominent agent at time 0 who plays *High* with probability $x_0 \in (0, 1)$, where x_0 is known to all agents who follow, and generates a signal for the first agent in the usual way. All agents after time 1 are not prominent.

In every case, all agents (including time 1 agents) are endogenous with probability $(1 - 2\pi)$ as before.

Scenario 1. The agent at time 1 is not prominent and his or her action is observed with the usual signal structure.

Scenario 2. The agent at time 1's action is observed perfectly by the period 2 agent, but not by future agents.

Scenario 2'. The agent at time 1 is only observed by the next agent according to a signal, but then is subsequently perfectly observed by all agents who follow from time 3 onwards.

Scenario 3. The agent at time 1 is prominent, and all later agents are viewed with the usual signal structure.

Clearly, as we move from Scenario 1 to Scenario 2 (or 2') to Scenario 3, we are moving from a non-prominent agent to a prominent one, with Scenarios 2 and 2' being hybrids, where the agent of generation $t = 1$ has greater visibility than a non-prominent agent but is not fully prominent in terms of being observed forever after.

We focus again on the greatest equilibrium and let $c^k(\lambda, \gamma, f_H, f_L, \pi)$ denote the cut-off signal above which the first agent (if endogenous) plays *High* under scenario k as a function of the underlying setting.

Proposition 5. *The cut-offs satisfy $c^2(\cdot) \geq c^3(\cdot)$ and $c^1(\cdot) \geq c^{2'}(\cdot) \geq c^3(\cdot)$, and there are settings $(\lambda, \gamma, f_H, f_L, \pi)$ for which the inequalities are strict.*

The intuition for this result is instructive. First, comparing Scenario 2 to Scenario 3, the former has the same observability of the action by the next generation (the only remaining generation that directly cares about the action of the agent) but not the common knowledge that future generations will also observe this action. This means that future generations will not necessarily coordinate on the basis of a choice of *High* by this agent, and this discourages *High* play by the agent at date $t = 2$, and through this channel, it also discourages *High* play by the agent at date $t = 1$, relative to the case in which there was full prominence. The comparison of Scenario 2' to Scenario 1 is perhaps more surprising. In Scenario 2', the agent at date $t = 1$ knows that his or her action will be seen by future agents, so if he or she plays *High*, then this gives agent 3 extra information about the signals that agent 2 is likely to observe. This creates strong feedback effects in turn affecting agent 1. In particular, agent 3 would choose a lower cut-off for a given cut-off of agent 2 when he or she sees *High* play by agent 1. But knowing that agent 3 is using a lower cut-off, agent 2 will also find it beneficial to use a lower cut-off. This not only feeds back to agent 3, making him or her even more aggressive in playing *High*, but also encourages agent 1 to play *High* as he or she knows that agent 2 is more likely to respond with *High* himself. In fact, these feedback effects continue and affect all future agents in the same manner, and in turn, the expectation that they will play *High* with a higher probability further encourages *High* play by agents 1 and 2. Thus, one can leverage things upwards even through delayed prominence.

Notably, a straightforward extension of the proof of Proposition 5 shows that the same comparisons hold if we replace “time 3” in Scenarios 2 or 2' with “time t ” for any $t \geq 3$.²³

23. There are two omitted comparisons: between scenarios 2 and 2' and between scenarios 1 and 2. Both of these are ambiguous. It is clear why the comparison between scenarios 2 and 2' is ambiguous as those information structures are not nested. The ambiguity between scenarios 1 and 2 is more subtle, as one might have expected that $c^1 \geq c^2$. The

5.2. Multiple agents within generations

We now return to a variation on the model of collective action outlined in Section 2 (now without the neighbourhood structure) to show how our results extend to an environment with multiple agents within each generation. For simplicity, we work directly with the same payoff matrix and notation as for the main results above, with the changes as indicated below.

There are n agents within each generation, and each interacts with all agents from the previous generation and from the next generation.²⁴ The expected utility of agent i from generation t is then

$$(1-\lambda) \sum_{j=1}^n u(A_{i,t}, A_{j,t-1}) + \lambda \sum_{j=1}^n u(A_{i,t}, A_{j,t+1}).$$

There is at most one prominent agent within a generation. If there is such a prominent agent, then his or her action is also taken by $m-1$ other (randomly chosen) agents of the same generation. Thus, a prominent agent is able to coordinate the actions of $m-1$ agents within their generation.

The information structure is as follows. If there is a prominent agent in generation $t-1$ then each agent of generation t observes the action of the prominent agent and nothing else about that generation. Moreover, prominent agents' behaviours are observed forever. If there is no prominent agent in generation $t-1$, agent i of generation t observes a signal $s_{i,t-1}$ generated from a randomly selected agent from the previous generation.

Any agent is exogenous with a probability 2π as before (except if some other agent is prominent within their generation, in which case their action may be directed by that prominent agent).

First consider the case in which $m=n$, so that a prominent agent is able to coordinate the actions of all other agents within his or her generation (which are thus all identical to his or her action). Then the results presented for the baseline model extend as follows.

The thresholds characterizing the structure of greatest/least equilibria extend from Proposition 2. To economize on space, we only discuss a couple of these thresholds and then move on to the case in which $m < n$.

Conditional upon seeing a signal s and given a prior belief that the probability that regular agents of the previous generation play *High* is x , the expected fraction of agents from the previous generation who play *High* is

$$\Phi_n(s, x) = \frac{1}{n} \Phi(s, x) + \frac{n-1}{n} x. \quad (8)$$

The threshold γ_H^n for *High* to be a best response when all future regular agents are expected to play *High* (independently of the last prominent play) is again $\gamma \leq (1-\lambda)\Phi_n(0, 1-\pi) + \lambda(1-\pi)$,

reason why this is not always the case is interesting. When signals are sufficiently noisy and x_0 is sufficiently close to 1, under scenario 1 agent 2 would prefer to choose *High* regardless of the signal he or she receives. This would in turn induce agent 1 to choose *High* for most signals. When agent 2 instead observes agent 1's action perfectly as in scenario 2, then (provided that λ is not too high) he or she will prefer to match this action, *i.e.* play *High* only when agent 1 plays *High*. The expectation that he or she will play *Low* in response to *Low* under scenario 2 then leads agents born in periods 3 and later to be more pessimistic about the likelihood of facing *High* and they will thus play *Low* with greater probability than they would do under scenario 1. This then naturally feeds back and affects the trade-off facing agent 2 and he or she may even prefer to play *Low* following *High* play; in response the agent born in period 1 may also choose *Low*. All of this ceases to be an issue if the play of the agent born at date 1 is observed by all future generations (as in scenarios 2' and 3), since in this case the ambiguity about agent 2's play disappears.

24. We could also allow the agents to interact within their own generation. This would require a third weight (in addition to λ and $1-\lambda$) and also introduces an extra layer of coordinations among the current generation. Nonetheless, there still exist greatest and least equilibria, and our main results are similar. For parsimony, we stick with the simpler model.

by the same reasoning as before. Thus,

$$\underline{\gamma}_H^n \equiv (1 - \lambda) \left[\frac{1}{n} \Phi(0, 1 - \pi) + \frac{n-1}{n} (1 - \pi) \right] + \lambda(1 - \pi). \tag{9}$$

This expression takes into account that signals are less informative about behaviour. Clearly, $\underline{\gamma}_H^n$ is increasing in n , which implies that the set of parameters under which *High* play will follow *High* prominent play is greater when there are more players within each generation. This is because the signal each one receives becomes less informative about the overall actions that a player faces, and thus they put less weight on the signal and more weight on the action of the last prominent agent.

This reasoning enables us to directly generalize the results of Proposition 2 and also determine how these thresholds vary with n (Proposition 6).

What happens if $m < n$? In this case, the analysis is more complicated, but similar results apply. In particular, a sufficient threshold for all endogenous agents to play *High* following a prominent *High* from the previous generation is given by the following reasoning. The worst posterior that an agent can have after a prominent *High* in the previous generation would be given by assuming that the history is such that all $n - m$ agents who have not been coordinated by the prominent agent are playing *Low*, which would be m/n . Therefore, we can write

$$\underline{\gamma}_H^{n,m} \equiv \min \left\{ (1 - \lambda) \frac{m}{n} + \lambda(1 - \pi), \underline{\gamma}_H^n \right\} \tag{10}$$

The intuition is clear. $\gamma \leq \underline{\gamma}_H^{n,m}$ is sufficient to ensure that all endogenous agents in the first generation following a prominent *High* play will also choose *High* if they expect the next generation to all play *High*, which is then guaranteed since $\gamma \leq \underline{\gamma}_H^n$. A similar argument extends to define a threshold $\bar{\gamma}_L^{n,m}$. This enables us to establish the following proposition.

Proposition 6. *Consider the model with n agents within each generation outlined in this subsection.*

1. *Suppose that $m = n$. Then, following a prominent play of Low, in the greatest equilibrium there is a Low social norm and all endogenous agents play Low if and only if $\bar{\gamma}_L^n < \gamma$. Following a prominent play of High, there is a High social norm and all endogenous agents play High if and only if $\gamma \leq \underline{\gamma}_H^n$. Moreover, the threshold $\underline{\gamma}_H^n$ is increasing in n . If, in addition, $\underline{\gamma}_H^n \geq \bar{\gamma}_L^n$ (which is satisfied when (7) holds), the threshold $\bar{\gamma}_L^n$ is also nonincreasing in n , so that both High and Low social norms following, respectively, High and Low prominent play, emerge for a larger set of parameter values. The same result also holds (i.e. the threshold $\bar{\gamma}_L^n$ is nonincreasing in n) when $q = 0$ so that there are no prominent agents after the initial period.*
2. *Suppose that $m < n$. Then, in the greatest equilibrium following a prominent play of Low, there is a Low social norm and all endogenous agents play Low if $\bar{\gamma}_L^{n,m} < \gamma$. Following a prominent play of High, there is a High social norm and all endogenous agents play High if $\gamma \leq \underline{\gamma}_H^{n,m}$.*

This proposition covers two of the thresholds, but leaves out the other two $\underline{\gamma}_L^n$ and $\bar{\gamma}_H^n$ (or $\underline{\gamma}_L^{n,m}$ and $\bar{\gamma}_H^{n,m}$), which are defined as the equivalents of $\underline{\gamma}_L$ and $\bar{\gamma}_H$ in the baseline model. As before, these thresholds depend on beliefs and do not have direct closed-form solutions. A complication here is that an agent’s prediction of what agents of the future generation will do involves interpreting past signals. The signal that an agent sees from the previous generation is

a noisy indicator about what those in his or her own generation have seen, which then translates into an indicator of what other agents in the next generation are likely to see. The extreme cases where λ is 0 or 1 decouple this relationship, but more generally equilibrium cut-offs can depend on these expectations.

Results about breaking the *Low* social norm can also be extended to this model. With a similar analysis, an endogenous prominent agent can choose *High* to change the social norm of the society, and in addition to the factors making such a choice more attractive in our baseline model, it will also be more attractive when m is greater.

Another observation is noteworthy. In this model, even if a player has information about the previous generation containing a prominent agent, but not about what action this agent took, this could still lead to a switch from a *Low* to *High* social norm in the greatest equilibrium. The reasoning is as follows: players of this generation can believe that an endogenous prominent agent will choose *High* with a high probability, and thus are more likely to respond with *High*. If so, it makes sense for endogenous prominent agent to choose *High* to coordinate with the next generation, making these beliefs self-fulfilling.

5.3. More general payoff structures

Our analysis can be extended straightforwardly to more general payoff structures. For this discussion and for the rest of this section, let us return to the model with a single agent per generation. First, consider the following general symmetric two-by-two game, where we retain the same labels on strategies for convenience:

$$\begin{array}{cc} & \textit{High} & \textit{Low} \\ \textit{High} & b_{11}, b_{11} & b_{12}, b_{21} \\ \textit{Low} & b_{21}, b_{12} & b_{22}, b_{22} \end{array}$$

Since the subtraction of the payoff vector for one action from the other (for a given player) generates a new payoff matrix that is strategically equivalent (*i.e.* leaves the set of best responses to any strategy profile unchanged), it follows that the current payoff matrix is equivalent to our baseline with:²⁵

$$\alpha = b_{22} - b_{12} > 0 \text{ and } \beta = b_{11} - b_{21} > 0,$$

provided that these inequalities are indeed satisfied. Therefore, all of our results so far identically generalize to this case.

The above generalization does not cover other interesting cases, for example, prisoner dilemma-type payoffs. Nevertheless, our general analysis also applies to such cases. For such cases it is the fact that the next generation will observe a signal of current actions that creates incentives for cooperation (even among regular players). This highlights the role of forward-looking behaviour even more clearly. Suppose, more specifically, that the payoff matrix is

$$\begin{array}{cc} & \textit{High} & \textit{Low} \\ \textit{High} & \beta, \beta & -\alpha, \kappa \\ \textit{Low} & \kappa, -\alpha & 0, 0 \end{array}$$

where $\kappa > \beta > 0$ and $\alpha > 0$.²⁶ It can be verified that all *Low* is an equilibrium provided that signals are sufficiently precise. More importantly, using the same reasoning as above, it can also be

25. To obtain this, simply subtract the payoff vector for the action *Low* for one player, (b_{21}, b_{22}) , from the payoff vector for *High*, (b_{11}, b_{12}) .

26. This can also be written more generally for the class of strategically equivalent payoff matrices, but we omit this step to simplify the discussion.

seen that with sufficiently precise signals and λ sufficiently large, there is also an equilibrium in which *High* is played in response to good signals from the past.²⁷ This is intuitive: with sufficiently precise signals, an agent will be fairly sure that the previous generation has played *High* following a good realization of the signal, and can also expect the next generation to receive a good signal. In this situation, deviating to *Low* would generate a payoff gain from this interaction with the previous generation, but will make the next generation switch to *Low* with sufficiently high probability, which is costly. This is enough to deter *Low* when λ is sufficiently high. Note that beliefs about whether *High* is being played with a high probability (e.g. as a function of past history of prominent agents) will again have a defining effect on the interpretation of signals from the last generation and thus on the willingness of an agent to go along with the prevailing social norms and play *High*.

It is also straightforward to verify that when $\kappa - \beta \leq \alpha$, the resulting dynamic game of incomplete information is one of strategic complements. This ensures that results similar to those presented above apply in this case.

5.4. Imperfect prominence

A natural question is whether our results on history-driven behaviour hinge on perfect observation of prominent agents. To investigate this question, consider a variation where all future generations observe the same imperfect signal concerning the action of past prominent agents. In particular, suppose that they all receive a public signal $r_t \in \{Low, High\}$ (in addition to the private signal s_t from the non-prominent agent in the previous generation) concerning the action of the prominent agent of time t (if there is indeed a prominent agent at time t). We assume that $r_t = a_t$ with probability η , where $a_t \in \{Low, High\}$ is the action of the prominent agent. Clearly, as $\eta \rightarrow 1$, we converge to our baseline environment.

An important observation in this case is that the third part of Proposition 1 no longer applies and the greatest and least equilibria are not necessarily semi-Markovian. This is because, given imperfect signals about the actions of prominent agents, the play of previous prominent agents is relevant for beliefs about the play of the last prominent agent. Nevertheless, when η is sufficiently large but still strictly < 1 , the greatest equilibrium is again semi-Markovian and can be driven by history; that is, the common signal generated by the action of the last prominent agent.²⁸ A similar analysis also leads to the conclusion that when η is sufficiently large, all endogenous agents playing *Low* following a prominent public signal of *Low* is the greatest equilibrium whenever $\gamma > \bar{\gamma}_L$, and our results on leadership-driven changes in social norms also generalize to this set-up. Notably, for these conclusions, η needs to be greater than a certain threshold that is strictly < 1 , and thus history-driven behaviour emerges even with signals bounded away from being fully precise.

27. There does not, however, exist an equilibrium in which *High* is played in response to all signals for obvious reasons.

28. In particular, it can be shown that following a signal of $r = H$, the probability that a prominent agent has indeed played *High* cannot be lower than

$$\eta' \equiv \frac{\pi \eta}{\pi \eta + (1 - \pi)(1 - \eta)}.$$

This follows because there is always a probability π that the prominent agent in question was exogenously committed to *High*. For η close enough to 1, η' is strictly $> \Phi(0, 1 - \pi)$. In that case, whenever $\gamma \leq \underline{\gamma}_H$, where $\underline{\gamma}_H$ is given by (4), the reasoning that established Proposition 2 implies that, when the public signal from the last prominent agent indicates that he or she played *High*, the greatest equilibrium involves all endogenous agents playing *High* (regardless of their signal).

5.5. *Endogenous prominence*

In practice certain agents, such as Nelson Mandela, George Washington, or Mahatma Gandhi, are prominent not exogenously, but because of the remarkable acts and self-sacrifices (as vividly described by the self-immolation of Mohamed Bouazazi, who became prominent and sparked the Arab Spring).

To capture these issues in the simplest possible way, we can extend our baseline model as follows. With independent probability χ , an agent has the opportunity to incur a cost of $\tau > 0$ to become prominent (have his or her actions be seen by all future generations). When, in addition, with probability $q > 0$, each agent may be prominent directly by chance, Proposition 4 becomes more interesting: now there are two threshold signals from the past generation, above one of them a directly prominent agent breaks the social norm of *Low*, and above the second (higher) one a regular agent with an opportunity to invest in prominence will do so to break the social norm of *Low*. However, this modified game no longer exhibits strategic complementarities because the expectation that the next generation will invest in prominence can discourage prominence this period and thus delay a potential switch to *High*. Moreover, higher expectations of past play reduce the incentive of a current agent to invest to become prominent. Nonetheless, the underlying structure that we have analysed here, such as the reversion of the play of non-prominent agents and the interaction between beliefs and play, can provide an important foundation for the future study of endogenous prominence.

6. CONCLUSION

In this article, we studied the emergence and dynamic evolution of the social norm of “cooperation”. Social norms shape beliefs and behaviour, but rather than being completely locked in, they change over time in response to individual behaviour and actions by prominent agents or “leaders”.

Our main contribution is to provide a tractable model to study the dynamics of social norms and the role of leadership and prominence in shaping social norms. An important aspect of our framework is that, rather than being shifts between multiple equilibria, changes in social norms in our model change along a given equilibrium path. In fact, these norms can be completely (uniquely) determined by history, but still change over time. Thus, history is more than a simple correlating device in our framework: behaviour today can be uniquely determined by distant history that is irrelevant to current payoffs. This is because past events provide information about how other agents will interpret their information. In particular, beginning from even a distant history of more cooperative play, current signals are more favourably interpreted.

This setup underlies our interpretation of social norms as “frames of reference” that shape how information from the past is interpreted because agents only receive noisy information about past play. History—shared, common knowledge past events—anchors these social norms. Consider, for instance, the example of norms of corruption within the police force, which we used as a motivating example above. If the history of the actions of prominent police officers or corruption cases indicates that there is a social norm of high corruption, then even moderately favourable signals of current behaviours are interpreted as being due to noise and individual police officers will be unwilling to curtail their corruption. This role of social norms as frames of reference (in shaping the interpretation of past signals) leads to a form of history-driven norm of corruption: high corruption persists partly because, given the prevailing social norm, the signals the police officers generate even with honest actions are misinterpreted and cannot be trusted, thus discouraging honest behaviour.

The impact of history is potentially countered by “prominent” agents, who create the opportunity for future generations to coordinate. Then social norms are no longer necessarily everlasting, because prominent agents exogenously committed to one or the other mode of behaviour may arrive and cause a switch in play—and thus in the resulting social norm. More interestingly, prominent agents can also endogenously leverage their greater visibility and play a leadership role by coordinating the expectations of future generations. In the police corruption example, starting from a social norm of pervasive corruption, a prominent police officer or commissioner may choose a highly visible honest action to break the social norm of corruption.

We also showed that in equilibria that are not completely driven by history, there is a pattern of “reversion” whereby, for example, play starting with one behaviour reverts towards another. For instance, in the police example, play starting with a prominent honest behaviour may revert back to greater levels of corruption. The reason for this is interesting. Although an agent immediately following a prominent honest behaviour knows that he or she is playing against an honest action in the past, an agent two periods after such an honest behaviour does not, as there may have been an exogenous corrupt agent in the previous period. However, must take into account that there may have been an exogenous non-prominent corrupt agent in the previous period. Three periods after a prominent honest behaviour, the likelihood of an intervening exogenous non-prominent agent switching to corruption is even greater, which then feeds back into the optimal responses. Indeed, these agents anticipate that even endogenous non-prominent agents now may switch to corrupt behaviour because they are unsure of with whom they are playing in the previous generation and an adverse signal makes them believe that they are playing against a corrupt agent. Secondly, they understand as well that even if they were honest, the signals generated may still be interpreted as if they were coming from corrupt practices, further discouraging honest behaviour.

There are several promising areas of future work based on our approach, and more generally based on the interplay between history, social norms, and interpretation of past actions. First, it would be useful to extend the analysis of the role of history, expectations and leadership to a more detailed, and empirically grounded, model of collective action, in which individuals care about how many people, from the past and future generations, will take part in some collective action, such as an uprising or demonstration against a regime.

Secondly, in some situations non-prominent agents in our model have an incentive to communicate their behaviour, since by doing so they can avoid the need to rely on social norms for forming accurate expectations of past and future play. It is also possible that a society might have asymmetric incentives to communicate history to future generations, for instance, perhaps erasing evidence of past prominent *Low* play and reporting past prominent *High* play. This has interesting consequences, making observed (perceived) histories less trustworthy. Another related direction is to study the evolution of social norms in situations where incentives are not fully aligned, so that communication does not fully circumvent the role of social norms in coordinating expectations.

Finally, it would be interesting to introduce an explicit network structure in the pattern of observation and interaction so that agents who occupy a central position in the social network—whose actions are thus known to be more likely to be observed by many others in the future—(endogenously) play the role of prominent agents in our baseline model. This will help us get closer to understanding which types of agents, and under which circumstances, can play a leadership role.

APPENDIX

A.1. Equilibrium Definition

Our definition of equilibrium is standard and requires that agents best respond to their beliefs conditional on any history and signal and given the strategies of others.²⁹ The only thing that we need to be careful about is defining those beliefs. In cases where $0 < q < 1$ and $\pi > 0$ those beliefs are easily derived from Bayes' rule (and an appropriate iterative application of equation (2)). We provide a careful definition that also allows for $q=0$ or $\pi=0$ even though in the text we have assumed $q > 0$ and $\pi > 0$. In these corner cases, some additional care is necessary since some histories off the equilibrium path may not be reached.³⁰

Consider any $t \geq 1$, any history h^{t-1} , and a strategy profile σ .

Let $\phi_{t+1}^t(\sigma_{t+1}, T_t, h^{t-1})$ be the probability that, given strategy σ_{t+1} , the next agent will play *High* if agent t plays *High* and is of prominence type $T_t \in \{P, N\}$. Note that this is well defined and is independent of the signal that agent t observes.

Let $\phi_{t-1}^t(\sigma, s_t, h^{t-1})$ denote the probability that agent t assigns to the previous agent playing *High* given signal s_t , strategy profile σ , and history h^{t-1} . In particular, if $h_{t-1} = \text{High}$ then set $\phi_{t-1}^t(\sigma, s_t, h^{t-1}) = 1$ and if $h_{t-1} = \text{Low}$ then set $\phi_{t-1}^t(\sigma, s_t, h^{t-1}) = 0$. If $h_{t-1} = N$ then define $\phi_{t-1}^t(\sigma, s_t, h^{t-1})$ via an iterative application Bayes' rule. Specifically, this is done via an application of (2) as follows. Let τ be the largest element of $\{1, \dots, t-1\}$ such that $h_\tau \neq N$ (i.e. the date of the last prominent agent). Then given $\sigma_{\tau+1}(h^\tau, N, s_{\tau+1})$ and π , there is an induced distribution on *High* and *Low* by generation $\tau+1$ and thus over $s_{\tau+2}$ (and note that $s_{\tau+1}$ is irrelevant since τ is prominent). Then given $\sigma_{\tau+2}(h^\tau, N, s_{\tau+2})$ and π , there is an induced distribution on *High* and *Low* by generation $\tau+2$, and so forth. By induction, there is an induced distribution on *High* and *Low* at time $t-1$, which we then denote by x_{t-1} . Then $\phi_{t-1}^t(\sigma, s_t, h^{t-1}) = \Phi(s_t, x_{t-1})$ where Φ is defined in (2).

From (3), it is a best response for agent t to play *High* if

$$(1-\lambda)\phi_{t-1}^t(\sigma, s_t, h^{t-1}) + \lambda\phi_{t+1}^t(\sigma_{t+1}, T_t, h^{t-1}) > \gamma, \quad (\text{A.1})$$

to play *Low* if

$$(1-\lambda)\phi_{t-1}^t(\sigma, s_t, h^{t-1}) + \lambda\phi_{t+1}^t(\sigma_{t+1}, T_t, h^{t-1}) < \gamma, \quad (\text{A.2})$$

and either if there is equality.

We say that σ forms an *equilibrium* if for each time $t \geq 1$, history $h^{t-1} \in \mathcal{H}^{t-1}$, signal $s_t \in [0, 1]$, and type $T_t \in \{P, N\}$ $\sigma_t(h^{t-1}, s_t, T_t) = 1$ if (A.1) holds and $\sigma_t(h^{t-1}, s_t, T_t) = 0$ if (A.2) holds, where $\phi_{t-1}^t(\sigma, s_t, h^{t-1})$ and $\phi_{t+1}^t(\sigma_{t+1}, T_t, h^{t-1})$ are as defined above.

A.2. Equilibria in games with strategic complementarities and infinitely many agents

We now establish a theorem that will be used in proving Proposition 1. This theorem is also of potential independent interest for this class of overlapping-generation incomplete information games.

Well-known results for games of strategic complements apply to finite numbers of agents (e.g. Topkis, 1979; Vives, 1990; Milgrom and Shannon, 1994; Zhou, 1994; and Van Zandt and Vives, 2007). The next theorem provides an extension for arbitrary sets of agents, including countably and uncountably infinite sets of agents.

Let us say that a game is a game of *weak strategic complements with a possibly infinite number of agents* if the agents are indexed by $i \in I$ and:

- each agent has an action space A_i that is a complete lattice with a partial ordering \geq_i and corresponding sup _{i} and inf _{i} ;
- for every agent i , and specification of strategies of the other agents, $a_{-i} \in \prod_{j \neq i, j \in I} A_j$, agent i has a non-empty set of best responses $BR_i(a_{-i})$ that is a closed sublattice of A_i (where "closed" here is in the lattice-sense, so that $\sup(BR_i(a_{-i})) \in BR_i(a_{-i})$ and $\inf(BR_i(a_{-i})) \in BR_i(a_{-i})$);

29. Definitions for perfect Bayesian equilibrium and sequential equilibrium are messy when working with continua of private signals, and so it is easiest to provide a direct definition of equilibrium here, which is relatively straightforward.

30. These beliefs can still be consequential. To see an example of why this matters in our context, consider a case where all agents are endogenous and prominent (so $\pi=0$ and $q=1$, which is effectively a complete information game). Let an agent be indifferent between *High* and *Low* if both surrounding generations play *Low*, but otherwise strictly prefer *High*. Begin with agent 0 playing *Low*. There is a (Bayesian) Nash equilibrium where all agents play *Low* regardless of what others do, but it is not perfect (Bayesian). This leads to different minimal equilibria depending on whether one works with Bayesian or perfect Bayesian equilibrium.

- for every agent i , if $a'_j \geq_j a_j$ for all $j \neq i, j \in I$, then $\sup_i BR_i(a'_{-i}) \geq_i \sup_i BR_i(a_{-i})$ and $\inf_i BR_i(a'_{-i}) \geq_i \inf_i BR_i(a_{-i})$.

For the next theorem, define $\mathbf{a} \geq \mathbf{a}'$ if and only if $a_i \geq_i a'_i$ for all i . The lattice of equilibria on $A = \prod_{i \in I} A_i$ can then be defined with respect to this partial ordering.³¹

Theorem 1. *Consider a game of weak strategic complements with a possibly infinite number of agents. An equilibrium exists, and the set of equilibria form a complete lattice.*³²

Proof of Theorem 1. Let $A = \prod_{i \in I} A_i$. Note that A is a complete lattice, where we say that $\mathbf{a} \geq \mathbf{a}'$ if and only if $a_i \geq_i a'_i$ for every $i \in I$, and where for any $S \subset A$ we define

$$\sup(S) = (\sup_i \{a_i : \mathbf{a} \in S\})_{i \in I}, \text{ and}$$

$$\inf(S) = (\inf_i \{a_i : \mathbf{a} \in S\})_{i \in I}.$$

Given the lattice A , we define the best response correspondence $f : A \rightarrow 2^A$ by

$$f(\mathbf{a}) = (BR_i\{\mathbf{a}_{-i}\})_{i \in I}$$

By the definition of a game of strategic complements, $BR_i(a_{-i})$ is a non-empty closed sublattice of A_i for each i and a_{-i} , and so it follows directly that $f(a)$ is a non-empty closed sublattice of A for every $a \in A$. Note that by the strategic complementarities f is monotone: if $\mathbf{a} \geq \mathbf{a}'$ then $\sup(f(\mathbf{a})) \geq \sup(f(\mathbf{a}'))$ and $\inf(f(\mathbf{a})) \geq \inf(f(\mathbf{a}'))$. This follows directly from the fact that if $a'_{-i} \geq a_{-i}$, then $\sup BR_i(a'_{-i}) \geq_i \sup BR_i(a_{-i})$ (and $\inf BR_i(a'_{-i}) \geq_i \inf BR_i(a_{-i})$) for each i .

Thus, by an extension of Tarski's (1955) fixed-point theorem due to Straccia *et al.* (2009) (see also Zhou, 1994),³³ f has a fixed point and its fixed points form a complete lattice (with respect to \geq). Note that a fixed point of f is necessarily a best response to itself, and so is an equilibrium, and all equilibria are fixed points of f , and so the equilibria are exactly the fixed points of f . \parallel

A.3. Proofs of Propositions 1-3

Proof of Proposition 1. Part 1. The result follows by showing that for any strategy profile there exists a best response that is in cut-off strategies. To see this, recall from (3) that *High* is a best response if and only if

$$(1 - \lambda)\phi'_{t-1} + \lambda\phi'_{t+1} \geq \gamma, \tag{A.3}$$

and is a unique best response if the inequality is strict. Clearly, $\phi'_{t-1}(\sigma, s, h^{t-1})$ (as defined in our definition of equilibrium) is increasing in s under the MLRP (and given that $\pi > 0$) in any period not following a prominent agent. Moreover, ϕ'_{t+1} is independent of the signal received by the agent of generation t . Thus, if an agent follows a non-prominent agent, the best responses are in cut-off strategies and are unique except for a signal that leads to exact indifference, that is (A.3) holding exactly as equality, in which case any mixture is a best response. An agent following a prominent agent has $\phi'_{t-1}(\sigma, s, h^{t-1})$ as either 0 or 1, and thus can set any cutoff that leads to a probability of *Low* corresponding to a best response - usually 0 or 1 except when indifferent when any cutoff works. This completes the proof of Part 1.

Also, for future reference, we note that in both cases the set of best responses are closed (either 0 or 1, or any mixture thereof).

Part 2. The result that there exists a semi-Markovian equilibrium in cut-off strategies follows from the proof of Part 3, where we show that the set of equilibria in cut-off strategies and semi-Markovian equilibria in cut-off strategies are non-empty and complete lattices.

Part 3. This part of the proof uses Theorem 1 applied to cut-off and semi-Markovian cut-off strategies to show that the sets of these equilibria are non-empty and complete lattices. We will then show that greatest and least equilibria are semi-Markovian. We thus first need to show that our game is one of weak strategic complements. We start with the following intermediate result.

Claim 1. The sets of cut-off and semi-Markovian cut-off strategies for a given player are complete lattices.

31. Note, however, that the set of equilibria is not necessarily a sublattice of A , as pointed out in Topkis (1979) and in Zhou (1994) for the finite case. That is, the sup in A of a set of equilibria may not be an equilibrium, and so sup and inf have to be appropriately defined over the set of equilibria to ensure that the set is a complete lattice. Nevertheless, the same partial ordering can be used to define the greatest and least equilibria.

32. These are pure strategy equilibria whenever the action spaces are pure actions.

33. The monotonicity of f here implies the EM-monotonicity in Proposition 3.15 of Straccia *et al.* (2009).

Proof. The cut-off strategies of a each prominence type of a player of generation t can be written as a vector in $[0, 1]^{3^t}$, where this vector specifies a cut-off for every possible history of the occurrence and play of prominent agents (and there are 3^t of them, including time $t=0$). This is a complete lattice with the usual Euclidean partial order. Semi-Markovian cut-off strategies, however, can be simply written as a pair of cut offs (depending on the current player's prominence type and as a function of the number of periods τ since the last prominent agent and what that prominent play was). \parallel

Next, we verify the strategic complementarities for cut-off strategies. Let $z_{t-1}(\sigma, h^{t-1})$ be the prior probability that this agent assigns to an agent of the previous period playing *High* conditional on h^{t-1} (and before observing s). Fix a cut-off strategy profile $c = (c_1^N(h^0), c_1^P(h^0), \dots, c_t^N(h^{t-1}), c_t^P(h^{t-1}), \dots)$ for every possible history. Suppose that $\sup BR_t^T(c)$ is the greatest best response of agent of generation t of prominence type T to the cut-off strategy profile c (meaning that it is the best response with the lowest cut-offs). Now consider:

$\tilde{c} = (\tilde{c}_1^N(h^0), \tilde{c}_1^P(h^0), \dots, \tilde{c}_t^N(h^{t-1}), \tilde{c}_t^P(h^{t-1}), \dots) \leq c = (c_1^N(h^0), c_1^P(h^0), \dots, c_t^N(h^{t-1}), c_t^P(h^{t-1}), \dots)$ for each infinite history. We show that $\sup BR_t^T(c) \geq \sup BR_t^T(\tilde{c})$ (the argument for $\inf BR_t^T(c) \geq \inf BR_t^T(\tilde{c})$ is analogous) and we can argue for any given history. First, cut-offs after $t+2$ do not affect $BR_t^T(c)$. Secondly, suppose that all cut-offs before $t-1$ remain fixed and c_{t+1}^N and c_{t+1}^P decrease (meaning that they are weakly lower for every history and at least one of them is strictly lower for at least one history). This increases $\phi_{t+1}^t(\sigma, T, h^{t-1})$ and thus makes (A.3) more likely to hold, so $\sup BR_t^T(c) \geq \sup BR_t^T(\tilde{c})$. Thirdly, suppose that all cut-offs before $t-2$ remain fixed, and c_{t-1}^N and c_{t-1}^P decrease. This increases $z_{t-1}(\sigma, h^{t-1})$ (or does not change beliefs if the $t-1$ agent was prominent) and thus $\phi_{t-1}^t(\sigma, s, h^{t-1})$ and thus makes (A.3) more likely to hold, so again $\sup BR_t^T(c) \geq \sup BR_t^T(\tilde{c})$. Fourthly, suppose that all other cut-offs remained fixed and c_{t-k-1}^N and c_{t-k-1}^P (for $k \geq 1$) decrease. By MLRP, this shifts the distribution of signals at time $t-k$ in the sense of first-order stochastic dominance and thus given c_{t-k}^N and c_{t-k}^P , it weakly increases $z_{t-k}(\sigma, h^{t-k-1})$, shifting the distribution of signals at time $t-k+1$ in the sense of first-order static dominance. Applying this argument iteratively k times (presuming no prominent dates in between, as otherwise it does not affect subsequent beliefs conditional on this history), we conclude that $\sup BR_t^T(c) \geq \sup BR_t^T(\tilde{c})$. This establishes that whenever $c \geq \tilde{c}$, $\sup BR_t^T(c) \geq \sup BR_t^T(\tilde{c})$. The same argument also applies to semi-Markovian cut-offs. Thus, from Theorem 1 the set of equilibria in cut-off strategies and set of semi-Markovian equilibria in cut-off strategies are non-empty complete lattices.

To complete the proof, we next show that greatest and least equilibria are semi-Markovian. We provide the argument for the greatest equilibrium and the argument for the least is analogous. It is clear that the overall greatest equilibrium is at least as high (with cut-offs at least as low) as the greatest semi-Markov equilibrium since it includes such equilibria, so it is sufficient to show that the greatest equilibrium is semi-Markovian. Thus, suppose to the contrary of the claim that the greatest equilibrium, say $c = (c_1^N(h^0), c_1^P(h^0), \dots, c_t^N(h^{t-1}), c_t^P(h^{t-1}), \dots)$, is not semi-Markovian. This implies that there exists some t (and $T \in \{P, N\}$) such that $c_t^T(h^{t-1}) > c_t^T(\tilde{h}^{t-1})$ where h^{t-1} and \tilde{h}^{t-1} have the same last prominent agent, say occurring at time $t-k$. Then consider:

$\tilde{c} = (c_1^N(h^0), c_1^P(h^0), \dots, c_{t-k+1}^N(h^{t-k}), c_{t-k+1}^P(h^{t-k}), \tilde{c}_{t-k+2}^N(h^{t-k+1}), \tilde{c}_{t-k+2}^P(h^{t-k+1}), \dots, \tilde{c}_t^N(h^{t-1}), \tilde{c}_t^P(h^{t-1}), c_{t+1}^N(h^t), c_{t+1}^P(h^t), \dots)$, where $\tilde{c}_{t-k+j+1}^T(h^{t-k+j}) = \min\{c_{t-k+j+1}^T(\tilde{h}^{t-k+j}), c_{t-k+j+1}^T(h^{t-k+j})\}$ with \tilde{h}^{t-k+j} and h^{t-k+j} are the truncated versions of histories \tilde{h}^{t-1} and h^{t-1} . Next, it is straightforward to see that \tilde{c} is also an equilibrium. In particular, following history \tilde{h}^{t-1} , c is an equilibrium by hypothesis. Since the payoffs of none of the players after $t-k$ directly depend on the action of the prominent agents before the last one, this implies that when all agents after $t-k$ switch their cut-offs after history h^{t-k} as in \tilde{c} , this is still an equilibrium. This shows that \tilde{c} is an equilibrium cut-off profile, but this contradicts that c is the greatest equilibrium. \parallel

Proof of Proposition 2.

Part 1. First, note that if the greatest equilibrium is for all endogenous agents to play *High* following a prominent play of *High* for some γ , then it is also the greatest equilibrium for all lower γ . This follows from the monotonicity of best responses in γ . Thus, the set of γ 's for which all endogenous agents playing *High* following the last prominent play being *High* is the greatest equilibrium is an interval.

By the argument preceding (4), the cut-off γ for all endogenous agents to play *High* following a prominent play of *High* is $\underline{\gamma}_H$ as defined in (4). It also follows that this is a closed interval, since it is an equilibrium for $\gamma = \underline{\gamma}_H$, but not for any higher γ since then *High* is not a best response conditional upon the lowest signal.

Thus, all endogenous agents playing *High* following the last prominent play being *High* is the greatest equilibrium if and only if $\gamma \in [0, \underline{\gamma}_H]$.

Consider next the set of γ 's for which all endogenous agents playing *Low* following the last prominent play being *High* is the greatest equilibrium. Again from (3), if γ belongs to this set, then any $\gamma' > \gamma$ also does. Let the cut-off γ be denoted $\overline{\gamma}_H$, so all playing *Low* is the greatest equilibrium if $\gamma > \overline{\gamma}_H$, but not if $\gamma < \overline{\gamma}_H$.

To complete the proof of Part 1, we need to show that the remaining interval $(\underline{\gamma}_H, \overline{\gamma}_H]$ is such that endogenous agents begin by playing *High* and then eventually play some *Low*. The fact that the initial play when $\gamma \in (\underline{\gamma}_H, \overline{\gamma}_H)$ must be *High*

follows from the proof of Proposition 3.³⁴ The fact that the intermediate interval must also involve some play of *Low* then follows from the proof that all play of *High* is only an equilibrium if $\gamma \leq \underline{\gamma}_H$.

Part 2. Proceeding similarly to Part 1, the set of γ 's such that in the greatest equilibrium all endogenous agents play *High* following a prominent *Low* is an interval of the form $[0, \underline{\gamma}_L]$, where in this case $\underline{\gamma}_L = \lambda(1 - \pi)$, since otherwise (3) would not be satisfied for an agent immediately following a prominent *Low*.

Now take $\gamma > \underline{\gamma}_L \equiv \lambda(1 - \pi)$. This is sufficient for *Low* to be a strict best response immediately following a prominent *Low*. But the next agent does not know for sure that the previous generation played *Low*. If $\gamma > \underline{\gamma}_L$, then he or she expects his or her previous generation agent to have played *Low* unless he or she was exogenously committed to *High*. This implies that it is sufficient to consider the expectation of ϕ_{t-1}^* under this assumption and ensure that even for the signal most favourable to the previous generation agent having played *High*, *Low* is a best response. The threshold for this is

$$\gamma_L^* \equiv (1 - \lambda)\Phi(1, \pi) + \lambda(1 - \pi). \tag{A.4}$$

Thus if $\gamma > \gamma_L^* > \underline{\gamma}_L$, this agent will also have a *Low* strict best response even in the greatest equilibrium. Now we proceed inductively and conclude that this threshold applies to all future agents. Thus, when $\gamma > \gamma_L^*$, all endogenous agents following a prominent *Low* will play *Low*.

Next consider the threshold $\bar{\gamma}_L$ for which, in the greatest equilibrium, all endogenous agents play *Low* following a prominent play of *Low*. This threshold satisfies $\bar{\gamma}_L \leq \gamma_L^*$. Once again with a similar argument as in Part 1, the set of γ 's for which the greatest equilibrium involves all endogenous agents playing *Low* following a prominent *Low* is an interval. An analogous argument as in Part 1 for the remaining interval concludes the proof of Part 2.

In addition, we show that if $\bar{\gamma}_L \leq \underline{\gamma}_H$ (and thus *a fortiori* if $\gamma_L^* \leq \underline{\gamma}_H$), then $\bar{\gamma}_L = \gamma_L^*$, thus establishing (6). We next prove this result.

Proof that $\bar{\gamma}_L \leq \underline{\gamma}_H$ implies $\bar{\gamma}_L = \gamma_L^*$. Suppose $\bar{\gamma}_L \leq \underline{\gamma}_H$ and consider the case where $\gamma = \bar{\gamma}_L$. Then following a *High* play of a prominent agent, all endogenous agents will play *High*. Therefore, for an endogenous prominent agent to have *Low* as best response for any signal and prior x , it has to be the case that $(1 - \lambda)\Phi(1, x) + \lambda(1 - \pi) \leq \bar{\gamma}_L$. Since $\bar{\gamma}_L \leq \gamma_L^*$, this implies

$$(1 - \lambda)\Phi(1, x) + \lambda(1 - \pi) \leq \bar{\gamma}_L \leq (1 - \lambda)\Phi(1, \pi) + \lambda(1 - \pi).$$

Therefore, $\Phi(1, x) \leq \Phi(1, \pi)$, or equivalently $x = \pi$ as π is the lowest possible prior of previous agent playing *High*. Hence $\bar{\gamma}_L = \gamma_L^*$. This result also implies that when $\gamma_L^* \leq \underline{\gamma}_H$, we also have $\bar{\gamma}_L = \gamma_L^*$. Then equation (7) is obtained by comparing the expressions for $\underline{\gamma}_H$ and γ_L^* . \square

Proof of Proposition 3. We prove Parts 2 and 3, and the proof involves proving Part 1.

Part 2. Consider play following a prominent *High*, and consider strategies listed as a sequence of cut-off thresholds $\{(c_\tau^P, c_\tau^N)\}_{\tau=1}^\infty$ for prominent and non-prominent players as a function of the number of periods τ since the last prominent agent. We first show by contradiction that $\{(c_\tau^P, c_\tau^N)\}_{\tau=1}^\infty$ must be non-decreasing. To do this, let us define a new sequence $\{(C_\tau^P, C_\tau^N)\}_{\tau=1}^\infty$ as follows:

$$C_\tau^T = \min\{c_\tau^T, c_{\tau+1}^T\}$$

for $T \in \{P, N\}$. The sequences $\{(c_\tau^P, c_\tau^N)\}_{\tau=1}^\infty$ and $\{(C_\tau^P, C_\tau^N)\}_{\tau=1}^\infty$ coincide if and only if $\{(c_\tau^P, c_\tau^N)\}_{\tau=1}^\infty$ is non-decreasing. Moreover, since $C_\tau^T \leq c_\tau^T$, if this is not the case, then there exist some τ, T such that $C_\tau^T < c_\tau^T$.

Suppose, to obtain a contradiction, that there exist some τ, T such that $C_\tau^T < c_\tau^T$ (and for the rest of the proof fix $T \in \{P, N\}$ to be this type). Define $B(\mathbf{C})$ be the lowest best response cut-off (for each τ, T) to the sequence of strategies \mathbf{C} . Since we have a game of weak strategic complements as established in the proof of Proposition 1, B is a non-decreasing function. We will first show that $B(\mathbf{C})_\tau^T \leq C_\tau^T$ for all τ and T (or that $B(\mathbf{C}) \leq \mathbf{C}$), and then we will show that in this case there exists an equilibrium $\mathbf{C}' \leq \mathbf{C}$. This will finally yield a contradiction since either $\mathbf{C} = \mathbf{c}$ or $\mathbf{C} \neq \mathbf{c}$, in which case \mathbf{c} is not greater than \mathbf{C}' , contradicting the fact that \mathbf{c} is the greatest equilibrium.

Let $\phi_{\tau-1}^T(\mathbf{C}, s_i)$ and $\phi_{\tau+1}^T(\mathbf{C}, s_i)$ denote the beliefs under \mathbf{C} of the last and next period agents, respectively, playing *High* if the agent of generation τ plays *High* conditional upon seeing signal s_i . Similarly, let $\phi_{\tau-1}^T(\mathbf{c}, s_i)$ and $\phi_{\tau+1}^T(\mathbf{c}, s_i)$ denote the corresponding beliefs under \mathbf{c} . If $C_\tau^T = c_\tau^T$, then since $\mathbf{C} \leq \mathbf{c}$ it follows that $\phi_{\tau-1}^T(\mathbf{C}) \geq \phi_{\tau-1}^T(\mathbf{c})$ and $\phi_{\tau+1}^T(\mathbf{C}) \geq \phi_{\tau+1}^T(\mathbf{c})$. This implies from (3) that

$$B(\mathbf{C})_\tau^T \leq B(\mathbf{c})_\tau^T = c_\tau^T = C_\tau^T,$$

where the second relation follows from the fact that \mathbf{c} is the cut-off associated with the greatest equilibrium. Thus, $B(\mathbf{c}) = \mathbf{c}$.

Hence, consider the case where $C_\tau^T = c_{\tau+1}^T < c_\tau^T$. We now show that also in this case $\phi_{\tau-1}^T(\mathbf{C}, s_i) \geq \phi_{\tau-1}^T(\mathbf{c}, s_i)$ and $\phi_{\tau+1}^T(\mathbf{C}, s_i) \geq \phi_{\tau+1}^T(\mathbf{c}, s_i)$. First, $\phi_{\tau+1}^T(\mathbf{C}) \geq \phi_{\tau+2}^T(\mathbf{c})$ follows directly from the fact that $C_{\tau+1}^T \leq c_{\tau+2}^T$. Next to establish that

34. The proof of Proposition 3 references this proposition, but there is no circularity since the reference is only to the result a play of all *High* is an equilibrium following a prominent *High* only in the case in which $\gamma \leq \underline{\gamma}_H$ — a result we have already established.

$\phi_{\tau-1}^\tau(\mathbf{C}) \geq \phi_\tau^{\tau+1}(\mathbf{c})$, it is sufficient to show that the prior probability of *High* at time $\tau-1$ under \mathbf{C} , $P_{\mathbf{C}}(a_{\tau-1} = \text{High})$, is no smaller than the prior probability of *High* at time τ under \mathbf{c} , $P_{\mathbf{c}}(a_\tau = \text{High})$. We next establish this:

Claim 2. $P_{\mathbf{C}}(a_{\tau-1} = \text{High}) \geq P_{\mathbf{c}}(a_\tau = \text{High})$.

Proof. We prove this inequality by induction. It is clearly true for $\tau=1$ (since we start with a prominent *High*). Next suppose it holds for $t < \tau$, and we show that it holds for τ . Note that

$$\begin{aligned} P_{\mathbf{C}}(a_{\tau-1} = \text{High}) &= (1 - F_H(C_{\tau-1}^N))P_{\mathbf{C}}(a_{\tau-2} = \text{High}) + (1 - F_L(C_{\tau-1}^N))(1 - P_{\mathbf{C}}(a_{\tau-2} = \text{High})), \\ P_{\mathbf{c}}(a_\tau = \text{High}) &= (1 - F_H(c_\tau^N))P_{\mathbf{c}}(a_{\tau-1} = \text{High}) + (1 - F_L(c_\tau^N))(1 - P_{\mathbf{c}}(a_{\tau-1} = \text{High})) \end{aligned}$$

Then we need to check that

$$\begin{aligned} (1 - F_H(C_{\tau-1}^N))P_{\mathbf{C}}(a_{\tau-2} = \text{High}) + (1 - F_L(C_{\tau-1}^N))(1 - P_{\mathbf{C}}(a_{\tau-2} = \text{High})) \\ \geq (1 - F_H(c_\tau^N))P_{\mathbf{c}}(a_{\tau-1} = \text{High}) + (1 - F_L(c_\tau^N))(1 - P_{\mathbf{c}}(a_{\tau-1} = \text{High})). \end{aligned}$$

By definition $C_{\tau-1}^N \leq c_\tau^N$, and therefore $1 - F_H(C_{\tau-1}^N) \geq 1 - F_H(c_\tau^N)$ and $1 - F_L(C_{\tau-1}^N) \geq 1 - F_L(c_\tau^N)$, so the following is a sufficient condition for the desired inequality:

$$\begin{aligned} (1 - F_H(c_\tau^N))P_{\mathbf{C}}(a_{\tau-2} = \text{High}) + (1 - F_L(c_\tau^N))(1 - P_{\mathbf{C}}(a_{\tau-2} = \text{High})) \\ \geq (1 - F_H(c_\tau^N))P_{\mathbf{c}}(a_{\tau-1} = \text{High}) + (1 - F_L(c_\tau^N))(1 - P_{\mathbf{c}}(a_{\tau-1} = \text{High})). \end{aligned}$$

This in turn is equivalent to

$$(1 - F_H(c_\tau^N))[P_{\mathbf{C}}(a_{\tau-2} = \text{High}) - P_{\mathbf{c}}(a_{\tau-1} = \text{High})] \geq (1 - F_L(c_\tau^N))[P_{\mathbf{C}}(a_{\tau-2} = \text{High}) - P_{\mathbf{c}}(a_{\tau-1} = \text{High})].$$

Since $P_{\mathbf{C}}(a_{\tau-2} = \text{High}) - P_{\mathbf{c}}(a_{\tau-1} = \text{High}) \geq 0$ by the induction hypothesis and $F_H(c_\tau^N) \leq F_L(c_\tau^N)$, this inequality is always satisfied, establishing the claim. \parallel

This claim thus implies that $\phi_{\tau-1}^\tau(\mathbf{C}) \geq \phi_\tau^{\tau+1}(\mathbf{c})$. Together with $\phi_{\tau+1}^\tau(\mathbf{C}) \geq \phi_{\tau+2}^{\tau+1}(\mathbf{c})$, which we established above, this implies that $B(\mathbf{C})_\tau^T \leq B(\mathbf{c})_{\tau+1}^T$. Then

$$B(\mathbf{C})_\tau^T \leq B(\mathbf{c})_{\tau+1}^T = c_{\tau+1}^T = C_\tau^T,$$

where the second relationship again follows from the fact that \mathbf{c} is an equilibrium and the third one from the hypothesis that $C_\tau^T = c_{\tau+1}^T < c_\tau^T$. This result completes the proof that $B(\mathbf{C}) \leq \mathbf{C}$. We next prove the existence of an equilibrium $\mathbf{C}' \leq \mathbf{C}$, which will finally enable us to establish the desired contradiction.

Claim 3. There exists an equilibrium \mathbf{C}' such that $\mathbf{C}' \leq \mathbf{C} \leq \mathbf{c}$.

Proof. Consider the (complete) sublattice of points $\mathbf{C}' \leq \mathbf{C}$. Since B is a non-decreasing function and takes all points of the sublattice into the sublattice (*i.e.* since $B(\mathbf{C}) \leq \mathbf{C}$), Tarski's (1955) fixed-point theorem implies that B has a fixed point $\mathbf{C}' \leq \mathbf{C}$, which is, by construction, an equilibrium. \parallel

Now the desired contradiction is obtained by noting that if $\mathbf{C} \neq \mathbf{c}$, then \mathbf{c} is not greater than \mathbf{C}' , contradicting the fact that \mathbf{c} is the greatest equilibrium. This contradiction establishes that $\mathbf{C} = \mathbf{c}$, and thus that $\{(c_\tau^P, c_\tau^N)\}_{\tau=1}^\infty$ is non-decreasing.

We next show that $\{(c_\tau^P, c_\tau^N)\}_{\tau=1}^\infty$ is increasing when $\gamma > \underline{\gamma}_H$. Choose the smallest τ such that $c_\tau^N > 0$. This exists from Proposition 2 in view of the fact that $\gamma > \underline{\gamma}_H$. By definition, an endogenous agent in generation $\tau-1$ played *High*, whereas the agent in generation $\tau+1$ knows, again by construction, that the previous generation will choose *Low* for some signals. This implies that $\phi_{\tau-1}^\tau > \phi_\tau^{\tau+1}$, and moreover, $\phi_{\tau+1}^\tau \geq \phi_{\tau+2}^{\tau+1}$ from the fact that the sequence $\{(c_\tau^P, c_\tau^N)\}_{\tau=1}^\infty$ is non-decreasing. This implies that $(c_{\tau+1}^P, c_{\tau+1}^N) > (c_\tau^P, c_\tau^N)$ (provided the latter is not already $(1,1)$). Now repeating this argument for $\tau+1, \dots$, the result that $\{(c_\tau^P, c_\tau^N)\}_{\tau=1}^\infty$ is increasing (for $\gamma > \underline{\gamma}_H$) is established, completing the proof of Part 2.

Part 3. In this case, we need to show that the sequence $\{(c_\tau^P, c_\tau^N)\}_{\tau=1}^\infty$ is non-increasing starting from a prominent agent choosing *Low*. The proof is analogous, except that we now define the sequence $\{(C_\tau^P, C_\tau^N)\}_{\tau=1}^\infty$ with

$$C_\tau^T = \min\{c_{\tau-1}^T, c_\tau^T\}.$$

Thus in this case, it follows that $\mathbf{C} \leq \mathbf{c}$, and the two sequences coincide if and only if $\{(c_\tau^P, c_\tau^N)\}_{\tau=1}^\infty$ is non-increasing. We define $B(\mathbf{C})$ analogously. The proof that $B(\mathbf{C}) \leq \mathbf{C}$ is also analogous. In particular, when $C_\tau^T = c_\tau^T$, the same argument establishes that

$$B(\mathbf{C})_\tau^T \leq B(\mathbf{c})_\tau^T = c_\tau^T = C_\tau^T.$$

Hence, consider the case where $C_\tau^T = c_{\tau-1}^T < c_\tau^T$. Then the same argument as above implies that $\phi_{\tau+1}^\tau(\mathbf{C}) \geq \phi_{\tau-1}^{\tau-1}(\mathbf{c})$. Next, we can also show that $\phi_{\tau-1}^\tau(\mathbf{C}) \geq \phi_{\tau-2}^{\tau-1}(\mathbf{c})$ by establishing the analogue of Claim 2.

Claim 4. $P_{\mathbf{C}}(a_\tau = \text{High}) \geq P_{\mathbf{c}}(a_{\tau-1} = \text{High})$.

Proof. The proof is analogous to that of Claim 2 and is again by induction. The base step of the induction is true in view of the fact that we now start with a *Low* prominent agent. When it is true for $t < \tau$, a condition sufficient for it to be also true for τ can again be written as

$$(1 - F_H(c_{\tau-1}^N))[P_C(a_{\tau-1} = \text{High}) - P_c(a_{\tau-2} = \text{High})] \geq (1 - F_L(c_{\tau-1}^N))[P_C(a_{\tau-1} = \text{High}) - P_c(a_{\tau-2} = \text{High})].$$

Since $P_C(a_{\tau-1} = \text{High}) - P_c(a_{\tau-2} = \text{High}) \geq 0$ and $F_H(c_{\tau-1}^N) \leq F_L(c_{\tau-1}^N)$, this inequality is satisfied, establishing the claim. \parallel

This result now implies the desired relationship

$$B(\mathbf{C})_\tau^T \leq B(\mathbf{c})_{\tau-1}^T = c_{\tau-1}^T = C_\tau^T.$$

Claim 3 still applies and completes the proof of Part 3. \parallel

Supplementary Data

Supplementary materials are available at *Review of Economic Studies* online.

Acknowledgments. We thank David Jimenez-Gomez and Pascual Restrepo for excellent research assistance and several useful suggestions, as well as Attila Ambrus, Sam Bowles, Jon Eguia, and Luca Sessa for comments on previous versions. We also thank various seminar participants, the editor, and referees for helpful comments. We gratefully acknowledge financial support from ARO MURI award W911NF-12-1-0509 and NSF grants SES-0961481 and SES-0729361.

REFERENCES

- ACEMOGLU, D. and WOLITZKY, A. (2014), "Cycles of Conflict: An Economic Model", *American Economic Review*, forthcoming.
- ANDERLINI, L., GERARDI, D. and LAGUNOFF, R. (2008), "A 'Super' Folk Theorem for Dynastic Repeated Games", *Economic Theory*, **37**, 357–394.
- ARGENZIANO, R. and GILBOA, I. (2010), "History as a Coordination Device" (Mimeo).
- AZUR, O. (2004), "What Sustains Social Norms and How do They Evolve? The Case of Tipping", *Journal of Economic Behavior and Organization*, **54**, 49–64.
- BANFIELD, E. (1958), *The Moral Basis of a Backward Society* (New York: Free Press).
- BELLOC, M. and BOWLES, S. (2013), "The Persistence of Inferior Cultural-Institutional Conventions", *American Economic Review*, *P&P*, **6**, 93–98.
- BIDNER, C. and FRANCOIS, P. (2013), "The Emergence of Political Accountability", *Quarterly Journal of Economics*, **128**, 1397–1448.
- BISIN, A. and VERDIER, T. (2001), "The Economics of Cultural Transmission and the Dynamics of Preferences", *Journal of Economic Theory*, **97**, 298–319.
- BLUME, L. (1995), "Evolutionary Equilibrium with Forward-Looking Players", (Unpublished Manuscript: Cornell University).
- BOROWIECKI, K. J. (Forthcoming), "History Matters: The Origins of Cultural Supply in Italy", Discussion Papers on Business and Economics, No.15/2012.
- BURDZY, K., FRANKEL, D. M. and PAUZNER, A. (2001), "Fast Equilibrium Selection by Rational Players Living in a Changing World", *Econometrica*, **69**, 163–189.
- CARLSSON, H. and VAN DAMME, E. (1993), "Global Games and Equilibrium Selection", *Econometrica*, **61**, 989–1018.
- CHAMLEY, C. (1999), "Coordinating Regime Switches", *Quarterly Journal of Economics*, **114**, 869–905.
- DIAMOND, P. and FUDENBERG, D. (1989), "Rational Expectations Business Cycles in Search Equilibrium", *Journal of Political Economy*, **97**, 606–619.
- DOEPKE, M. and ZILIBOTTI, F. (2008), "Occupational Choice and the Spirit of Capitalism", *Quarterly Journal of Economics*, **123**, 747–793.
- ELLISON, G. (1997), "Learning from Personal Experience: One Rational Guy and the Justification of Myopia", *Games and Economic Behavior*, **19**, 180–210.
- FRANKEL, D. M. and PAUZNER, A. (2000), "Resolving Indeterminacy in Dynamic Settings: The Role of Shocks", *Quarterly Journal of Economics*, **115**, 285–304.
- GALOR, O. (2011), *Unified Growth Theory* (Princeton, NJ: Princeton University Press).
- HERMALIN, B. E. (2012), "Leadership and Corporate Culture", in R. Gibbons and J. Roberts (eds) *Handbook of Organizations* (Princeton University Press).
- JACKSON, M. O. and PECK, J. (1991), "Speculation and Price Fluctuations with Private, Extrinsic Signals", *Journal of Economic Theory*, **55**, 274–295.
- KANDORI, M., MAILATH, G. J. and ROB, R. (1993), "Learning, Mutation, and Long run Equilibria in Games", *Econometrica*, **61**, 29–56.
- KRUGMAN, P. (1991), "History Versus Expectations", *Quarterly Journal of Economics*, **106**, 651–667.

- LAGUNOFF, R. and MATSUI, A. (1997). Asynchronous Choice in Repeated Coordination Games”, *Econometrica*, **65**, 1467–1477.
- LOCKE, R. M. (Forthcoming), “Building Trust”, in A. Flora (ed.) *Mezzogiorno e Politiche di Sviluppo*, Naples: Edizioni Scientifiche Italiane. 109–133.
- MAILATH, G. J. and SAMUELSON, L. (2006), *Repeated Games and Reputations* (Oxford University Press).
- MATSUI, A. (1996), “On Cultural Evolution: Social Norms, Rational Behavior, and Evolutionary Game Theory”, *Journal of the Japanese and International Economies*, **10**, 262–294.
- MATSUYAMA, K. (1991), “Increasing Returns, Industrialization, and Indeterminacy of Equilibrium”, *Quarterly Journal of Economics*, **106**, 617–650.
- MILGROM, P. and SHANNON, C. (1994), “Monotone Comparative Statics”, *Econometrica*, **62**, 157–180.
- MORRIS, S. and SHIN, H. S. (1998), “Unique Equilibrium in a Model of Self-Fulfilling Currency Attacks”, *American Economic Review*, **88**, 587–597.
- MYERSON, R. (2011), “Towards a Theory of Leadership and State-Building” (Mimeo, University of Chicago).
- PUTNAM, R. D. (1993), *Making Democracy Work* (Princeton, NJ: Princeton University Press).
- SABETTI, F. (1996), “Path Dependency and Civic Culture: Some Lessons from Italy about Interpreting Social Experiments”, *Politics and Society*, **24**, 19–44.
- STEINER, J. and STEWART, C. (2008), Contagion through Learning”, *Theoretical Economics*, **3**, 431–458.
- STRACCIA, U., OJEDA-ACIEGO, M. and DAMÁSIO, C. V. (2009), “On Fixed-Points of Multivalued Functions on Complete Lattices and their Application to Generalized Logic Programs”, *SIAM Journal on Computing*, **38**, 1881–1911.
- TABELLINI, G. (2008), “The Scope of Cooperation: Values and Incentives”, *Quarterly Journal of Economics*, **123**, 905–950.
- TARSKI, A. (1955), “A Lattice-Theoretical Fixpoint Theorem and its Applications”, *Pacific Journal of Mathematics*, **5**, 285–309.
- TIROLE, J. (1996), “A Theory of Collective Reputations (with Applications to the Persistence of Corruption in the Firm Quality)”, *Review of Economic Studies*, **63**, 1–22.
- TOPKIS, D. M. (1979), “Equilibrium Points in Nonzero-Sum N-Person Submodular Games”, *SIAM Journal on Control and Optimization*, **17**, 773–787.
- VAN ZANDT, T. and VIVES, X. (2007), “Monotone Equilibria in Bayesian Games of Strategic Complementarities”, *Journal of Economic Theory*, **134**, 339–360.
- VIVES, X. (1990), “Nash Equilibrium with Strategic Complementarities”, *Journal of Mathematical Economics*, **19**, 305–321.
- WOOD, G. (2009), *Empire of Liberty: A History of the Early Republic, 1789-1815* (Oxford, UK: Oxford University Press).
- YOUNG, H. P. (1993), The Evolution of Conventions”, *Econometrica*, **61**, 57–84.
- YOUNG, H. P. (Forthcoming), “Social Norms”, *New Palgrave Dictionary of Economics*.
- ZHOU, L. (1994), “The Set of Nash Equilibria of a Supermodular Game is a Complete Lattice”, *Games and Economic Behavior*, **7**, 295–300.