

Understanding Norm Change: An Evolutionary Game-Theoretic Study

Soham De^{*}
Department of Computer Science
University of Maryland
College Park, MD, USA
sohamde@cs.umd.edu

Dana Nau
Department of Computer Science
and Institute for Systems Research
University of Maryland
College Park, MD, USA
nau@cs.umd.edu

Michele Gelfand
Department of Psychology
University of Maryland
College Park, MD, USA
mgelfand@umd.edu

ABSTRACT

Human societies around the world interact with each other by developing and maintaining social norms, and it is critically important to understand how such norms emerge and change. In this work, we define an evolutionary game-theoretic model to study how norms change in a society, based on the idea that different strength of norms in societies translate to different game-theoretic interaction structures and incentives. We use this model to study, both analytically and with extensive agent-based simulations, the evolutionary relationships of the *need for coordination* in a society (which is related to its norm strength) with two key aspects of norm change: *cultural inertia* (whether or how quickly the population responds when faced with conditions that make a norm change desirable), and *exploration rate* (the willingness of agents to try out new strategies). Our results show that a high need for coordination leads to both high cultural inertia and a low exploration rate, while a low need for coordination leads to low cultural inertia and high exploration rate. This is the first work, to our knowledge, on understanding the evolutionary causal relationships among these factors.

Keywords

Agent-based analysis of human interactions, Evolutionary algorithms, Emergent behavior

1. INTRODUCTION

Human societies around the world are unique in their ability to develop, maintain, and enforce social norms. Social norms enable individuals in a society to coordinate actions, and are critical in accomplishing different tasks. Neuroscience, field, and experimental research have all established that there are marked differences in the strength of social norms around the globe [1, 10, 11, 14, 17, 18, 19, 27]. Some cultures (e.g., some middle-eastern countries, India, South Korea, etc.) are *tight*, in the sense that they tend

to have strong social norms, with a high degree of norm-adherence and higher punishment directed towards norm-violators. Other cultures (e.g., Netherlands, New Zealand, Australia, etc.) are *loose*, i.e., individuals tend to develop weaker norms with more tolerance for deviance [11, 14, 27]. This indicates that the nature of human interaction and influence is vastly different across different cultures or societies around the world.

To date, there has been little research on the evolutionary processes of norm maintenance and the processes that lead to norm change, and how these processes are substantially different in societies around the world. However, recent world events (e.g., recent social uprisings and turmoil) show that it is critically important to develop such an understanding. In this paper, we draw ideas from recent social science research to build culture-sensitive models that provide insights into the substantial societal differences that exist in how individuals interact and influence each other.

Although evolutionary game theory (EGT) was first developed to model biological evolution [20, 28, 31], it also has become useful as a way to model cultural evolution (for examples, see Section 2). In this paper, we use EGT to examine the relationships of the amount of *need for coordination* (which psychological and sociological studies show is related to norm strength [27]), with two key aspects of norm change in societies: (i) the amount of *cultural inertia*, i.e., the amount of resistance to changing a cultural norm, and (ii) the *exploration rate*, i.e., the extent to which agents are willing to try out new behaviors. More specifically, our primary contributions in this paper are as follows:

- We provide a novel way to (i) model a society's strength of norms by using an agent's need for coordination in the society, and (ii) model the desirable/undesirable norms in a society. This is done by characterizing how they affect the payoffs in a game-theoretic payoff matrix, leading to different interaction structures and incentives in a society.
- We investigate cultural evolution of norm change in this model using two well-known models of change in evolutionary game theory (the replicator dynamic [29] and the Fermi rule [2]). Using mathematical analyses and extensive agent-based simulations, we establish

An extended version of this paper, with proofs included, is available in [9].

Code for all the simulations presented in this paper is available at: <https://github.com/sohamde/inertia-aamas>

^{*}Corresponding Author

that: *the higher the need for coordination is, the higher the cultural inertia will be, and vice versa.* When a population faces conditions that make a norm change desirable, a high need for coordination will make them slower to change to the new norm compared to a society with a lower need for coordination. Further, if the need for coordination is high enough, the existing norm will not change at all.

- In order to understand how norms change in different cultures, we also examine whether the need for coordination in a society has a causal evolutionary relationship to an agent’s tendency to learn socially (i.e., adopt a behavior that is being used by other agents in the population) versus innovate/explore new random behaviors. In order to be able to do so, we propose a novel way to model this, where we let the exploration rate, i.e., the probability that an agent tries out a new action at random, *evolve* over time as part of the agent’s strategy, rather than stay fixed as in previous work [30].
- The cultural differences in the distribution of agent strategies favoring social learning versus innovation or exploration can have a critical impact on how attitudes, beliefs and behaviors spread throughout the population, and thus, is vital to understanding norm change. At a societal level, such differences can affect the rate at which new technologies, languages, moral traditions, and political institutions are adopted, while at local levels, they can alter the effectiveness of persuasion methods at the individual level. Using the above model of evolving exploration rates, we verify this by establishing, via extensive agent-based simulations, that: *the higher the need for coordination is, the lower the exploration rate will be, and vice versa.*

These results provide insight into the reasons why tight societies are less open to change, and why cultural inertia and high levels of social learning develop in such societies. To our knowledge, this is the first work to provide a culturally-sensitive model of norm change and to show how the processes of norm propagation differ across societies.

The rest of the paper is organized as follows. Section 2 includes background and related work. Section 3 provides our model of the need for coordination, and mathematical analyses and agent-based simulations showing how it affects cultural inertia. Section 4 describes our model of evolving exploration rates, and shows how the degree of need for coordination affects the evolution of exploration rates. In Section 5 we discuss the significance of our results.

2. BACKGROUND AND RELATED WORK

EGT offers a simple framework for dealing with large populations of interacting individuals, where individuals interact using different strategies, leading to game-theoretic payoffs that denote an individual’s evolutionary fitness. EGT was first developed to model biological evolution [20, 28, 31]. In such models, high-fitness individuals are more likely to reproduce than low-fitness individuals, and hence, the strategies used by those high-fitness individuals become more prevalent in the population over time. Thus, EGT studies the evolution of populations, without requiring the usual

$$M_c = \begin{array}{|c|c|c|} \hline & A & B \\ \hline A & a_c, a_c & 0, 0 \\ \hline B & 0, 0 & b_c, b_c \\ \hline \end{array} \quad M_f = \begin{array}{|c|c|c|} \hline & A & B \\ \hline A & a_f, a_f & a_f, b_f \\ \hline B & b_f, a_f & b_f, b_f \\ \hline \end{array}$$

Figure 1: Individual payoff matrices. M_c denotes the coordination game and M_f denotes the fixed-payoff game used in our model.

decision-theoretic ‘rationality’ assumptions typically used in classical game theory models.

EGT models have also been used in studies of a wide variety of social and cultural phenomena [8], e.g., cooperation [3, 12, 23, 22, 25], punishment [4, 5, 6, 24, 26], ethnocentrism [7, 13, 15], etc. In EGT models of cultural evolution, biological reproduction is replaced by social learning: if an individual uses some strategy that produces high payoffs, then others are more likely to adopt the same strategy.

EGT models of cultural evolution use highly simplified abstractions of complex human interactions, designed to capture only the essential nature of the interactions of interest. These models do not give exact numeric predictions of what would happen in real life; but they are helpful for studying the underlying dynamics of different social processes, by establishing causal relationships between various factors and observed evolutionary outcomes. Since the evolution of a human culture over time is virtually impossible to study in laboratory settings or field studies, EGT modeling provides a useful tool to apply to the study of culture and norms.

3. PROPOSED MODEL

Past field and experimental research have shown that ‘tight’ societies have stronger norms, where individuals adhere to norms much more than ‘loose’ societies, and face higher punishment when deviating. On the other hand, individuals in ‘loose’ societies typically have more tolerance for deviant behavior [11, 14, 27]. Past EGT studies have shown that a society’s exposure to *societal threat* is a key mediating factor in its strength of norms [27], where threats can be either ecological like natural disasters and scarcity of resources, or manmade such as threats of invasions and conflict. In high-threat situations, societies tend to develop strong norms for coordinating social interaction, (i.e., to become tighter), since coordination is vital for the society’s survival. In low-threat situations, there is less need for coordination, which affords weaker norms and looser societies.

Using this intuition, we hypothesize that individuals in different societies interact using *different* payoff structures and incentives. Tight societies tend to have a high need for coordination, and we can model the extreme case as a *coordination game* M_c , where one only gets a payoff if playing the same action as the agent one is interacting with. In loose societies, on the other hand, individuals’ payoffs are less affected by others’ actions, and we can model the extreme case as a *fixed-payoff game* M_f , in which an agent’s payoff depends only on the action played by that agent, and not on the actions of the other agent. For cases in between the two extremes, we use a game in which the payoff matrix is a weighted combination of a coordination game and a fixed-payoff game, with the weighting factor $0 \leq c \leq 1$ denoting the need for coordination.

As is done in many EGT studies, we consider games in which individuals have two possible actions to choose from.

$$M = \begin{array}{c|cc} & A & B \\ \hline A & ca_c + (1-c)a_f, ca_c + (1-c)a_f & (1-c)a_f, (1-c)b_f \\ \hline B & (1-c)b_f, (1-c)a_f & cb_c + (1-c)b_f, cb_c + (1-c)b_f \end{array}$$

Figure 2: Weighted payoff matrix M used in our model defined as $M = cM_c + (1-c)M_f$.

$$M' = \begin{array}{c|cc} & A & B \\ \hline A & a, a & (1-c)a, (1-c)b \\ \hline B & (1-c)b, (1-c)a & b, b \end{array}$$

Figure 3: Updated payoff matrix after assuming $a_c - b_c = a_f - b_c$ and performing a positive affine transformation on M in Figure 2.

In our case, the two actions A and B correspond to possible norms that the society could settle on. As shown in Figure 1, the coordination game has a payoff matrix M_c in which a_c and b_c are the payoff parameters; and the fixed-payoff game has a payoff matrix M_f in which a_f and b_f are the payoff parameters. The weighted combination of the two games, shown in Figure 2, is $M = cM_c + (1-c)M_f$, where $0 \leq c \leq 1$ is the need for coordination.

We first present a lemma that shows that under a mild assumption, the payoff matrix M can be much simplified under a constant translation.

LEMMA 1. *Consider the game matrix M defined in Figure 2, and assume that $a_c - b_c = a_f - b_f$. Then, under a suitable constant translation, and using $a_c = a$ and $b_c = b$, the game matrix M reduces to the matrix M' shown in Figure 3.*

PROOF. On adding the constant value of $(1-c) * (a_c - a_f) = (1-c) * (b_c - b_f)$ (where equality holds under the assumption) to all payoffs in M , the payoff matrix M reduces to M' , shown in Figure 3, where we denote $a_c = a$ and $b_c = b$. \square

The assumption $a_c - b_c = a_f - b_f$ is very reasonable, since this just ensures that switching from one norm to the other always results in the same change in payoffs, regardless of the weight c on the coordination game. Otherwise, there would be an added causal factor for the dynamics of norm change. Also note that, from Lemma 1, under constant translations, this assumption reduces to just setting $a_c = a_f$ and $b_c = b_f$. For the rest of the section, we will work with payoff matrix M' where we set $a_c = a_f = a$ and $b_c = b_f = b$. In subsequent sections, we will show why simplifying the payoff matrix by adding a constant positive value to all payoffs (as shown in Lemma 1) is a perfectly reasonable step to take.

From payoff matrix M' , we see that whenever $b < a$, the better action for the society to settle on (in terms of payoff) is A , while if $a < b$ then it is B . Let M'_{AB} be the payoff that an agent receives when they play action A and their opponent plays action B . Let M'_{AA} , M'_{BA} and M'_{BB} be defined similarly. Studying the Nash equilibrium of the game M' , we get the following lemma.

LEMMA 2. *Consider the game matrix M' defined in Figure 3, where all payoff values are positive, i.e., $a, b > 0$. Then we have:*

- (i) *If $b > a$, the strategy profile (B, B) is a Nash Equilibrium. Further, if $c \geq \frac{b-a}{b}$, then (A, A) is also a*

Nash equilibrium. Further, the strategy profile $((q, 1-q), (q, 1-q))$ is a Nash Equilibrium only when $c \geq \frac{b-a}{b}$, where $q = \frac{b-(1-c)a}{c(a+b)}$. Note that the mixed strategy $(q, 1-q)$ denotes playing action A with probability q and action B with probability $1-q$.

- (ii) *Similarly, if $a > b$, the strategy profile (A, A) is a Nash Equilibrium. Further, if $c \geq \frac{a-b}{a}$, then the strategy profile (B, B) , as well as $((q, 1-q), (q, 1-q))$ are also Nash Equilibria, with $q = \frac{b-(1-c)a}{c(a+b)}$.*

PROOF. We present a short proof for the mixed strategy Nash equilibrium when $b > a$. The other proofs follow similarly. Note that $((q, 1-q), (q, 1-q))$ is a mixed-strategy Nash Equilibrium when:

$$qM_{AA} + (1-q)M_{BA} = qM_{AB} + (1-q)M_{BB}.$$

Simplifying this, we get:

$$q = \frac{b - (1-c)a}{c(a+b)}.$$

Since $0 \leq q \leq 1$, we see that $((q, 1-q), (q, 1-q))$ is a Nash Equilibrium when both the following conditions are satisfied:

$$c \geq \frac{b-a}{b} \quad \text{and} \quad c \geq \frac{a-b}{a}.$$

When $b > a$, $c \geq \frac{a-b}{a}$ is always satisfied. Thus, when c is large enough such that $c \geq \frac{b-a}{b}$, $((q, 1-q), (q, 1-q))$ is a mixed-strategy Nash Equilibrium. \square

From Lemma 2, we see that only when c is high enough, the *sub-optimal* action pair becomes a Nash Equilibrium, where *sub-optimal* refers to the fact that the action pair that has lower payoff than the optimal action pair. This means that when $b > a$, (A, A) is the sub-optimal action pair. Thus, from Lemma 2, we see that if the need for coordination c is high, then the population may converge to either of two different equilibria, one of which is sub-optimal in terms of overall payoff. When c is low, on the other hand, the society will converge to a single globally-optimal equilibrium.

In the next two sub-sections we introduce two models for studying norm change, using two well-known models of evolutionary change (the replicator dynamic [29] and the Fermi rule [2]). We show that both models of evolutionary change are invariant to constant translations, and thus the results from this section carry forward. We derive results for how different societies respond to a need for norm change using both mathematical analysis on infinite *well-mixed* populations (where well-mixed denotes that any agent can interact with any other agent in the population), and extensive agent-based simulations on finite structured populations (where agents are placed on a network and can interact with only their neighbors).




Figure 4: Simulations with a well-mixed infinite population where reproduction is determined by the replicator dynamic with $b > a$. The figures show the change in the proportion of B agents with time.

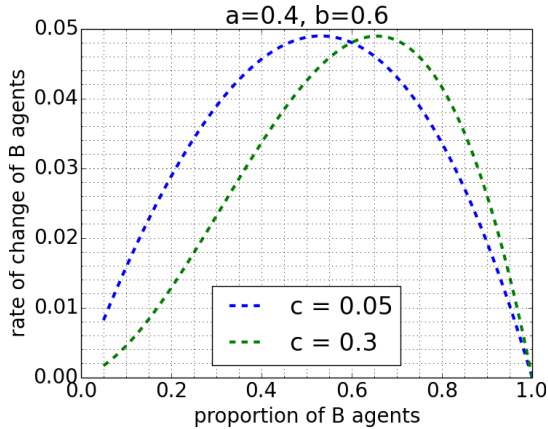


Figure 5: Simulations with a well-mixed infinite population where reproduction is determined by the replicator dynamic with $b > a$. The figure shows the rate of change of B agents versus the proportion of B agents.

3.1 Replicator dynamic on infinite well-mixed populations

Consider a well-mixed infinite population of agents. This is a standard setting used in evolutionary game theory, since a well-mixed infinite population is usually analytically tractable. Let the agents be interacting with each other using game matrix M' defined in Figure 3, and the proportion of agents playing each strategy be denoted by $x = (x_A, x_B)$, i.e., x_A proportion of agents with strategy A , and proportion $x_B = 1 - x_A$ with strategy B . Also, let $u_A(x)$ and $u_B(x)$ denote the payoffs received by an agent playing actions A and B respectively. The expected payoff for an agent is given by interacting with a randomly chosen agent in the population. Thus, we get the following:

$$\begin{aligned}\mathbb{E}[u_A(x)] &= x_A M'_{AA} + x_B M'_{AB}, \\ \mathbb{E}[u_B(x)] &= x_A M'_{BA} + x_B M'_{BB}.\end{aligned}$$

On analyzing the Nash Equilibria of this system, we observe the following lemma.

LEMMA 3. *Consider a well-mixed infinite population where agents interact using the game M' in Figure 3. Assuming all payoff values are positive, i.e., $a, b > 0$, and using Lemma 2, we have:*

- (i) *When $b > a$, $x_A = 0$ is a Nash Equilibrium. If $c \geq \frac{b-a}{b}$, then $x_A = 1$ and $x_A = \frac{b-(1-c)a}{c(a+b)}$ (which corresponds to the mixed-strategy Nash Equilibrium in Lemma 2) are also Nash Equilibria.*
- (ii) *Similarly, when $a > b$, $x_A = 1$ is a Nash Equilibrium, while if $c \geq \frac{a-b}{a}$, then $x_A = 0$ and $x_A = \frac{b-(1-c)a}{c(a+b)}$ also are Nash Equilibria.*

PROOF. We present a short proof for the intermediate case where $x_A = \frac{b-(1-c)a}{c(a+b)}$. Consider $x_A = p$ with $0 < p < 1$. For $x_A = p$ to be a NE, the following two conditions need to be simultaneously satisfied:

$$\begin{aligned}pM'_{AA} + (1-p)M'_{AB} &\geq pM'_{BA} + (1-p)M'_{BB}, \\ \text{and } pM'_{BA} + (1-p)M'_{BB} &\geq pM'_{AA} + (1-p)M'_{AB}.\end{aligned}$$

Both of these conditions are satisfied only when:

$$pM'_{AA} + (1-p)M'_{AB} = pM'_{BA} + (1-p)M'_{BB}.$$

This simplifies to:

$$p = \frac{b - (1-c)a}{c(a+b)},$$

and similar to Lemma 2, the results follow. \square

We assume that on each iteration, agents interact with other randomly chosen agents, and the population evolves according to the replicator dynamic. The replicator dynamic is based on the idea that the proportion of agents of a type (or strategy) increases when it achieves expected payoff higher than the average payoff, and decreases when achieving lower payoff than the average payoff. Thus, over time, the proportion of agents of a type that achieves payoff higher than the average payoff starts increasing in the population, and eventually take over. More formally, the replicator dynamic is given by the differential equation

$$\dot{x}_A = \frac{dx_A}{dt} = x_A \cdot (\mathbb{E}[u_A(x)] - \theta(x)), \quad (1)$$

where $\theta(x) = x_A \mathbb{E}[u_A(x)] + x_B \mathbb{E}[u_B(x)]$ is the average payoff received by all agents in the population. From (1), it is clear that the rate of change remains the same under constant translations of the payoff matrix, since these translations would cancel each other out. Thus, the assumption made in the previous section, follows through to this section as well.

Using the game matrix M' , the rate of change in the proportion x_A is given by:

$$\dot{x}_A = x_A(1 - x_A)(c(a + b)x_A - (b - (1 - c)a)). \quad (2)$$

The fixed points of this rate of change are given by:

$$x_A = 0, \quad x_A = 1, \quad \text{and} \quad x_A = \frac{b - (1 - c)a}{c(a + b)}. \quad (3)$$

These correspond to the Nash Equilibria derived earlier. Next, we study the stability of the Nash equilibria derived above, where we define a stable Nash equilibrium to be one where: if an infinitesimal proportion of agents change their strategy from the Nash equilibrium, the population, using the replicator dynamic, reverts back to the original Nash equilibrium. If an infinitesimal change in the proportion of agents playing each strategy leads the population to converge to a different Nash equilibrium, the original Nash equilibrium is called unstable. Thus, we state the following corollary.

COROLLARY 1. *From Lemma 3 and Eq. (2) and Eq. (3), we see that the Nash Equilibria $x_A = 0$ and $x_A = 1$ are stable, while the Nash Equilibrium $x_A = \frac{b - (1 - c)a}{c(a + b)}$ is unstable.*

PROOF. Let $\phi = \frac{b - (1 - c)a}{c(a + b)}$. We notice that, if $x_A = \phi + \epsilon$, then $\dot{x}_A > 0$, while if $x_A = \phi - \epsilon$, then $\dot{x}_A < 0$, for any small $\epsilon > 0$. Thus, $x_A = \frac{b - (1 - c)a}{c(a + b)}$ represents an unstable fixed point, while $x_A = 0$ and $x_A = 1$ represent stable fixed points. \square

There is a further notion of equilibrium used in EGT called evolutionarily stable strategies (ESS). A strategy is called an ESS when, if adopted by the population, it cannot be taken over by an initially small proportion of agents with a different strategy. Thus, we state the following theorem.

THEOREM 1. *From Lemma 3 and Corollary 1, we see:*

- (i) *When $b > a$, B is an ESS. If $c \geq \frac{b - a}{b}$, then A is also an ESS.*
- (ii) *When $a > b$, A is an ESS. If $c \geq \frac{a - b}{a}$, then B is also an ESS.*

Due to lack of space, we omit the proof for Theorem 1. We observe that the strategies A and B are Evolutionary Stable Strategies (ESS), when adopted by everyone in the population (corresponding to the stable Nash equilibria $x_A = 1$ and $x_A = 0$). The unstable Nash Equilibrium, on the other hand, does not correspond to an ESS, since even a small group is able to invade and change the equilibrium. Thus, only stable Nash Equilibria correspond to evolutionarily stable strategies.

Theorem 1 indicates that a society is bound to end up at one of the evolutionarily stable strategies: with every individual on action A or everyone on action B , since even a small perturbation moves the society away from the unstable Nash equilibrium. When c is low, there exists only a single ESS, and thus the society adapts itself and settles on the ESS. When c is high, there are two ESSs, and thus the society might settle on either one, depending on the starting point of the society.

Let us consider two societies: one with a lower need for coordination is c_1 , and one with a high need for coordination

$c_2 > c_1$. To avoid some awkward phrasing, we'll call these the "looser" and "tighter" societies, respectively. Suppose a majority of both societies are playing norm A , and suppose they evolve according to the replicator dynamic given in Eq. (2). We are interested in how these two societies would respond to the action B , when the payoff of action B is higher than A , i.e., when $b > a$, or equivalently, $M'_{BB} > M'_{AA}$. First notice that if $c_2 > (b - a)/b$, and $c_1 < (b - a)/b$, it follows from Theorem 1 that the tighter society remains on norm A while the looser one switches to the globally optimal norm B .

Now suppose the difference in norm payoffs is large enough such that $c_2 < (b - a)/b$ (and thus also, $c_1 < (b - a)/b$). This ensures that there is only a single equilibrium for both societies at $x_A = 0$. Thus, both societies would switch to norm B , and we are interested in the rate at which this change occurs. Let \dot{x}_{B1} and \dot{x}_{B2} denote the rate of change when the need to conform is c_1 or c_2 , respectively. Then we can show that:

$$\dot{x}_{B2} - \dot{x}_{B1} = x_B(1 - x_B)(c_2 - c_1)((a + b)x_B - b).$$

This simplifies to

$$\dot{x}_{B2} - \dot{x}_{B1} \begin{cases} \leq 0, & \text{when } x_B \leq \frac{b}{a + b}; \\ > 0, & \text{when } x_B > \frac{b}{a + b}. \end{cases} \quad (4)$$

Thus, $\dot{x}_{B2} < \dot{x}_{B1}$ in the initial stages when $x_B < b/(a + b)$. However, once the proportion of B agents become big enough such that $x_B > b/(a + b)$, then the higher the value of c , the higher the rate of change will be. Thus, when c is high, the switch from A to B takes time to speed up, with more cultural inertia than when c is low, even when the payoff of the new norm is arbitrarily large compared to the previous norm. The initial cultural inertia results in the society with a higher c value to take longer overall to switch to the new norm.

Figure 4 illustrates these properties of well-mixed populations using the replicator dynamic. In all the simulations, we start off the society at $x_A = 0.95$. In the first of the three graphs, the tighter society (again using "tighter" as shorthand for "higher need for coordination") has $c > (b - a)/b$. Thus, while the less-tight society switches to the more beneficial norm B , the tighter society is resistant to the change (since the difference in payoffs is small) and stays with norm A . The second and third graphs show situations where both societies switch to norm B . We observe that the tighter society switches more slowly towards changing to norm B , but the difference in speed decreases as the difference in payoffs between B and A increases.

As derived in (4), the rate of change for a society with higher c grows larger than a society with lower c only after $x_B > \frac{b}{a + b}$. This is shown in Figure 5. This also indicates the initial inertia that societies with a higher need for coordination experience towards changing norms. The need for coordination in these societies lead to individuals being reluctant to try out new norms, which in turn leads to inertia.

3.2 Agent simulations on finite networks

A limitation of the above model is that it assumes that the population is *infinite* and *well-mixed*. While the assumption that a population is infinite is not a bad approximation for very large populations (which is the scale that we are interested in), the assumption that agents are well-mixed, i.e.,

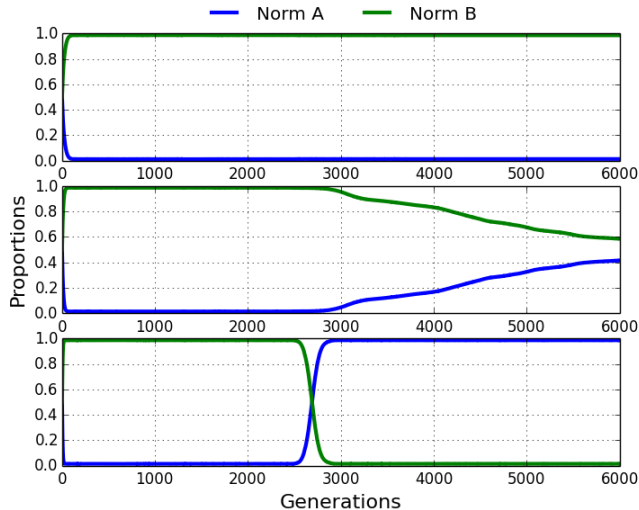


Figure 6: From top to bottom: $c = 1.0$, $c = 0.75$, $c = 0.5$. Initially: $a = 1.0, b = 1.15$. Structural shock at 2500 iterations. After structural shock, payoffs become: $a = 1.15, b = 1.0$. We see that the higher the need for coordination, the more the cultural inertia.

where any agent can interact with any other agent, is often inaccurate. In this section, we show that the results derived in the previous section, also extend to cases where agents are structured on the nodes of a graph/network, where agents can only interact with another agent if they are connected by an edge in the graph.

More specifically, we now consider a structured population where agents are arranged on the nodes of a toroidal (wrap-around) grid, such that each agent can interact only with the 4 other agents they are connected to. We consider toroidal grids as a convenient example, however, the results we describe below also extend to other network structures like small-world networks, and preferential attachment models. Mathematical analysis of evolutionary games on structured populations is not yet a well-developed field, and thus we perform simulations of our model as follows.

Initially, we arrange agents with random strategies (A or B) on each node of the grid. In each iteration, each pair of agents connected by an edge interact in a two-player game defined by the payoff matrix M . The total payoff of each agent is computed by summing over the payoffs received by an agent for each game that they play. Since the population is finite, we use dynamics defined on finite populations. After each interaction phase, agents use the Fermi rule to update its strategy for the next iteration. Under the Fermi rule, an agent ψ_a picks a random neighbor ψ_n and observes its payoff, and the agent then decides to switch to the neighbor's strategy with probability $p = (1 + \exp(-s(u_a - u_n)))^{-1}$, where u_a and u_n are the payoffs of the agent and the neighbor, and s is a user-defined parameter (in all our experiments, we set $s = 5$). With probability $1 - p$, the agent retains its old strategy. With a small probability μ , called the exploration rate, an agent also tries out an action completely randomly. This repeats for every iteration of the simulation. Note that the Fermi rule also only depends on a difference between payoff values, and thus, like the replica-

tor dynamic, is also invariant to constant translations. Thus, for all our experiments in this section, we use the simplified game matrix M' from Figure 3.

To study cultural inertia (i.e., resistance to changing a cultural norm) or rapid cultural change in different societies, we use a game-theoretic model of a *structural shock*. A structural shock represents a catastrophic incident in a society, where suddenly there is abrupt change in the payoffs for actions A and B . We are interested in studying how societies with different needs for coordination react to such an abrupt and drastic shift in the payoffs of the possible actions. In our EGT model, we implement a structural shock by simply interchanging the payoffs of actions A and B , thus, denoting a sudden change in the globally optimal action in a society. This is equivalent to interchanging the payoff values a and b . Thus, if initially, we have $b > a$, after a structural shock, we get $a > b$.

Now consider that, initially, the action with a higher utility (and the current norm) in a society is B , i.e., $b > a$ with $x_A = 0$. Suppose, the society experiences a structural shock, where now action A becomes more desirable with $a > b$. On introducing a small proportion of agents playing norm A (say $x_A = 0.01$), if the need for coordination is low then the population will switch to the new norm with $x_A = 1$. This is because, after the structural shock, the Nash Equilibrium (and ESS) is $x_A = 1$, as shown above. However, if the need for coordination is high (i.e., $c \geq \frac{a-b}{a}$), then $x_A = 0$ is still a Nash Equilibrium (and ESS) and the population will remain on the sub-optimal norm B even after the structural shock.

All experiments were run on a grid with 2500 nodes, and the simulation goes on for 6000 iterations, with a structural shock implemented at 2500 iterations. 100 independent simulations are run for each setting and the results are averaged over the 100 runs. Figure 6 shows the results of our simulations. The plots show the proportion of agents playing norm A vs norm B . As before, the parameter c denotes the need for coordination. When c is low, very little cultural inertia develops and agents are more willing to innovate by exploring behaviors other than the current societal norms. In this case, the population will change more quickly to a different norm if the new norm will be beneficial. By contrast, when c is high, we see the evolutionary emergence of higher levels of cultural inertia, with agents less willing to innovate or to violate established cultural norms. In this case, the population is slower to change to the new norm, and if c is high enough it may not change at all. Thus, qualitatively, the results with a structured populations match those from the infinite well-mixed populations in Section 3.1, and the mechanics that lead to the above results can be explained using the same equilibrium results derived above.

4. EVOLVING EXPLORATION RATES

In addition to the amount of cultural inertia, another key aspect to understanding how norms change in different cultures is to study an agent's tendency to learn socially (i.e., adopt a behavior that is being used by other agents in the population) or innovate and explore new random behaviors. Such tendencies are critical in understanding the rate at which new technologies, languages, or moral traditions are adopted in a population, and help us understand the processes of influence and persuasion at the individual level.

In the model presented in Section 3.2 for finite structured populations, we kept the exploration rate (i.e., the small

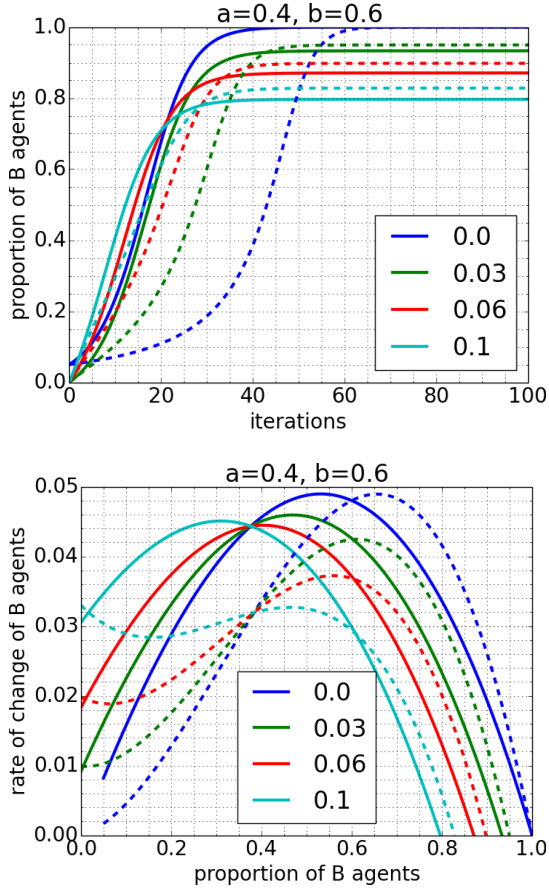


Figure 7: The solid lines denote $c = 0.05$, while the dotted lines denote $c = 0.3$. The colors denote the mutation rates.

probability with which an agent tries out a new strategy at random) at a constant low value. This exploration rate denotes how much an agent is open to change and trying out new actions at random. Thus, it seems that the need for coordination in a society might affect how likely an individual is to try out different actions, instead of conforming to their neighbors. Particularly, it seems natural to assume that individuals in tight societies are much less likely to try out random actions than individuals in loose societies [11, 14]. In this section, we test this hypothesis by presenting a model to study the evolution of exploration rates in different societies.

To get some intuition about the hypothesis, we go back to our setting of a well-mixed infinite population. Note that the replicator dynamic does not have a provision for exploration rates. Thus, we use a variant of the replicator dynamic called the replicator-mutator equation [21]. Using this variant, one can include exploration rates into the replicator dynamic. Thus, if we fix μ to be the exploration rate, we can write the replicator-mutator equation as:

$$\begin{aligned} \dot{x}_A &= (1 - \mu)x_A\mathbb{E}[u_A(x)] + \mu x_A\mathbb{E}[u_B(x)] - x_A\theta(x), \\ &= x_A(\mathbb{E}[u_A(x)] - \theta(x)) + \mu x_A(\mathbb{E}[u_B(x)] - \mathbb{E}[u_A(x)]). \end{aligned}$$

Thus, like the replicator dynamic, we can write the rate of change in terms of payoff differences, which makes the dynamic invariant to constant translations. Thus, in this

section, we again use the simplified game matrix M' from Figure 3. Simplifying the equation for \dot{x}_A , we get:

$$\begin{aligned} \dot{x}_A &= x_A(1 - x_A)(c(a + b)x_A - (b - (1 - c)a)) \\ &\quad + \mu(x_A x_B(1 - c)(b - a) + (x_B^2 b - x_A^2 a)). \end{aligned} \quad (5)$$

Figure 7 plots the replicator-mutator equation (Eq. (5)) with a well-mixed infinite population. The solid lines are for a low need for coordination ($c = 0.05$), while the dotted lines are for a high need for coordination ($c = 0.3$), and we plot the proportion of B agents, as well as the rate of change, for various exploration rate μ values. From the figure, we see that for all exploration rates μ , when the need for coordination is high then there is higher cultural inertia.

To study the evolution of explorative behaviors, we let the exploration rate (referred to as the mutation rate in biological models) *evolve*. The exploration rate is the probability μ with which an agent chooses a random new strategy at each iteration ($0 \leq \mu \ll 1$). In biological evolution, mutation occurs so rarely that game-theoretic biological models often omit it. In cultural evolution, however, exploration is an important step since individuals try out new behaviors much more frequently [30]. Studying the evolution of exploration rates helps us get insights about a society's openness to change. Low exploration rates suggest that individuals are less likely to try out new strategies and are more likely to coordinate with their neighbors. On the other hand, high exploration rates mean that individuals are more open to change and innovation.

To model the evolution of exploration rates, we first create a set L of possible exploration rates. These can be a finite discrete set of exploration rates. For all our experiments, we use the set of exploration rates: $L = \{0.0, 0.1, 0.2, 0.3, 0.4, 0.5\}$. The exploration rate is added as part of the strategy of an agent, and each individual now chooses an exploration rate in addition to the game action (A or B). Thus, an agent now copies the exploration rate of a neighbor, along with the game action, when updating its strategy using the Fermi rule.

Note that, a regularly changing environment is essential for studying the evolution of exploration rates since, if the environment is not changing frequently enough, an exploration rate of 0 would always be evolutionarily stable. To model the changing environment, we will use the same switch in dominant norms (structural shock) that we used in our earlier experiments, except now we apply the structural shock multiple times at much shorter and regular intervals. We use a fixed interval of 75 iterations to apply the structural shock. We run the simulation for a total of 2000 iterations. For these experiments, an agent's strategy set now becomes a size of 10: 5 possible exploration rates in $L \times 2$ possible game actions (norm A or norm B). We use the same toroidal grid as described before. Figure 8 shows the experimental results. Each column in Figure 8 shows, for a specific c , the proportion of agents playing norm A vs norm B (top plot), and the proportion of agents with each exploration rate (bottom plot).

We see that when the need for coordination is high, an exploration rate of $\mu = 0.0$ is adopted by the majority of the society. Individuals in such a society are more likely to adopt the strategies of their neighbors, and this leads to high cultural inertia. In loose societies, however, higher proportions of exploration rates $\mu > 0$ evolve, and individuals are more open towards change, leading to lower cultural inertia.

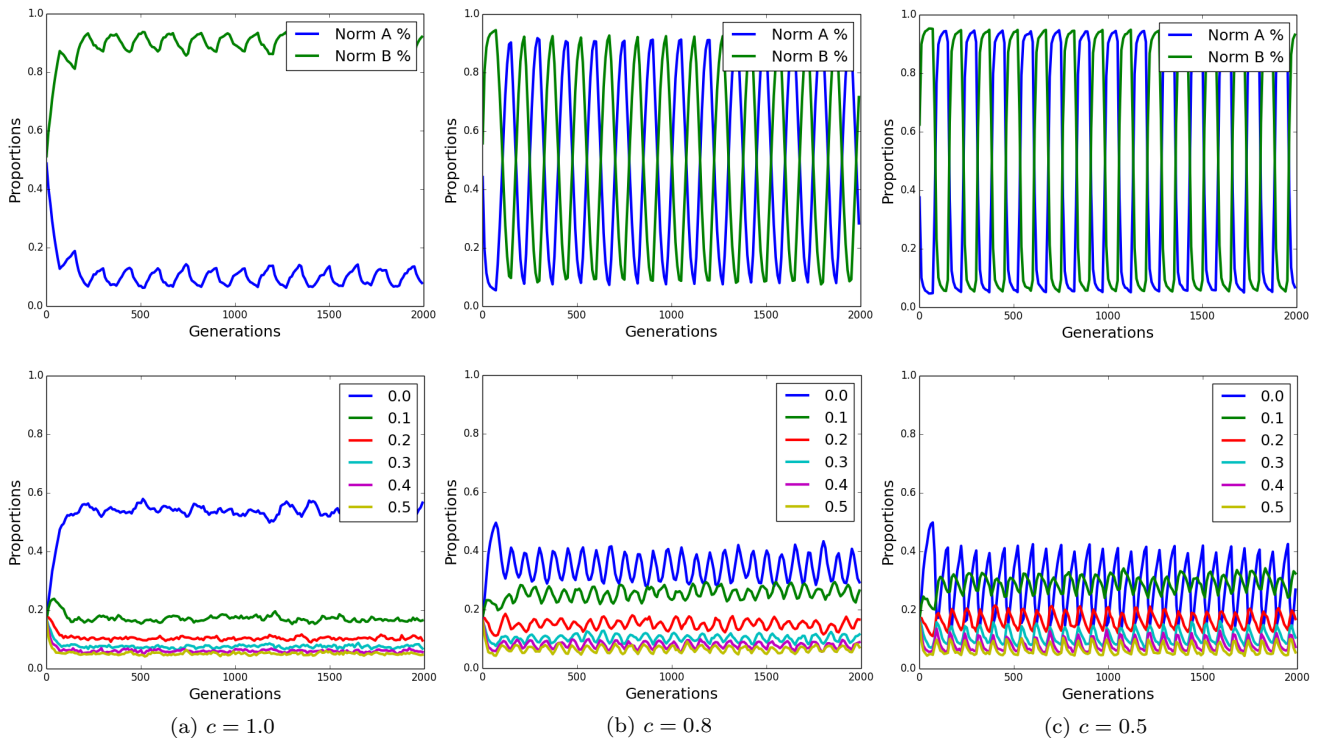


Figure 8: Initially: $a = 1.0, b = 1.15$.

This fits well with our results in Section 3, and provides insights into why cultural inertia develops in societies with a higher need for coordination.

5. DISCUSSION

In this paper, we study cultural inertia and norm change in human societies. We build evolutionary game-theoretic models that show that societies that have a higher need for coordination have higher cultural inertia, with individuals being less likely to switch to the new norm even when it might have a larger payoff. Societies with a lower need for coordination, on the other hand, have low cultural inertia, with individuals more willing to innovate and open to change.

In prior EGT models, the exploration rate (i.e., the small probability with which an agent tries out a new strategy at random) has been kept constant. However, prior work indicates that the exploration rate might be adaptive to the structural properties of the society. In this work, we propose a model that studies how exploration rates evolve in different human societies. Our model shows that when the need for coordination is high, an exploration rate of 0 is adopted by the majority, and individuals are more likely to adopt the strategies of their neighbors. When the need for coordination is low, higher exploration rates evolve leading to lower cultural inertia, and more openness to change. This explains why tight cultures tend to have less deviant behavior among individuals with more norm adherence.

To our knowledge, we propose the first EGT model that predicts the effects of the need for coordination on norm change and cultural inertia. In the future, it would be interesting to study how network structures differ between tight and loose cultures. This would provide the results neces-

sary to extend our work and study the dynamics of the rate at which norm change occurs in different cultures. Further, while all current work on information diffusion has concentrated on general models, our work indicates that the structure of interaction and incentives are vastly different across societies. Thus, it would be interesting to study culture-sensitive models for information diffusion and propagation, that would be more accurate in predicting cascades and epidemics in different societies around the world.

By studying how socio-structural factors such as the need for coordination affect cultural inertia, this work aims to establish a culturally-sensitive model of norm change. With this model, we identify the conditions that lead to stability or instability in established population norms in different cultural contexts. Such knowledge is critical in providing us the ability to identify early markers of impending drastic shifts in populations' norms and thus enable tools providing alerts to potential social uprisings and turmoil.

Acknowledgments

This research was funded by the U.S. Air Force Grant FA9550-14-1-0020. The authors acknowledge the University of Maryland supercomputing resources (<http://www.it.umd.edu/hpcc>) made available for conducting the research reported in this paper.

REFERENCES

- [1] D. Balliet and P. A. Van Lange. Trust, punishment, and cooperation across 18 societies a meta-analysis. *Perspectives on Psychological Science*, 8(4):363–379, 2013.
- [2] L. E. Blume. The statistical mechanics of strategic

- interaction. *Games and economic behavior*, 5(3):387–424, 1993.
- [3] S. Bowles and H. Gintis. The evolution of strong reciprocity: cooperation in heterogeneous populations. *Theoretical population biology*, 65(1):17–28, 2004.
- [4] R. Boyd, H. Gintis, S. Bowles, and P. J. Richerson. The evolution of altruistic punishment. *Proceedings of the National Academy of Sciences*, 100(6):3531–3535, 2003.
- [5] H. Brandt, C. Hauert, and K. Sigmund. Punishment and reputation in spatial public goods games. *Proceedings of the Royal Society of London B: Biological Sciences*, 270(1519):1099–1104, 2003.
- [6] H. Brandt, C. Hauert, and K. Sigmund. Punishing and abstaining for public goods. *Proceedings of the National Academy of Sciences*, 103(2):495–497, 2006.
- [7] S. De, M. J. Gelfand, D. Nau, and P. Roos. The inevitability of ethnocentrism revisited: Ethnocentrism diminishes as mobility increases. *Scientific reports*, 5, 2015.
- [8] S. De, D. Nau, and M. Gelfand. Using game theory to study the evolution of cultural norms. *arXiv preprint arXiv:1606.02570*, 2016.
- [9] S. De, D. Nau, and M. Gelfand. Understanding norm change: An evolutionary game-theoretic study (extended version). *arXiv preprint*, 2017.
- [10] J. Ensminger and J. Henrich. *Experimenting with social norms: Fairness and punishment in cross-cultural perspective*. Russell Sage Foundation, 2014.
- [11] M. J. Gelfand, J. L. Raver, L. Nishii, L. M. Leslie, J. Lun, B. C. Lim, L. Duan, A. Almaliach, S. Ang, J. Arnadottir, et al. Differences between tight and loose cultures: A 33-nation study. *science*, 332(6033):1100–1104, 2011.
- [12] W. D. Hamilton and R. Axelrod. The evolution of cooperation. *Science*, 211(27):1390–1396, 1981.
- [13] R. A. Hammond and R. Axelrod. The evolution of ethnocentrism. *Journal of Conflict Resolution*, 50(6):926–936, 2006.
- [14] J. R. Harrington and M. J. Gelfand. Tightness–looseness across the 50 united states. *Proceedings of the National Academy of Sciences*, 111(22):7990–7995, 2014.
- [15] M. Hartshorn, A. Kaznatcheev, and T. Shultz. The evolutionary dominance of ethnocentric cooperation. *Journal of Artificial Societies and Social Simulation*, 16(3):7, 2013.
- [16] J. Henrich and R. Boyd. Why people punish defectors: Weak conformist transmission can stabilize costly enforcement of norms in cooperative dilemmas. *Journal of theoretical biology*, 208(1):79–89, 2001.
- [17] J. Henrich, J. Ensminger, R. McElreath, A. Barr, C. Barrett, A. Bolyanatz, J. C. Cardenas, M. Gurven, E. Gwako, N. Henrich, et al. Markets, religion, community size, and the evolution of fairness and punishment. *science*, 327(5972):1480–1484, 2010.
- [18] J. Henrich, R. McElreath, A. Barr, J. Ensminger, C. Barrett, A. Bolyanatz, J. C. Cardenas, M. Gurven, E. Gwako, N. Henrich, et al. Costly punishment across human societies. *Science*, 312(5781):1767–1770, 2006.
- [19] B. Herrmann, C. Thöni, and S. Gächter. Antisocial punishment across societies. *Science*, 319(5868):1362–1367, 2008.
- [20] J. Hofbauer and K. Sigmund. *Evolutionary games and population dynamics*. Cambridge university press, 1998.
- [21] S. Levin. Complex adaptive systems: exploring the known, the unknown and the unknowable. *Bulletin of the American Mathematical Society*, 40(1):3–19, 2003.
- [22] M. A. Nowak. Five rules for the evolution of cooperation. *science*, 314(5805):1560–1563, 2006.
- [23] M. A. Nowak and K. Sigmund. Tit for tat in heterogeneous populations. *Nature*, 355(6357):250–253, 1992.
- [24] D. G. Rand and M. A. Nowak. The evolution of antisocial punishment in optional public goods games. *Nature communications*, 2:434, 2011.
- [25] R. L. Riolo, M. D. Cohen, and R. Axelrod. Evolution of cooperation without reciprocity. *Nature*, 414(6862):441–443, 2001.
- [26] P. Roos, M. Gelfand, D. Nau, and R. Carr. High strength-of-ties and low mobility enable the evolution of third-party punishment. *Proceedings of the Royal Society of London B: Biological Sciences*, 281(1776):20132661, 2014.
- [27] P. Roos, M. Gelfand, D. Nau, and J. Lun. Societal threat and cultural variation in the strength of social norms: An evolutionary basis. *Organizational Behavior and Human Decision Processes*, 129:14–23, 2015.
- [28] J. M. Smith. *Evolution and the Theory of Games*. Cambridge university press, 1982.
- [29] P. D. Taylor and L. B. Jonker. Evolutionary stable strategies and game dynamics. *Mathematical biosciences*, 40(1-2):145–156, 1978.
- [30] A. Traulsen, C. Hauert, H. De Silva, M. A. Nowak, and K. Sigmund. Exploration dynamics in evolutionary games. *Proceedings of the National Academy of Sciences*, 106(3):709–712, 2009.
- [31] J. W. Weibull. *Evolutionary game theory*. MIT press, 1997.