The Evolution of Social Norms

H. Peyton Young

Department of Economics, Oxford University, Oxford OX1 3UQ, United Kingdom;
email: peyton.young@economics.ox.ac.uk

Keywords

Abstract

Social norms are patterns of behavior that are self-enforcing within
a group: Everyone conforms, everyone is expected to conform, and ev-
eryone wants to conform when they expect everyone else to conform.
Social norms are often sustained by multiple mechanisms, including
a desire to coordinate, fear of being sanctioned, signaling membership
in a group, or simply following the lead of others. This article shows
how stochastic evolutionary game theory can be used to study
the resulting dynamics. I illustrate with a variety of examples drawn
from economics, sociology, demography, and political science. These
include bargaining norms, norms governing the terms of contracts,
norms of retirement, dueling, foot binding, medical treatment, and
the use of contraceptives. These cases highlight the challenges of ap-
plying the theory to empirical cases. They also show that the modern
theory of norm dynamics yields insights and predictions that go be-
yond conventional equilibrium analysis.

1. OVERVIEW

Social norms govern our interactions with others. They are the unwritten codes and informal understandings that define what we expect of other people and what they expect of us. Norms establish standards of dress and decorum, obligations to family members, property rights, contractual relationships, conceptions of right and wrong, notions of fairness, and the meanings of words. They are the building blocks of social order. Despite their importance, however, they are so embedded in our ways of thinking and acting that we often follow them unconsciously and without deliberation; hence we are sometimes unaware of how crucial they are to navigating social and economic relationships.

There is a substantial literature on social norms in philosophy, sociology, anthropology, law, political science, and economics. Given space limitations, it is impossible to provide a comprehensive account of this literature here.[1] Instead I shall focus on the question of how social norms evolve and how norm shifts take place using evolutionary game theory as the framework of analysis. Although this framework is relatively recent, many of the key concepts originate with Hume (1888 [1739], p. 490):

> I observe, that it will be for my interest to leave another in the possession of his goods, *provided* he will act in the same manner with regard to me. He is sensible of a like interest in the regulation of his conduct. When this common sense of interest is mutually express'd, and is known to both, it produces a suitable resolution and behaviour. And this may properly enough be call'd a convention or agreement betwixt us, tho' without the interposition of a promise; since the actions of each of us have a reference to those of the other, and are perform'd upon the supposition, that something is to be perform'd on the other part. Two men, who pull the oars of a boat, do it by an agreement or convention, tho' they have never given promises to each other. Nor is the rule concerning the stability of possession the less deriv'd from human conventions, that it arises gradually, and acquires force by a slow progression, and by our repeated experience of the inconveniences of transgressing it. On the contrary, this experience assures us still more, that the sense of interest has become common to all our fellows, and gives us a confidence of the future regularity of their conduct. . . . In like manner are languages gradually establish'd by human conventions without any promise. In like manner do gold and silver become the common measures of exchange. . .

In this remarkable passage, Hume identifies three key factors in the evolution of norms: (*a*) They are equilibria of repeated games; (*b*) they evolve through a dynamic learning process; and (*c*) they underpin many forms of social and economic order. In recent years, these ideas have been formalized using the tools of evolutionary game theory (Foster & Young 1990; Young 1993a,b 1995, 1998a,b; Kandori et al. 1993; Skyrms 1996, 2004; Binmore 1994, 2005; Bowles 2004).[2] Although the theory is well-developed, there has been relatively little research that brings it into contact with empirical examples. I have therefore structured this article around seven case studies, some of which are based on experiments, some on historical narratives, and some on field data. My aim is to show how evolutionary game theory can illuminate the evolution of norms in real

---

[1]See, in particular, Schelling (1960, 1978), Lewis (1969), Ullmann-Margalit (1977), Akerlof (1980), Axelrod (1984, 1986), Sugden (1986), Coleman (1987, 1990), Elster (1989a,b), North (1990), Ellickson (1991), Kandori (1992), Young (1993a,b, 1998a,b), Bernheim (1994), Skyrms (1996, 2004), Binmore (1994, 2005), Posner (2000), Hechter & Opp (2001), Boyd & Richerson (2002, 2005), Bowles (2004), Bicchieri (2006), Burke & Young (2011), and Myerson & Weibull (2015).

[2]Other textbook treatments of evolutionary game theory without a particular focus on norms include Vega-Redondo (1996), Samuelson (1997), Hofbauer & Sigmund (1998), and Sandholm (2010).

situations despite limitations in the data. The discussion will also highlight some shortcomings of current theory and point out directions for future research.

Let me be clear at the outset that I shall not attempt to distinguish between various categories of social norms, such as customs, conventions, conceptions of right and wrong, notions of propriety, and regularities of behavior. Instead I shall use the term "social norm" to cover a constellation of behaviors that range from fine points of etiquette to strong conceptions of moral duty. In all of their incarnations, social norms exhibit certain key features that have important implications for the ways in which they evolve. These include the following:

1. Norms are behaviors that are self-enforcing at the group level: People want to adhere to the norm if they expect others to adhere to it. The precise mechanisms and motivations that create this positive feedback loop differ substantially from one context to another, as I shall show in subsequent sections.
2. Norms typically evolve without top-down direction through a process of trial and error, experimentation, and adaptation. They illustrate how social order is constructed through interactions of individuals rather than by design.
3. Norms can take alternative forms; that is, they govern interactions that have multiple equilibria. Consequently, they are contingent on context, social group, and historical circumstances. Consider the implications of having an illegitimate child, ignoring a challenge to a duel, binding your daughter's feet, leaving all your property to your eldest son, practicing contraception, keeping a mistress, burping at the end of a meal, or dancing at a funeral. In some societies, these would represent serious norm violations; in others, they are perfectly normal practice.

2. MECHANISMS SUPPORTING NORMATIVE BEHAVIOR

Norms come in different guises, serve a variety of purposes, and are sustained as equilibria by multiple mechanisms. The literature provides a rich account of different types of norms and the social and psychological mechanisms that support them.[3] Here I shall briefly summarize some of the main mechanisms that support social norms without attempting to be exhaustive.

2.1. Coordination

Examples of pure coordination include using words with their conventional meanings, carrying standard forms of money to make purchases, and negotiating the terms of economic contracts. The motive is to coordinate with others in a particular type of interaction; there is no need to punish deviants because failure to coordinate is inherently costly.

2.2. Social Pressure

Norms often prescribe behaviors that run counter to an individual's immediate self-interest, in which case they are sustained by the prospect of social disapproval, ostracism, loss of status, and other forms of social punishment. Numerous forms of cooperation are enforced in this manner, but so are dysfunctional norms, such as dueling, blood feuds, and foot binding.[4]

---

[3]See, in particular, Homans (1950), Coleman (1987, 1990), Elster (1989a,b), Hechter & Opp (2001), and Bicchieri (2006).

[4]For other examples of dysfunctional norms, see Edgerton (1992). There is an ongoing debate about what motivates third parties to inflict costly punishments, but laboratory experiments suggest that they are in fact willing to do so (Fehr et al. 2002; Fehr & Fischbacher 2004a,b).

2.3. Signaling and Symbolism

Norms can signal intentions, aspects of personal character, or membership in a group. Although the behaviors themselves are of little consequence, they have important reputational implications. Dress codes are often used to signal membership in specific groups or the holding of particular preferences, such as veiling by Muslim women (Carvalho 2013), "hanky codes" among gay men, and tattoos in criminal gangs (Gambetta 2009). A current example is the use of flag lapel pins by US politicians to signal their patriotism. Other signals connote general rather than specific traits: Observing fine points of etiquette, showing up on time, speaking in turn, and displaying appropriate degrees of deference are often taken to be signs of reliability and trustworthiness (Posner 2000). The mechanism that sustains such symbolic norms is the negative interpretation that others place on deviations.

2.4. Benchmarks and Reference Points

Social regularities in behavior sometimes result from the use of benchmarks or reference points to make decisions. Schelling (1960) emphasized the importance of such "focal points" to solve coordination problems; a prominent example is 50-50 division in bargaining. However, there are other types of socially constructed reference points that serve as heuristics for making individual decisions rather than coordinating with someone else.[5] These often take the form of targets, benchmarks, or thresholds, such as what time of day to have the first drink, what fraction of one's income to give to charity, how much to save for retirement, how old one should be before getting married, or how much one should weigh.[6] Such benchmarks derive their salience from the fact that they are widely used; they are a convenient way of making (or justifying) a decision. Thus, as in the previous cases, there is a positive feedback effect between the behaviors of the group and the actions of the individual.

A particularly interesting example of such a reference point is the age at which people customarily retire. In the United States, the social security system was originally designed to encourage retirement at age 65. In the early 1960s, the law was changed so that the expected benefits remained essentially the same for anyone retiring between 62 and 65, but many people did not take advantage of the change. Indeed the distribution of retirement ages continued to have a spike at 65 until well into the 1990s (Axtell & Epstein 1999, Burtless 2006). It appears that once 65 became fixed in the public mind as the "normal" age to retire, it continued to serve as a heuristic for many years after it had ceased to have any particular economic significance.[7]

The preceding discussion does not cover all of the mechanisms that can support a social norm, nor are these mechanisms mutually exclusive. Foot binding one's daughter enhances her marriage prospects (a coordination motive) and also signals the family's status. Retiring at 65 may be largely a heuristic decision, but it may also have elements of coordination (retiring when one's friends retire). Using bad table manners signals a lack of sensitivity to social norms (symbolic), and it may preclude future profitable dealings with others at the table (coordination failure).

---

[5]For more on heuristics, see Kahneman et al. (1982), Gigerenzer et al. (1999), and Boyd & Richerson (2005).

[6]Survey data suggest that people's choice of body weight is subject to social influence; that is, it depends in part on the distribution of body weights in their social reference group (Burke & Heiland 2007, Burke et al. 2010b, Hammond & Ornstein 2014).

[7]The origin of 65 as the standard retirement age can be traced to the pension system devised by Bismarck in the 1880s. In fact, the original age in the German system was 70; later it was reduced to 65 (Cohen 1957, Myers 1973). Many European countries and the United States subsequently copied various features of the German system in their own retirement arrangements. Thus 65 served as a convenient reference point for those designing the retirement systems, and also for the individuals within these systems once 65 became embedded in the public consciousness.

I shall not restrict the concept of social norm to behaviors that have a moral or injunctive character, as is common in much of the literature (see, among others, Homans 1950; Elster 1989a,b; Coleman 1990; Bicchieri 2006). Some of the examples discussed under the previous headings are of this nature, but many—perhaps most—are held in place by a combination of mechanisms. The key to the analysis of norm dynamics is the existence of a positive feedback loop between social and individual behaviors, whether or not this feedback arises from the prospect of social punishment. This feedback loop has important consequences for the evolution of norms, as we shall see in the following sections.

3. KEY FEATURES OF NORM DYNAMICS

As mentioned at the outset, the purpose of this article is to review key elements of the theory of social norms and to discuss recent attempts to bring the theory to data. In the case of field data, it is often difficult to identify the effects of a norm due to spurious correlation and the presence of multiple feedback channels (Manski 1993, Brock & Durlauf 2001, Moffitt 2001). In spite of these difficulties, much can be said about the qualitative features of the dynamics that hold irrespective of the specific channels that are governing the feedback effects. These include the following.

3.1. Persistence

Norms tend to persist for long periods, and to respond very sluggishly to changes in external conditions that alter the benefits and costs of adhering to the norm. Persistence is a striking feature of certain contractual norms, such as customary shares for landowners and tenants in agricultural contracts, percentages charged by tort lawyers in malpractice and product liability cases, performance fees charged by fund managers, commissions for real estate agents, and so forth (see Section 7). Once established by precedent, such shares can persist for many years in spite of changing economic conditions. Norms of cooperation and mutual trust, or the lack thereof, can persist for centuries once they are established in a given society, as can norms that support persistent inequalities between social classes (Banfield 1958, Putnam 1993, Tabellini 2008, Guiso et al. 2009, Acemoglu & Robinson 2012, Belloc & Bowles 2013).

3.2. Tipping

When norm shifts do occur, the transition tends to be sudden rather than incremental (Schelling 1960, 1978). Once a crucial threshold is crossed and a sufficient number of people have made the change, positive feedback reinforces the new way of doing things, and the transition is completed rapidly.

3.3. Punctuated Equilibrium

In combination, persistence and tipping create a characteristic signature in the historical evolution of norms: There are long periods of no change punctuated by occasional bursts of activity in which an old norm is rapidly displaced by a new one. This is the *punctuated equilibrium effect* (Young 1998b).[8] The foot binding of women was a norm in China for many centuries, yet it disappeared

---

[8]In biology this term has a more specific interpretation; here I use it to describe the characteristic look of a dynamical path.

almost entirely within a single generation (see Section 8.1 below). In our own times, attitudes toward homosexual behavior have undergone a radical shift within the space of a few decades.

3.4. Compression

An important implication of positive feedback is that individual choices tend to exhibit less variation than would otherwise be expected. Burke & Young (2011) refer to this effect as "conformity warp"; here I shall call it *compression*. For example, if people did not use a socially constructed retirement age as a heuristic, the distribution of retirement ages would be more diffuse than it actually is. If there were no standard performance fee in the hedge fund industry, there would be a greater range of negotiated fees. If landowners and tenants did not gravitate toward normative contracts, such as the 50-50 division of output, one would see more diversity in contract terms, depending on underlying fundamentals such as the quality of land and labor inputs, risk aversion, and the value of the parties' outside options (Young & Burke 2001; see also Section 7).[9]

3.5. Local Conformity/Global Diversity

Norms are population-level equilibria in interactions that typically have multiple equilibria. They evolve through a process in which chance events play an important role; hence the norms in a given community at a given point in time are inherently unpredictable. When communities do not interact, or interact only occasionally, they may follow different evolutionary trajectories and thus have different norms even though in other respects they are similar (Young 1998b). In the United States, for example, the standard treatment for certain medical conditions differs markedly from one region to another even when there are no significant differences in the cost of performing these procedures (see Section 8.4).

4. EVOLUTIONARY MODELS OF NORM DYNAMICS

4.1. The Basic Framework

Consider a population of players who interact over time, where each interaction entails playing a game G. For expositional simplicity, I shall usually assume that G is a two-person coordination game, but much of the theory extends to other classes of games. Depending on the context, I shall sometimes assume that G is symmetric and that pairs of players are drawn at random from a single population to play the game. At other times, I shall assume that G is asymmetric and players are matched at random from disjoint populations that occupy different roles in the game. A background assumption is that the populations are sufficiently large that the actions of any single individual will typically have a small impact on the dynamics of the process as a whole. More generally, I shall assume the following:

1. People do not have full information about what is going on in society at large. Their information is limited to their own prior experience and a sample of experiences of others in their social group.
2. People typically choose myopic best responses given their information, but they may occasionally deviate for a variety of reasons, including inattention, experimentation, idiosyncratic beliefs, and unobserved payoff shocks.

---

[9]For econometric tests of the compression effect, see Graham (2008).

3. People interact at random with others in their social group, usually with some bias toward their geographical or social "neighbors."

This framework occupies a middle ground in the level of rationality attributed to the agents. It presumes that people are purposeful and that they adapt their behavior based on expectations of others' behavior. It does not presume that they know the overall process in which they are embedded, or that they attempt to manipulate the process through strategic, forward-looking behavior. In a large population where people interact at random, it would be rare for a given individual to be in a position to alter the dynamics within a reasonable length of time.[10]

Nevertheless, there are exceptions. When people interact in small close-knit groups, such as villages, religious communities, and small professional organizations, a single individual may be able to steer the evolutionary process toward a new norm within a fairly short period of time. This is particularly true if the individual is in a position to set a public example, such as a religious leader or village elder. Such "norm entrepreneurs" stand to enhance their status if the new norm is beneficial to the community (Sunstein 1996, Ellickson 2001, Acemoglu & Jackson 2015). Note, however, that being prominent can also be a liability, because such individuals have a lot to lose should their efforts fail: Norm entrepreneurship is a risky undertaking for the well-connected. For this reason, politicians, religious leaders, and village elders may be among the least willing to induce norm shifts even though they are the ones most capable of doing it. (One of the reasons that dueling persisted for so long is that it was entrenched among the upper classes; see Section 8.2.)

Although the small-group situation is important, there are many settings where the relevant group is sufficiently large and the evolutionary time frame sufficiently long that individuals are not in a position to have much influence on the dynamics. In these cases, myopic best response behavior serves as a plausible baseline assumption. This is not to say that all players behave in the same way: Differences in risk aversion, in amounts of information, and in the degree of sophistication (k-level reasoning) can all be treated within the same general framework (Young 1993b, 1998b; Hurkens 1995; Saez-Marti & Weibull 1999).

Before leaving the topic of rationality, I should remark that there is an important branch of the evolutionary games literature that attributes even less rationality to the agents. This is the biological model of the evolution of cooperation (Axelrod 1984, 1986). In this framework, people do not consciously optimize or form beliefs; rather they are endowed with particular strategies whose reproductive success hinges on how well they do in competition with alternative strategies. Given space limitations, I shall not elaborate on this approach here, except to note that the stochastic evolutionary framework can also be applied to this case (Foster & Young 1990, 1991).

In the remainder of this review, I shall use the perturbed best response framework to examine the following questions. First, under what conditions do the interactions of many dispersed individuals eventually coalesce into a social norm, that is, to a situation in which people's actions and expectations constitute an equilibrium at the level of the group? Second, are some norms more likely to emerge than others? Third, what are the characteristic features of the dynamics, both in the intermediate and in the long run? Fourth, what are the welfare implications of the outcomes they produce?

---

[10]Ellison (1997) explores conditions under which myopic versus forward-looking behavior is justified in large populations.

4.2. Stochastic Evolutionary Game Theory

To address these questions, we need to formalize the process by which individuals form expectations and update their behavior. To simplify the exposition, let us begin by considering a symmetric two-person game G and a single population of players who interact in pairs and play G. Later we shall see how to extend the framework to asymmetric games and multiple populations.

The set of available actions to each player will be denoted by X. Given a pair of actions (x, x'), let $u(x, x')$ denote the payoff to the first player and $u(x', x)$ the payoff to the second player. These payoffs define a symmetric two-person game in normal form. In most applications, the players will be embedded in a social or geographical space that determines the probability with which each pair will interact, and also the importance or "salience" of their interaction. To model these effects, let us assume that each pair of players (i, j) has an importance weight $w_{ij} \geq 0$. We shall also assume that the weights are symmetric, that is, $w_{ij} = w_{ji}$ for all $i \neq j$, and that $w_{ii} = 0$. Symmetry holds, for example, if w_{ij} measures the proximity of i and j in a suitably defined space. If the weights are binary, say $w_{ij} \in \{0, 1\}$, the pairs $\{i, j\}$ such that $w_{ij} = 1$ can be interpreted as the edges of an undirected network, all of which are equally weighted.

The dynamics can be modeled as follows. There is a population consisting of n players with identical utility functions. Time periods are discrete and denoted by $t = 1, 2, 3 \ldots$. At the end of period t, the state is an n-vector $\mathbf{x}^t$, where $x_i^t \in X$ is the current choice of action by each player i, $1 \leq i \leq n$. The state space is denoted by $\mathbf{X} = X^n$. Assume that the players update their strategies asynchronously: At the start of period $t + 1$, one agent, say i, is selected uniformly at random to update. Given the current choices of everyone else, $\mathbf{x}_{-i}^t$, define the utility of agent i from choosing action x as follows:

$$U_i\left(x, \mathbf{x}_{-i}^t\right) = \sum_j w_{ij} u\left(x, x_j^t\right). \tag{1}$$

This expression can be interpreted in several ways. Suppose, for example, that in each period one pair $\{i, j\}$ is drawn from the set of all pairs with probability $p_{ij} = w_{ij} / \sum_{1 \leq h < k \leq n} w_{hk}$ and they play the game G. In this case, $U_i(x, \mathbf{x}_{-i}^t)$ is proportional to i's expected utility from action x given the current actions of the other members of the population and the probability of encountering them. Alternatively, we could suppose that all pairs of individuals interact once in each period, and that w_{ij} is the weight that i attaches to payoffs from interacting with j. In either case, the functions $U_i(x, \mathbf{x}_{-i}^t)$ define an n-person game, which we shall call the *social game* induced by G and the weights w_{ij}.

To complete the specification of the model, assume that when agent i updates his or her choice, he or she chooses a new action $x_i^{t+1} = x$ with a probability that is increasing in $U_i(x, \mathbf{x}_{-i}^t)$. In particular, we shall usually assume that i's choice maximizes $U_i(x, \mathbf{x}_{-i}^t)$ with high probability, and that he or she chooses other actions with low probability. We shall call this a *perturbed best response model*.

In the evolutionary games literature, there are two benchmark models of perturbed best response behavior. The *uniform error model* posits a small error rate $\varepsilon \in [0, 1]$ such that with probability $1 - \varepsilon$, agent i chooses an action that maximizes $U_i(x, \mathbf{x}_{-i}^t)$, and with probability ε, he or she chooses an action uniformly at random (Ellison 1993; Kandori et al. 1993; Young 1993a,b). [If there are several actions that maximize $U_i(x, \mathbf{x}_{-i}^t)$, they are chosen with equal probability.]

An alternative approach is the *logit* or *log-linear response model* (Blume 1995; Durlauf 1997; Young 1998b; Blume & Durlauf 2001, 2003; Brock & Durlauf 2001; Young 2011; Kreindler & Young 2013, 2014). Let β be a nonnegative real number and suppose that i's probability of choosing action x at time $t + 1$ is given by the expression

$$P\left[x_i^{t+1} = x\right] = e^{\beta U_i\left(x, \mathbf{x}_{-i}^t\right)} / \sum_{y \in X_i} e^{\beta U_i\left(y, \mathbf{x}_{-i}^t\right)}. \tag{2}$$

This is a perturbed best response model in which the probability of deviating from the best response decreases the greater the corresponding loss in utility. As β becomes large, the probability of choosing a best response approaches 1. In the limiting case ($\beta = \infty$), a best response is chosen with probability 1.[11]

The preceding framework can be modified to allow for the possibility that players respond to the history of actions taken in prior periods, not just those taken in the immediately preceding period. We can also allow for the case of asymmetric games by supposing that players are drawn from multiple populations. To illustrate these extensions, consider a two-person normal-form game G with action space X for the row players and a possibly different action space Y for the column players. In each period t, one row player and one column player are drawn at random from two disjoint populations R and C to play the game G. Let $(x^t, y^t) \in X \times Y$ be the pair of actions they choose in period t. The *history* through period t can be represented as follows:

$$\mathbf{h}^t = \left(\left(x^1, y^1\right), \left(x^2, y^2\right), \ldots \left(x^t, y^t\right)\right). \tag{3}$$

In practice, people tend to put more weight on recent events than on distant ones. This *recency effect* can be modeled by truncating the history to the last m periods, where m (memory) is a positive integer. The updating process can be modeled as follows. When a player is selected to play at time $t + 1$, he or she draws a random sample from the actions of members of the other population in the previous m periods, then chooses a perturbed best response to the sample frequency distribution. For example, under the uniform error model, the player would choose a best response with probability $1 - \varepsilon$ and choose an action randomly with probability ε. Alternatively, the player might choose an action according to the log-linear response model discussed earlier.

4.3. Ergodic Distributions and Stochastic Stability

The models described above can be analyzed by adapting the theory of large deviations to the case of finite Markov chains (Freidlin & Wentzell 1984, Young 1993a). Here I shall restrict attention to the case where the state space is finite; in particular, the action spaces are finite, historical memory is finite, and the populations are finite. In this case, we obtain an irreducible Markov chain that has a unique ergodic distribution. When the departures from best response have very low probability, this distribution will typically be concentrated on a unique state (or, in the case of ties, on a small subset of states). These are called the *stochastically stable states* of the evolutionary process (Foster & Young 1990, Kandori et al. 1993, Young 1993a).

There are various techniques for computing the stochastically stable states (Kandori et al. 1993, Young 1993a, Blume 1995, Ellison 2000). An especially tractable case arises when (*a*) players respond only to the prior-period actions of others using a log-linear response function as in Equation 2 and (*b*) the game G is a two-person symmetric potential game. A potential game has the property that there is a function $\rho : X^2 \to R$ such that

---

[11]An alternative interpretation is that the agent chooses a best response with certainty and is subject to an unobserved payoff shock. When these shocks are extreme-value distributed, the resulting probabilities are given by the expression in Equation 2. This is a standard model in the discrete choice literature (McFadden 1974).

$$\forall x, x', y \in X, \rho(x', y) - \rho(x, y) = u(x', y) - u(x, y). \tag{4}$$

In other words, a unilateral change of action by one player leads to a change in that player's payoff that exactly equals the change in the potential function ρ (Monderer & Shapley 1996). A potential function for G can be used to define a potential function for the n-person game with payoff functions $U_i(\mathbf{x}) = \sum_j w_{ij} u(x_i, x_j)$ as follows:

$$\tilde{\rho}(\mathbf{x}) = \sum_{i,j} w_{ij} \rho(x_i, x_j). \tag{5}$$

Note that the maximum of the potential function $\tilde{\rho}(\mathbf{x})$ need not occur at a welfare-maximizing state. In particular, if G is a symmetric 2×2 coordination game, the potential is maximized when everyone plays the risk-dominant equilibrium, which need not be Pareto optimal.

The preceding framework can be extended to include heterogeneous preferences for taking different actions. Suppose that player i's payoff consists of two parts: i's expected payoff from playing a game G against others, $U_i(\mathbf{x}) = \sum_j w_{ij} u(x_i, x_j)$, and i's enjoyment of the action itself, say $v_i(x_i)$. Then i's total payoff in state $\mathbf{x}$ is

$$U_i^*(\mathbf{x}) = \sum_j w_{ij} u(x_i, x_j) + v_i(x_i). \tag{6}$$

It is straightforward to verify that if G has potential function ρ; then Equation 6 defines a game with potential function

$$\rho^*(\mathbf{x}) = \sum_{i,j} w_{ij} \rho(x_i, x_j) + \sum_i v_i(x_i). \tag{7}$$

> **Theorem 1:** The evolutionary process with the potential function in Equation 7 has a unique ergodic distribution, and the long-run probability of each state $\mathbf{x}$ is given by

$$\mu(\mathbf{x}) = e^{\beta \rho^*(\mathbf{x})} / \sum_{\mathbf{y} \in X} e^{\beta \rho^*(\mathbf{y})}. \tag{8}$$

> **Corollary 1:** The stochastically stable states of this process are precisely those states $\mathbf{x}$ that maximize the potential function $\rho^*(\mathbf{x})$.

I shall show how this framework can be applied to an empirical case in Section 7.

4.4. Qualitative Features of the Dynamics

Although the models discussed above differ in their details, they have similar qualitative properties. First, all of them exhibit the dichotomy of persistence and tipping: There are long periods in which people follow a norm (with some idiosyncratic variation), but every so often these tranquil periods are interrupted by a burst of activity in which one norm is displaced by another (the punctuated equilibrium effect). Second, the existence of a norm creates excess uniformity within a given community; in other words, people make less diverse choices than they otherwise would (the compression effect). Third, the evolutionary process is inherently stochastic and unpredictable; hence communities that do not interact may operate under different norms even though they

are exposed to the same influences and have similar socioeconomic characteristics. In other words, controlling for all community-level fundamentals, there may still be significant variance between communities (the local conformity/global diversity effect).

In the remainder of the article, I shall illustrate these ideas through a series of seven case studies. Two of them are based on historical narratives (foot binding and dueling), two on laboratory experiments (naming and bargaining games), and three on field data (norms of fertility, norms of medical practice, and contractual norms in agriculture). I treat three of these cases in some detail (naming, bargaining, and contracts) and the remaining four in synoptic form. There are, of course, many other empirical studies of social norms, some of which I list in the concluding section. I have chosen to concentrate on these few in order to highlight the complexities that must be dealt with in applying the theory to actual cases, and to illustrate how one can obtain useful insights in spite of data limitations. The cases also point to ways in which current theory can be extended.

5. NAMING

Our first case study is an experiment designed to show how naming conventions can arise in a large population through a trial-and-error learning process (Centola & Baronchelli 2015; see also Baronchelli et al. 2006). The basic setup is a pure coordination game: Two people are shown a picture of a face, and they simultaneously and independently suggest names for it. If they provide the same name, they earn a reward; otherwise they pay a small penalty. There is no restriction on the names that subjects can provide—this is left to their imagination.

Each trial consists of 25 rounds in which the same face is shown to all subjects in all rounds. The number of subjects in each trial ranges from 24 to 96. At the start of a given round, each subject is paired anonymously with one other subject, and they are given 20 seconds in which to provide a name. Pairs are drawn at random from the edges of a fixed, undirected network. The subjects do not know the structure of the network or the identities of the people they are paired with; they know only the names that their partners provided in previous rounds. Thus they accumulate information about the names that are currently popular among the people they are being paired with. This fact permits naming conventions to emerge spontaneously. Of particular interest is the effect the network structure has on the evolutionary process.

The dynamics can be modeled as follows. Each subject is located at the node of a fixed network (which is not known to the subjects). In each round, one edge of the network is selected at random, and the corresponding two subjects play the name game. At the end of the round, each learns the name chosen by his or her partner. The history through round t is a sequence of the form $h^t = (x^1, \delta^1), (x^2, \delta^2) \ldots (x^t, \delta^t)$, where x_i^t is the name player i offered in round t, $\delta^t = \left(\delta_{ij}^t \right)_{1 \le i < j \le n}$ is an indicator vector such that $\delta_{ij}^t = 1$ if players i and j were paired in round t, and $\delta_{ij}^t = 0$ otherwise. At the end of round t, player i knows the names provided by all partners with whom he or she was paired in rounds $1 \le t' \le t$.

Let us compare the results from two different types of networks: (a) a ring in which each player is adjacent to his or her nearest four neighbors and (b) the complete graph in which everyone is adjacent to everyone else. The top row of **Figure 1** shows a representative evolutionary path for the ring network. By round 10, several distinct local conventions have emerged in which small groups of nearby agents choose the same name. There are five local conventions altogether with contested border regions separating them. The bottom row of **Figure 1** shows a representative path for the complete graph. Compared to the ring network, there are more failures to coordinate in the initial rounds, but by round 10, a dominant convention has become established, and by round 24, this convention has displaced all of the others. Note that these patterns emerge without the subjects knowing anything about the overall structure of the network or the state of the system.
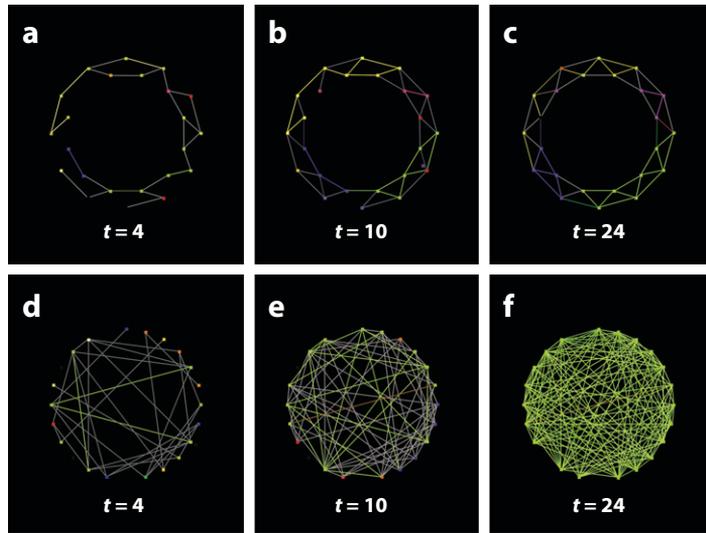
6. BARGAINING

The next case is an experimental study of the evolution of bargaining norms (Gallo 2014). There is a population of subjects called "buyers" who repeatedly play the Nash demand game against a single "Seller." In each period, each buyer plays the game once against the Seller. Simultaneously, the buyer and Seller demand shares of a fixed pie. If their demands do not exceed the total size of the pie, they receive rewards proportional to their demands; otherwise they get nothing (Nash 1950). Before making a demand, the buyer receives information about the demands made by the Seller in previous encounters with other buyers. The Seller is programmed to play a perturbed best response strategy to a sample of previous demands of the buyers, and the buyers know this.[12]

A key feature of the experiment is to vary the network through which buyers obtain information from other buyers, and to see what effect the network topology has on the trajectory of play. Specifically, the experiment permits an examination of the following questions. How often does the process converge to a norm, that is, a state in which all or almost all buyers demand the same amount irrespective of their position in the network? How many rounds on average does it take for convergence to occur? Does the network topology affect the resulting distribution of norms?

Gallo's (2014) experimental setup is as follows. A given "trial" involves a fixed group of six buyers plus the Seller.[13] The buyers are located at the nodes of a fixed network, which determines the channels of communication between them. Each trial consists of 50 rounds of play. In every round, each buyer is matched once with the Seller in the Nash demand game. The buyer learns only

---

[12]This is a variant of the evolutionary bargaining model proposed by Young (1993b). For extensions and variations of this model, see Saez-Marti & Weibull (1999), Binmore et al. (2003), and Bowles et al. (2010).

[13]The experiments were conducted at the Centre for Experimental Social Science at Nuffield College, University of Oxford. Subjects consisted of both undergraduate and graduate students from the university.
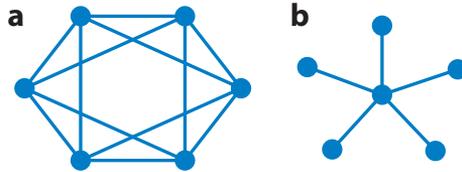
Figure 2

(*a*) Regular network and (*b*) star network.

whether his or her demand was compatible with the Seller's; the buyer is not told how much the Seller demanded. Demands are constrained to be nonnegative integers, and the total size of the pie is 17 units. Hence a pair of demands (x, y) is compatible if and only if x and y are nonnegative integers and $x + y \leq 17$. The number 17 was chosen to reduce the focal qualities of certain solutions; in particular, simple fractions such as one-half and one-third cannot be implemented in whole numbers.

When a given buyer b is about to make a demand, the buyer is told what demands were made by the Seller in a random sample of prior matches with b's neighbors in the previous six rounds. The idea is that each buyer learns about some of the Seller's prior demands through his or her network of contacts. The size of the buyer's sample is $2d_b$, where d_b is the number of neighbors to whom b is connected. From this information, the buyer can draw inferences about the demands the Seller is currently making; the buyer also knows from prior matches which of his or her own demands led to successful outcomes. When matched against any given buyer, the Seller samples from the prior demands made by all buyers within the last six rounds; the Seller then chooses a perturbed best response. Specifically, the Seller is programmed to choose a myopic best response to the frequency distribution of demands in the Seller's sample with 95% probability, and to make a demand at random with 5% probability.[14]

Consider the situation where the network is regular of degree 4 (see **Figure 2a**). Fix a particular buyer b. At the start of a given round $t \geq 7$, b receives a random sample of demands made by the Seller in interactions with b's neighbors over the previous six periods. Thus there are 24 such interactions in all, and the size of each sample is $2d_b = 8$. Given this information, the buyer chooses a demand. Simultaneously, the Seller generates a demand as follows: The Seller chooses a random sample of size 8 from all demands made by all buyers in the last six periods ($6 \times 6 = 36$). With probability 95%, the Seller chooses a best response to the sample frequency distribution, and with probability 5% he or she chooses a demand uniformly at random from the set of integers $\{1, 2, \ldots, 17\}$.

We shall say that the process converges to a bargaining norm if at least five out of six buyers demand the same amount for at least four consecutive periods. Convergence occurred in 11 out of 13 trials, and the resulting share for the buyers ranged from 9 to 14 (see **Figure 3**). Moreover, convergence was typically quite rapid: In the 11 trials where convergence occurred, the average waiting time was 33 rounds. The stochastic evolutionary models discussed in Section 4 do not speak directly to the question of how quickly convergence to a norm will occur in practice. This depends on the amount of "noise" in the subjects' responses, which was not the focus of this

---

[14]The Seller follows this procedure separately for each buyer; that is, the Seller resamples from the demands of all buyers and chooses a perturbed best response. The best response is computed from the sample distribution under the assumption of risk neutrality.
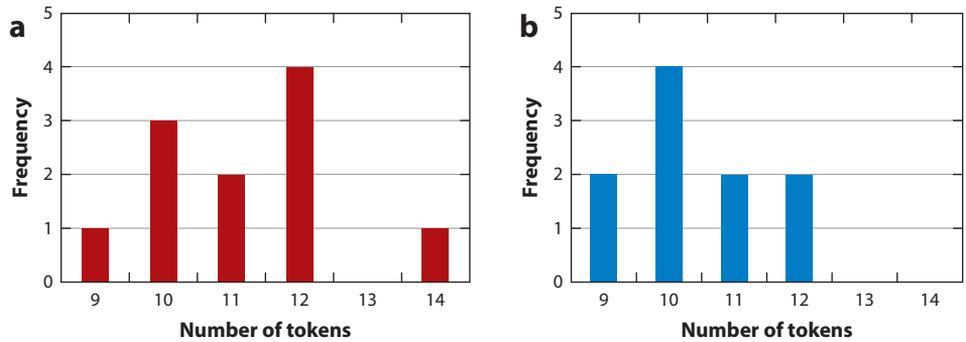
Figure 3

Frequency distribution of bargaining norms expressed in number of tokens for the buyers: (*a*) regular network and (*b*) star network.

experiment. Rather, the aim was to examine whether the structure of the buyers' network influences the norms to which they converge.

Stochastic evolutionary models make specific predictions about the effect of network structure in this situation. When subjects within a population have different sample sizes (and thus different amounts of information about the other side's actions), the stochastically stable norm is particularly sensitive to the minimum sample size. The smaller the minimum sample size within a given population, the less its members can expect to get, all else being equal. The reason is the following. The long-run stability of a norm depends on how easily it can be dislodged through random perturbations that cause a shift in expectations by one or both populations. An agent with a small sample size can be induced to lower his or her expectations given only a few instances of higher demands by the opposite population. By contrast, changing the expectations of an agent with a large sample size takes more high demands by the other side, which is less likely to occur.[15]

The results of Gallo's (2014) experiment agree with this prediction. Consider the star network shown in **Figure 2***b*. There are five agents with one neighbor each; hence in each period, they receive a sample of size 2. The central agent has degree 5 and a sample size of 10. The Seller's sample size remains as before (8). The theory suggests that, as a group, the buyers in this situation are less advantaged than in the regular network of degree 4. To test this prediction, Gallo estimates the amount by which buyers' demands in the regular network exceeded buyers' demands in the star network, averaged over the last 10 periods of play. He finds that the average demand in the regular network was approximately 10% larger than in the star network (significant at the 1% level). Moreover, this result continues to hold if one compares the average demands that were compatible, that is, did not lead disagreement.

These results suggest a number of directions for further theoretical and empirical research. In particular, it would be useful to know how frequently subjects actually deviate from myopic best reply, and whether this error rate tends to diminish as players gain more experience over the course of the game. Relatedly, it would be useful to extend the theory to accommodate situations where the error rates are time-varying and/or bounded away from zero, and to estimate how this affects the expected time it takes to reach equilibrium. For recent work along these lines, readers are referred to Young (2011) and Kreindler & Young (2013, 2014).

---

[15]This prediction continues to hold even when agents are risk averse (Young 1993b, 1998a,b).

7. CONTRACTING

The next case uses field data to examine contractual norms in agriculture, which is a leading example of principal-agent contracting. The logic of a share contract is that it spreads the risk between the contracting parties. Theory predicts that the terms of such a contract will depend on such factors as the agent's inherent skill, the principal's ability to monitor the agent's effort, their attitudes toward risk, and the value of their outside options (Cheung 1969, Stiglitz 1974, Bolton & Dewatripont 2005). To the extent that there is heterogeneity in these factors, there should be variation in the shares that the parties negotiate. In practice, however, share contracts often exhibit a high degree of uniformity, a phenomenon that has attracted the attention of economists since the foundation of the discipline (Smith 1776, Mill 1848). This "excess uniformity" is a feature not only of agricultural contracts, but of many other principal-agent contracts that are based on "usual and customary" shares, such as real estate commissions (Fisher & Yavas 2010) and contingency fees for tort lawyers (Kritzer 2004).[16]

The case of performance fees in the hedge fund industry is particularly well-documented (Mallaby 2010). Until very recently, the standard in the industry was to charge "two and twenty": 2% per year in management fees and 20% of the profits as a performance fee irrespective of the size of the fund. The 20% rate can be traced to Alfred Winslow Jones, who originated the concept of hedge funds in the 1940s and claimed as justification that this was the rate charged by Phoenician merchants in ancient times (Mallaby 2010, p. 30).

The evolutionary theory of norms provides a framework for understanding the phenomenon of customary shares. The thesis is that, once a particular way of sharing the output becomes established in a given business and a given locale, it provides a powerful focal point that suppresses much of the variation that would arise from idiosyncratic factors. Here I shall discuss the application of this idea to agricultural share contracts in the Midwestern United States, for which particularly good data are available.

An agricultural share contract is an arrangement in which a landowner and a tenant split the gross proceeds of each year's harvest in fixed proportions or shares. A key advantage of such a contract is that it shares the risk of an uncertain outcome while offering the tenant an incentive to increase the expected value of that outcome. When contracts are competitively negotiated, one would expect the size of the share to reflect such factors as the expected yield per acre, the risk aversion of the parties, their outside options, and so forth. In practice, however, the shares tend to cluster around "usual and customary" levels even when there are substantial and observable differences in the quality (expected yield) of different parcels of land, and there are different outside options for the tenants. Furthermore, these contractual norms tend to be anchored at prominent focal points, such as 1/2-1/2, 2/5-3/5, or 1/3-2/3. The importance of such focal points, and the fact that they tend to be specific to particular regions, has been documented in many parts of the world, including India, Africa, and the United States (Bardhan & Rudra 1980, Bardhan 1984, Robertson 1987, Young & Burke 2001).

Here I shall summarize a study of share contracts in the state of Illinois (Young & Burke 2001). This study is based on survey data collected by the Illinois Cooperative Extension Service (1995), which provide detailed information on the terms of contracts on several thousand farms in different parts of the state. This includes information on the size of the farm, the terms on which all inputs and outputs are shared, gross output, and the net incomes to the tenant and landlord. A key additional feature is a measure of each farm's inherent productivity (i.e., expected yield per acre)

---

[16]Based on a survey of contingency fees in Wisconsin, Kritzer (2004, p. 39) finds that in about 60% of the cases the shares are one-third for the lawyer and two-thirds for the client (except in situations where the fees are regulated by statute).

based on the soil types found on that farm. In particular, each farm is assigned a soil quality index, which is proportional to the expected yield per acre under standard management conditions.

There are three striking features of these data. First, over 98% of the share contracts were based on 1/2-1/2, 2/5-3/5, or 1/3-2/3. Second, there are strong regional differences in their frequency of use: In the northern part of the state, the customary share is 1/2-1/2, whereas in the southern part of the state, the customary share is either 1/3-2/3 or 2/5-3/5 for the landlord and tenant, respectively.[17] Furthermore, there is a high degree of uniformity within each region despite the fact that there are substantial differences in the productivities of individual farms in the region that would seem to call for different shares. This is an example of the local conformity/global diversity effect discussed in Section 3. Third, there are many instances in which farms have essentially the same soil quality, but when they are located in different regions, they often operate under substantially different shares that reflect differences in regional norms. Compression prevents the contractual shares from fully reflecting the underlying heterogeneity among farms.

These features are captured in **Figure 4**, which shows the distribution of shares in two representative counties (one in the north and one in the south), as a function of the soil quality of the farms in these counties. In the north (Tazewell County), the share is almost always 1/2-1/2 despite the fact that the highest-quality farms are almost twice as productive as the lowest-quality farms. In the south (Effingham County), 2/3 for the tenant is the most frequent contract, and there is somewhat greater variation in the contract terms; nevertheless, only three contracts are used with appreciable frequency.

The evolutionary models discussed in Section 3 can shed some light on these apparent anomalies. Let us identify each farm i with the vertex of a graph. Assume that the contract adopted at a given farm depends on three factors: the farm's inherent productivity (output per acre), local wages from nonagricultural employment, and the frequency with which the contract is used by others in the vicinity. Let s_i denote the expected output of farm i under a standard management regime, and let y_i denote the annual income available to the tenant from alternative employment in the vicinity of farm i. A share contract at i specifies a proportional share x_i for the tenant, and the complementary share $1 - x_i$ for the landlord, where $x_i \in [0, 1]$. Thus the tenant's expected annual income on farm i is $x_i s_i$. The *incentive-compatible* contracts are those that induce the tenant to accept farm employment instead of the alternative income y_i; that is, they satisfy the constraint $x_i s_i \geq y_i$. The rent to the landlord is $v_i(x_i) = (1 - x_i)s_i$.

To model the impact of local custom, let us suppose that there are n farms situated at the nodes of a graph Γ. Let E denote the set of edges in Γ. Two farms are said to be *neighbors* if there is an edge joining them. Assume for simplicity that all edge weights $w_{ij} = 1$; that is, all neighbors of a given farm i have an equal degree of social influence on the contracting parties at i. For analytical convenience, we shall assume that the set of feasible contracts $X \subset (0, 1)$ is finite. The subset $X_i = \{x_i \in X : x_i s_i \geq y_i\}$ consists of the incentive-compatible contracts at i.

A state of the process is a vector $\mathbf{x}(t) \in \prod_i X_i$ that specifies the contract in force at each farm at a given point in time t. For simplicity, we shall omit the dependence on t in the following discussion. Given a state $\mathbf{x} \in \prod_i X_i$, let $\delta_{ij}(\mathbf{x}) = \delta_{ji}(\mathbf{x}) = 1$ if i and j are neighbors and $x_i = x_j$; otherwise let $\delta_{ij}(\mathbf{x}) = 0$. Thus $\delta_{ij}(\mathbf{x})$ identifies the pairs of neighbors that are coordinated on

---

[17]This north-south division corresponds roughly to the southern boundary of the last major glaciation. In both regions, farming techniques are similar and the same crops are grown—mainly corn, soybeans, and wheat. In the north, the land tends to be flatter and more productive than in the south, though there is substantial variability within each of the regions.

a Tazewell
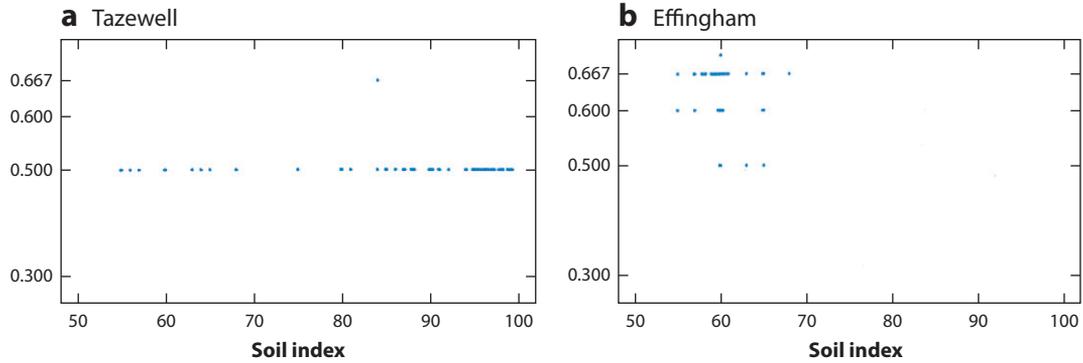
b Effingham

Figure 4

Distribution of tenant shares in (*a*) a northern county (Tazewell) and (*b*) a southern county (Effingham). Figure adapted from Young & Burke (2001).

the same contract in state **x**. Let us posit that the *propensity* to adopt a particular contract at *i* is an increasing function of its expected rent to the landlord and the degree to which it conforms to other contracts in *i*'s neighborhood subject to the incentive-compatibility constraint. Define an *adoption propensity function* as follows:

$$\forall x_i \in X_i, \quad \pi_i(x_i, \mathbf{x}_{-i}) = (1 - x_i)s_i + \gamma \sum_j \delta_{ij}(x_i, \mathbf{x}_{-i}). \tag{9}$$

The number $\gamma \geq 0$ is the *social conformity parameter*, which can be estimated from data. It captures the possibility that the terms of a given contract will tend to be pulled in the direction of customary practice in the area. The model also allows for the possibility that there is no such effect ($\gamma = 0$).

Note that the model specifies that conformity must be precise to have any benefit; in other words, if two contracts offer slightly different shares (say 48% versus 50%), the conformity payoff is nil. This is consistent with the hypothesis that these fractions are serving as focal points. Focal points have the peculiar property that nearby solutions are among the least focal (Schelling 1960, pp. 111–12). Thus the focal point hypothesis suggests that the distribution of shares will be sharply concentrated on particular values; the distribution will not be smooth as would be expected if they arise from unobserved heterogeneity.

Consider the following evolutionary dynamic: From time to time, contracts come up for renegotiation on each farm. Assume that these renegotiation times are described by Poisson arrival processes that are independently and identically distributed (i.i.d.) among farms. Thus, with probability 1, no two renegotiations occur simultaneously, and we can think of the process as evolving in discrete time periods with a single renegotiation occurring in each period. Let $\mathbf{x}^t$ denote the state of the process in period *t*, and suppose that the next renegotiation occurs at farm *i* at the start of period $t + 1$. We shall suppose that the probability of adopting different contracts at *i* is a log-linear function of the propensity to adopt; that is, for some $\beta > 0$,

$$\forall x \in X_i, \quad P[x_i^{t+1} = x] = e^{\beta \pi_i(x, \mathbf{x}^t_{-i})} / \sum_{y \in X_i} e^{\beta \pi_i(y, \mathbf{x}^t_{-i})}. \tag{10}$$

This is observationally equivalent to assuming that x_i^t maximizes the following perturbed propensity function:

$$\tilde{\pi}_i\left(x_i^t, \mathbf{x}_{-i}^t\right) = \left(1 - x_i^t\right)s_i + \gamma \sum_{j \neq i} \delta_{ij}\left(x_i^t, \mathbf{x}_{-i}^t\right) + \varepsilon_i^t, \tag{11}$$

where the shocks ε_i^t are i.i.d. extreme-value distributed with the cumulative distribution function $P(\varepsilon_i^t \leq z) = e^{-e^{-\beta z}}$. The resulting process has the following potential function:[18]

$$\forall \mathbf{x}, \ \rho(\mathbf{x}) = \sum_i (1 - x_i)s_i + (\gamma/2) \sum_i \sum_{j \neq i} \delta_{ij}(\mathbf{x}). \tag{12}$$

The first term, $\sum_i (1 - x_i)s_i$, represents the total rent to land, which we shall abbreviate by $r(\mathbf{x})$. The double sum $\sum_i \sum_{j \neq i} \delta_{ij}(\mathbf{x})$ is twice the total number of edges that are coordinated on the same contract. Denoting the number of coordinated edges by $c(\mathbf{x})$, we can write the potential function in the compact form

$$\rho(\mathbf{x}) = r(\mathbf{x}) + \gamma c(\mathbf{x}). \tag{13}$$

The ergodic distribution of this process has the Gibbs-Boltzmann form; that is, the long-run probability $\mu(\mathbf{x})$ of each state $\mathbf{x}$ is given by

$$\mu(\mathbf{x}) \propto e^{\beta[r(\mathbf{x}) + \gamma c(\mathbf{x})]}. \tag{14}$$

In particular, the log probability of each state $\mathbf{x}$ is a linear function of the total rent to land plus the degree of local conformity.

Given sufficiently rich event data, one could compute the probability that a given contract is adopted at a given farm i, conditional on the frequency of contracts in force at other nearby farms, and from this one could estimate γ and β (assuming one could control for common shocks). Note that this framework does not presume that conformity matters, but if $\gamma > 0$ at a high level of significance, this would provide support for the hypothesis that it does matter.

Young & Burke (2001) did not attempt such an estimation due to data limitations. However, the model makes other predictions that are corroborated by the Illinois data. One of these predictions is that the most likely state is one in which customs vary from one region to another based on the average soil quality in each region. Within each region, however, the same customary share will be used on farms with markedly different soil qualities.

What is the evidence that social norms are playing a role in the choice of share contracts, as opposed to a common unobserved factor? This issue bedevils all of the studies that rely on cross-sectional data. Complete identification in such cases is difficult (Manski 1993, Brock & Durlauf 2001, Moffitt 2001). However, there are several features of the data that strongly suggest that social norms are present. First, the distribution is concentrated on a small number of inherently focal fractions. If this concentration were due to a common factor, such as the reservation wage in a given region, it would be a strange coincidence that the outcomes are precisely 1/2-1/2, 2/5-3/5, or 1/3-2/3. Second, using cross-sectional data, one can test for compression: If the same share applies to farms of different quality, there should be an upward bias in the net income of tenants who work

---

[18]To verify that this is a potential function, it suffices to check that whenever an agent i undertakes a unilateral change of action, the change in i's propensity exactly equals the change in the potential.

on high-quality land. This prediction is confirmed by the Illinois data (Burke 2015). It is, of course, possible that the market equilibrates through assortative matching rather than through variation in contract terms; that is, high-quality farms attract high-quality tenants. If this is the case, total output per acre on high-quality farms should be higher than is predicted by their soil quality index, which was calibrated using a fixed level of capital and labor inputs (Odell & Oschwald 1970). The data do not support this hypothesis either: It appears that tenants on high-quality land are able to capture a portion of the land rent without a concomitant increase in labor productivity (Burke 2015).

The analytical framework described above is quite general and could be applied to many other types of contractual relationships—such as real estate commissions, lawyers' contingency fees, and hedge fund managers' bonuses—where norms may well play a role in determining contract terms.

8. OTHER CASES

This section provides an overview of four case studies of norm dynamics that further illustrate how the evolutionary theory of norms can be brought to bear on empirical cases. They differ in the types of data they use—historical, cross-sectional, experimental—and in the phenomena they study. In Section 8.1, we discuss the first case, Mackie's (1996) study of foot binding and infibulation, which focuses on the types of interventions that have successfully induced norm shifts, as well as top-down attempts that have failed to induce such shifts. The second case is an evolutionary model of dueling due to Jindani (2014), who shows how norms respond to exogenous changes in objective costs (in this case mortality rates) as well as changing attitudes that weaken the force of social sanctions (see Section 8.2). The third case, discussed in Section 8.3, is concerned with the impact of community norms on the use of contraception in developing countries, and with strategies for identifying norms using network data (Kohler et al. 2001, Munshi & Myaux 2006). The fourth case (Section 8.4) examines regional differences in standards of medical practice in the United States (Wennberg & Gittelsohn 1973, 1982; Chandra & Staiger 2007; Burke et al. 2010a). Although these differences are substantial, there is an ongoing debate about whether they result from different norms that become established in different medical communities, or whether they result from local productivity spillover effects. Both mechanisms produce similar dynamic feedback effects, but the implications for welfare are different.

8.1. Foot Binding

Foot binding is an ancient Chinese practice in which a young girl's toes are bent backward toward the heel and her feet are wrapped in tight bandages to prevent them from growing to normal size.[19] After 5–10 years of painful treatment, the result is a pair of tiny bowed feet that are only three to four inches long. A woman with bound feet cannot engage in most forms of manual labor and cannot walk very far. The custom appears to have originated under the Sung dynasty (960–1279). It was initially applied to concubines in the imperial court and subsequently spread to the upper classes as a sign of gentility. By the Ming dynasty (1368–1644), it had become common practice except among the lower classes where women were needed as laborers. The practice was thought to promote female fidelity, because a woman with bound feet was also housebound. From the standpoint of the girl's family, however, a key consideration was that it enhanced her marriage prospects. Foot binding was the equilibrium of a game in which boys' families preferred them to

---

[19]The discussion in this section is based on Mackie (1996).

marry girls with bound feet, and the girls' families dared not give up the practice for fear that they would not find good husbands.

This social norm remained in place for many centuries, yet it disappeared almost entirely within a single generation. Mackie (1996) argues that the norm unraveled due to a deliberate campaign by reformers to eradicate the practice. A key part of the story is that reformers did not rely on top-down edicts prohibiting the practice; these had been tried repeatedly in earlier times and had come to naught. Instead they organized "natural foot societies" in which families pledged not to bind their daughters and not to allow their sons to marry bound women. In addition, they conducted public campaigns explaining the adverse consequences of foot binding for health, mobility, and employability. In other words, they recognized that the key problem was to simultaneously shift the expectations of a group of interacting families, so that it would become rational for them not to continue the practice given its harmful effects and its decreased benefit in the local marriage market. Mackie argues that this approach could serve as a template for displacing other harmful norms, including the practice of female genital cutting in sub-Saharan Africa (Mackie 1996, Mackie & LeJeune 2009).

8.2. Dueling

From the Renaissance to the nineteenth century, dueling was a common practice among the upper classes in Europe (Nye 1993, Hopton 2007). A gentleman risked losing his honor if he failed to challenge someone for making offensive statements, or if he failed to accept such a challenge. Both parties were willing to risk their lives rather than face the prospective damage to their reputations by violating the norm. It is a stark example of how an extremely costly norm can be held in place by social pressure. Although dueling was practiced initially by the nobility, in the nineteenth century it had become common among members of the professional classes, especially politicians and journalists, who were very much in the public eye.

The norm of dueling is particularly interesting for several reasons. First, it is a highly complex norm with many interlocking parts, each of which can be construed as a subnorm operating within a larger normative framework. Once a challenge had been made, seconds were named whose function was first to try to reconcile the parties. Failing that, they negotiated a meeting time and place, the choice of weapons, and various other conditions (such as the number of paces if pistols were used). Second, the practice persisted despite repeated attempts to ban it. According to Hopton (2007), there were nine separate attempts to enforce a ban on dueling in France during the nineteenth century, none of which was successful. Thus it is an excellent example of how difficult it can be to change a norm—even a very costly one—from the top down. Indeed the very people who were best positioned to exercise leadership—legislators and opinion makers—were also highly susceptible to public pressure and had much to lose by attempting to change the status quo.

Third, the costliness of the norm (i.e., the fatality rate) was affected by changes in weapons technology. Before 1500, it was customary in England and France to use slashing broadswords and to protect oneself with bucklers. Although the swords caused injuries, they were mostly superficial flesh wounds, and the fatality rate was quite low. By the late 1500s, broadswords had been replaced by rapiers, which were more deadly, and the fatality rate increased substantially. Then in the eighteenth century, rapiers were replaced by pistols, which further increased the fatality rate.

A particularly interesting feature of the dynamics is that in France the norm adjusted by a return to rapiers together with rules of engagement that limited the risk of serious injury, whereas pistols continued to be the weapon of choice in England and the United States. Jindani (2014) proposes an evolutionary model of this process in which norm shifts are driven in part by changes in their

objective cost, and in part by broader changes in society that determine the social cost of flouting the norm. The model suggests that the differential evolution of the norm in England and France can be attributed to differences in objective costs (lower fatality rates in France), and to differences in social attitudes, which in France continued to support a code of honor that was a vestige of the ancien régime (Nye 1993, Hopton 2007).

8.3. Fertility

Birthrates in developing countries have been steadily declining due to the greater availability of contraceptives, family planning clinics, and enhanced educational and economic opportunities for women. These changes can be understood as a rational response by individuals to the perceived benefits and costs of contraception. Nevertheless, the timing and pace of the fertility transition in different communities have led many demographers to conclude that other factors are at work, including differences in social norms (Bongaarts & Watkins 1996; Montgomery & Casterline 1996; Entwhisle et al. 1996; Kohler 1997, 2000; Kohler et al. 2001; Behrman et al. 2002; Billari et al. 2009; Liefbroer & Billari 2010). In practice, however, it is extremely difficult to disentangle the impact of community norms from other factors, such as unobserved common effects, learning from others, and imitation.

Here I shall briefly outline the results of several recent papers that address these issues. Munshi & Myaux (2006) examine the rate of contraceptive use by groups that are located within the same community and that have similar socioeconomic characteristics, but that differ in their religious affiliation. The data report contraceptive use over an 11-year period by women in 70 villages in rural Bangladesh, where each village is composed of significant numbers of both Muslims and Hindus. Munshi & Myaux express each individual's decision to practice contraception as a function of (*a*) the proportion currently using contraception in their own religious group and (*b*) the proportion currently using contraception in the other religious group, controlling for individual characteristics such as age, education level, spouse's occupation, spouse's education level, and household assets. Their key finding is that a woman's contraception decision is positively associated with its prevalence within her own religious group, but there is no statistically significant relationship with its prevalence in the other group.

Although this result is consistent with a social norms explanation, there are other plausible possibilities. For example, the absence of cross-group effects could arise from a "social learning" process in which women gain information about the benefits and costs from the experiences of other women in their social network. Even though family planning clinics and contraceptives might have been available to all women in a given village at the same time, there could be differences in the rate of uptake due to the stochastic nature of the adoption process, which would result in different amounts of experiential information within the two groups at a given point in time.

Kohler et al. (2001) suggest a way of distinguishing between social learning and alternative mechanisms by examining the effect of network structure on the rate of contraceptive use. They define a woman's network as the set of women with whom she informally discusses family planning issues. The density of the network is the proportion of links that the members have with each other. Thus if a woman's network contains n individuals, the maximum possible number of links between them is $n(n-1)/2$. Kohler et al. hypothesize that, for a fixed network size, increasing density provides greater scope for applying social pressure to members of the group while providing relatively little new information due to redundancy.

These hypotheses lead to the following predictions. Holding the network size fixed and controlling for socioeconomic characteristics (age, education, number of children, etc.), women in denser networks will be less likely to adopt contraception given the proportion of other network

a Obisa

b Owich, Kawadhgone, and Wakula South
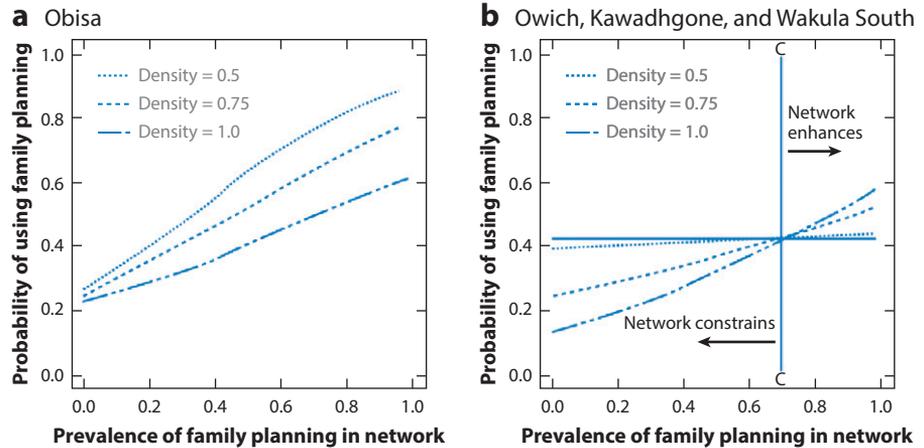
Figure 5

Effect of contraceptive prevalence on the probability of adoption as a function of network densities: (*a*) Obisa and (*b*) Owich, Kawadhgone, and Wakula South. Figure reproduced with permission from Kohler et al. (2001).

members who have adopted (because of redundancy in information), and women in denser networks will show a greater tendency to conform to the dominant practice in their group, whether the dominant practice is to use contraception or to avoid using it.

The authors test these two predictions using a detailed survey of approximately 500 women living in rural Kenya. They find a significant difference in behavior between those living in regions with isolated villages and low market activity (Owich, Kawadhgone, Wakula South) and another region (Obisa) where women have access to a large and active market. **Figure 5** shows the probability that a given woman is an adopter conditional on the proportion in her network who are adopters and on the density of the network. In all regions, sparser networks are associated with higher probabilities of adoption when the proportion of adopters is less than two-thirds. This is consistent with the first hypothesis, namely, that for a given level of adoption, a sparse network provides more independent information than does a dense network. In the nonmarket regions, however, a reversal occurs when the proportion of adopters exceeds a critical threshold that is slightly above two-thirds (the estimated value is 71%). Above this value, higher densities increase the probability that a given individual will adopt, all else being equal. In other words, dense networks initially inhibit the adoption of contraception, but beyond a certain threshold they solidify support for it.

Of course, this analysis does not rule out other factors, such as social learning, nor does it rule out the possibility that dense networks result from a tendency of individuals to interact with others like themselves (homophily). Stronger inferences could be drawn from longitudinal data that give the timing of individual adoption decisions conditional on adoptions by members of one's social network and controlling for individual fixed effects, but such data are difficult to obtain on a large scale.[20]

---

[20]Behrman et al. (2002) analyze longitudinal data on contraceptive use in Kenya. They show that social interactions matter after controlling for individual fixed effects, but they do not analyze network density data as in the study by Kohler et al. (2001). Other methods for distinguishing between social learning and social conformity are discussed in Young (2009) and Brown (2013).

8.4. Medical Treatment

A long-standing puzzle for public health researchers in the United States is that the treatment of certain medical conditions varies substantially between states, and even between counties within the same state. Differences have been documented for a wide variety of treatments, including Caesarean sections, tonsillectomies, beta-blockers, and stents (Wennberg & Gittelsohn 1973, 1982; Chandra & Staiger 2007; Burke et al. 2010a). These variations remain substantial even after controlling for differences in the availability and cost of medical technologies as well as other factors that might predict differences in treatment intensity (Phelps & Mooney 1993).

It has been suggested that these local norms of medical practice are due in part to peer effects (Burke et al. 2007, 2010a). Physicians tend to respond more readily to practice recommendations made by local "opinion leaders" than to impersonal recommendations such as professional practice guidelines and results of clinical trials (Soumerai et al. 1998, Bhandari et al. 2003). Burke et al. (2010a) incorporate these peer effects into a dynamic model of treatment choice as follows. They assume that a physician's choice of treatment for a given patient is guided in part by an assessment of its efficacy for patients of that type (depending on age, sex, and medical condition), and in part by its frequency of use by others in the local medical community. They show that in such a setting regional treatment norms emerge. Moreover, the model predicts that the norm in each region will typically reflect the (objectively) best treatment for the dominant patient type in that region. For example, if treatment A is better suited to older patients than treatment B, one would expect to see A as the norm in regions with a high proportion of older people, and B in regions with a high proportion of younger people. The logic is analogous to contractual norms in agriculture, as discussed in Section 7: The customary share for the tenant is higher in regions where the land quality is poor on average, and is lower where the land quality is high on average. Within each region, however, the existence of a norm suppresses variation that might otherwise arise from differences in the land quality of individual farms (the compression effect).

It should be stressed, however, that compression does not necessarily imply a loss of welfare. In the case of medical practice, positive feedbacks may result from learning and productivity spillovers, in which case the welfare effects of local treatment norms may be positive (Chandra & Staiger 2007; Wennberg & Gittelsohn 1973, 1982). As physicians become particularly adept in the dominant local treatment, patients who are well-suited to the treatment will benefit from their doctors' added expertise. However, patients who are better suited to an alternative treatment face the prospect of receiving either a substandard application of that alternative or an expert application of the ill-suited, dominant treatment. Unless these patients can be transferred to a location where there are experts in the alternative treatment, as a second best they are better off receiving the dominant local treatment.

More generally, this case illustrates the importance of disentangling the various channels through which positive feedback effects may be operating. It would not be surprising if both social conformity and productivity spillovers are factors in the evolution of regional norms of medical practice. A challenge for future research is to tease out the relative strength of these two effects, and to examine their welfare implications.

9. DIRECTIONS FOR FUTURE RESEARCH

My aim in this article has been to outline the main ideas in evolutionary game theory and how they relate to the evolution of social norms. The theory provides a comprehensive framework for analyzing the dynamics that result when individuals make choices based on conventional economic factors as well as on the norms of behavior within their social group. The dynamics exhibit certain

features that are not captured by conventional equilibrium models that omit social feedback effects. One such feature is persistence, that is, the tendency of a norm to stay in place for long periods of time in spite of exogenous changes in incentives. A second key feature is local conformity/global diversity: Similar communities may exhibit different norms due to chance events. Third, the very existence of a norm diminishes the variation in behaviors that would otherwise arise. This compression effect can be tested with sufficiently rich cross-sectional data, as we saw in the case of agricultural contracts. A similar analysis could be carried out for other types of contracts based on customary percentages, such as the fees earned by fund managers, real estate agents, and tort lawyers.

The theory shows how social norms can coalesce from out-of-equilibrium conditions with no top-down direction. It also demonstrates that the outcomes need not be optimal: Evolutionary forces can lead to dysfunctional norms that persist for long periods of time, a prediction that is confirmed by historical evidence. On a more positive note, the theory suggests that certain kinds of interventions will be more effective than others for promoting norm shifts. Traditional instruments, such as the blanket application of taxes or subsidies, are often insufficient to overcome the social penalties for violating a prevailing norm. In some cases, it might be more effective to alter the behaviors of a few key actors in small close-knit groups, thus leveraging the positive feedback effects from social interactions.

Although the cases provide considerable support for the predictions of theory, they also point to areas where the theory needs further refinement. For example, current models generally assume that the error rates are held at a low and constant level. Experimental results suggest that error rates can be quite high initially, and that they change over time depending on how successful people are in coordinating. These effects have important implications for the rate at which an evolutionary process approaches equilibrium and for the equilibria that are most likely to be selected. Another shortcoming of current theory is that it focuses almost exclusively on long-run behavior (for which powerful selection results exist) and tends to neglect the intermediate-run dynamics, which are also of great interest.

I have selected a few cases illustrating how evolutionary game theory can be brought to bear on empirical and experimental evidence, but there are many other examples. These include norms stigmatizing welfare and unemployment (Lindbeck et al. 1999, Brügger et al. 2009), norms stigmatizing out-of-wedlock births (Akerlof et al. 1996), norms of revenge (Elster 1990), tax compliance (Wenzel 2004, 2005; Nicolaides 2014), corruption (Fisman & Miguel 2006), veiling (Carvalho 2013), littering (Cialdini et al. 1990), executive compensation (Brown & Young 2014), body weight (Burke & Heiland 2007, Burke et al. 2010b, Hammond & Ornstein 2014), number of children (Liefbroer & Billari 2010), property rights (Platteau 2000, Bowles & Choi 2013), and marriage (Kanazawa & Still 2001).

DISCLOSURE STATEMENT

ACKNOWLEDGMENTS

LITERATURE CITED

Acemoglu D, Jackson MO. 2015. History, expectations and leadership in the evolution of social norms. *Rev. Econ. Stud.* 82:423–56

Acemoglu D, Robinson JA. 2012. *Why Nations Fail: The Origins of Power, Prosperity, and Poverty*. New York: Crown Bus.

Akerlof GA. 1980. A theory of social custom, of which unemployment may be one consequence. *Q. J. Econ.* 94:749–75

Akerlof GA, Yellen JL, Katz L. 1996. An analysis of out-of-wedlock childbearing in the United States. *Q. J. Econ.* 111:277–317

Axelrod R. 1984. *The Evolution of Cooperation*. New York: Basic

Axelrod R. 1986. An evolutionary approach to norms. *Am. Polit. Sci. Rev.* 80:1095–111

Axtell R, Epstein JM. 1999. Coordination in transient social networks: an agent-based computational model of the timing of retirement. In *Behavioral Dimensions of Retirement Economics*, ed. HJ Aaron, pp. 161–83. Washington, DC: Brookings Inst.

Banfield EC. 1958. *The Moral Basis of a Backward Society*. Chicago: Free

Bardhan P. 1984. *Land, Labor, and Rural Poverty: Essays in Development Economics*. New York: Columbia Univ. Press

Bardhan P, Rudra A. 1980. Terms and conditions of sharecropping contracts: an analysis of village survey data in India. *J. Dev. Stud.* 16:287–302

Baronchelli A, Felici M, Loreto V, Caglioti E, Steels L. 2006. Sharp transition towards shared vocabularies in multi-agent systems. *J. Stat. Mech.* 2006:P06014

Behrman JR, Kohler H-P, Watkins SC. 2002. Social networks and changes in contraceptive use over time: evidence from a longitudinal study in rural Kenya. *Demography* 39:713–38

Belloc M, Bowles S. 2013. The persistence of inferior cultural-institutional conventions. *Am. Econ. Rev.* 103:93–98

Bernheim D. 1994. A theory of conformity. *J. Polit. Econ.* 102:841–77

Bhandari M, Devereauz PJ, Swiontkowski MF, Schemitsch EH, Shankardass K, et al. 2003. A randomized trial of opinion leader endorsement in a survey of orthopedic surgeons: effect on primary response rates. *Int. J. Epidemiol.* 32:634–36

Bicchieri C. 2006. *The Grammar of Society: The Nature and Dynamics of Social Norms*. Cambridge, UK: Cambridge Univ. Press

Billari F, Philipov D, Testa MR. 2009. Attitudes, norms and perceived behavioural control: explaining fertility intentions in Bulgaria. *Eur. J. Popul.* 25:439–65

Binmore K. 1994. *Game Theory and the Social Contract*, Vol. 1: *Playing Fair*. Cambridge, MA: MIT Press

Binmore K. 2005. *Natural Justice*. New York: Oxford Univ. Press

Binmore K, Samuelson L, Young HP. 2003. Equilibrium selection in bargaining models. *Games Econ. Behav.* 45:296–328

Blume LE. 1995. The statistical mechanics of best-response strategy revision. *Games Econ. Behav.* 11:111–45

Blume LE, Durlauf SN. 2001. The interactions-based approach to socioeconomic behavior. See Durlauf & Young 2001, pp. 15–44

Blume LE, Durlauf SN. 2003. Equilibrium concepts for social interaction models. *Int. Game Theory Rev.* 5:193–209

Bolton P, Dewatripont M. 2005. *Contract Theory*. Cambridge, MA: MIT Press

Bongaarts J, Watkins SC. 1996. Social interactions and the contemporary fertility transitions. *Popul. Dev. Rev.* 22:639–82

Bowles S. 2004. *Microeconomics: Behavior, Institutions, and Evolution*. Princeton, NJ: Princeton Univ. Press

Bowles S, Choi J-K. 2013. Coevolution of farming and private property during the Early Holocene. *PNAS* 110:8830–35

Bowles S, Hwang H-S, Naidu S. 2010. Evolutionary bargaining with intentional idiosyncratic play. *Econ. Lett.* 109:31–33

Boyd R, Richerson PJ. 2002. Group beneficial norms can spread rapidly in a structured population. *J. Theor. Biol.* 215:287–96

Boyd R, Richerson PJ. 2005. *The Origin and Evolution of Cultures*. New York: Oxford Univ. Press

Brock W, Durlauf SN. 2001. Discrete choice with social interactions. *Rev. Econ. Stud.* 68:235–60

Brown LM. 2013. *The diffusion of innovations: empirical and experimental evidence*. PhD Diss., Univ. Oxford

Brown LM, Young HP. 2014. *The diffusion of a social innovation: executive stock options from 1936–2005*. Work. Pap., Dep. Econ., Univ. Oxford

Brügger B, Lalive R, Zweimüller J. 2009. *Does culture affect unemployment? Evidence from the Röstigraben*. Discuss. Pap. 4283, IZA, Bonn

Burke MA. 2015. *The distributional effects of contractual norms: the case of cropshare agreements*. Work. Pap., Fed. Reserve Bank Boston

Burke MA, Fournier GM, Prasad K. 2007. The diffusion of a medical innovation: Is success in the stars? *South. Econ. J.* 73:588–603

Burke MA, Fournier GM, Prasad K. 2010a. Geographic variations in a model of physician treatment choice with social interactions. *J. Econ. Behav. Organ.* 73:418–32

Burke MA, Heiland F. 2007. Social dynamics of obesity. *Econ. Inq.* 45:571–91

Burke MA, Heiland F, Nadler CM. 2010b. From 'overweight' to 'about right': evidence of a generational shift in body weight norms. *Obesity* 18:1226–34

Burke MA, Young HP. 2011. Social norms. In *Handbook of Social Economics*, Vol. 1A, ed. A Bisin, J Benhabib, MO Jackson, pp. 311–38. Amsterdam: Elsevier

Burtless G. 2006. Social norms, rules of thumb, and retirement: evidence for rationality in retirement planning. In *Social Structures, Aging, and Self-Regulation in the Elderly*, ed. KW Schaie, LL Carstensen, pp. 123–60. New York: Springer

Carvalho J-P. 2013. Veiling. *Q. J. Econ.* 128:337–70

Centola D, Baronchelli A. 2015. The spontaneous emergence of conventions: an experimental study of cultural evolution. *PNAS* 112:1989–94

Chandra A, Staiger D. 2007. Productivity spillovers in health care: evidence from the treatment of heart attacks. *J. Polit. Econ.* 115(1):103–40

Cheung SNS. 1969. *The Theory of Share Tenancy*. Chicago: Univ. Chicago Press

Cialdini R, Kallgren C, Reno R. 1990. A focus theory of normative conduct: a theoretical refinement and reevaluation of the role of norms in human behavior. *Adv. Exp. Soc. Psychol.* 24:201–34

Cohen WJ. 1957. *Retirement Policies Under Social Security*. Berkeley: Univ. Calif. Press

Coleman JS. 1987. Norms as social capital. In *Economic Imperialism: The Economic Approach Applied Outside the Field of Economics*, ed. G Radnitzky, P Bernholz, pp. 133–55. New York: Paragon

Coleman JS. 1990. *Foundations of Social Theory*. Cambridge, MA: Harvard Univ. Press

Durlauf SN. 1997. Statistical mechanical approaches to socioeconomic behavior. In *The Economy as a Complex Evolving System*, Vol. 2, ed. WB Arthur, SN Durlauf, D Lane, pp. 81–104. Menlo Park, CA: Addison-Wesley

Durlauf SN, Young HP, eds. 2001. *Social Dynamics*. Cambridge, MA: MIT Press

Edgerton RB. 1992. *Sick Societies: Challenging the Myth of Primitive Harmony*. New York: Free

Ellickson RC. 1991. *Order Without Law: How Neighbors Settle Disputes*. Cambridge, MA: Harvard Univ. Press

Ellickson RC. 2001. The evolution of social norms: a perspective from the legal academy. See Hechter & Opp 2001, pp. 35–75

Ellison G. 1993. Learning, local interaction and coordination. *Econometrica* 61:1047–71

Ellison G. 1997. Learning from personal experience: one rational guy and the justification of myopia. *Games Econ. Behav.* 19:180–210

Ellison G. 2000. Basins of attraction, long run stochastic stability and the speed of step-by-step evolution. *Rev. Econ. Stud.* 67:17–45

Elster J. 1989a. *The Cement of Society*. Cambridge, UK: Cambridge Univ. Press

Elster J. 1989b. Social norms and economic theory. *J. Econ. Perspect.* 3(4):99–117

Elster J. 1990. Norms of revenge. *Ethics* 100:862–85

Entwhisle B, Rindfuss RD, Guilkey A, Chamratrithirong A, Curran SR, Sawangdee Y. 1996. Community and contraceptive choice in rural Thailand: a case study of Nang Rong. *Demography* 33:1–11

Fehr E, Fischbacher U. 2004a. Third party punishment and social norms. *Evol. Hum. Behav.* 25:63–87

Fehr E, Fischbacher U. 2004b. Social norms and human cooperation. *Trends Cogn. Sci.* 8:185–90

Fehr E, Fischbacher U, Gächter S. 2002. Strong reciprocity, human cooperation, and the enforcement of social norms. *Hum. Nat.* 13:1–25

Fisher LM, Yavas A. 2010. A case for percentage commission contracts: the impact of a "race" among agents. *J. Real Estate Finance* 40:1–13

Fisman R, Miguel E. 2006. *Cultures of corruption: evidence from diplomatic parking tickets*. NBER Work. Pap. 12312

Foster DP, Young HP. 1990. Stochastic evolutionary game dynamics. *Theor. Popul. Biol.* 38:219–32

Foster DP, Young HP. 1991. Cooperation in the short and in the long run. *Games Econ. Behav.* 3:145–56

Freidlin M, Wentzell A. 1984. *Random Perturbations of Dynamical Systems*. Berlin: Springer-Verlag

Gallo E. 2014. *Communication networks in markets*. Work. Pap. Econ. 1431, Univ. Cambridge

Gambetta D. 2009. *Codes of the Underworld: How Criminals Communicate*. Princeton, NJ: Princeton Univ. Press

Gigerenzer G, Todd PM, ABC Res. Group. 1999. *Simple Heuristics That Make Us Smart*. New York: Oxford Univ. Press

Graham BS. 2008. Identifying social interactions through conditional variance restrictions. *Econometrica* 76:643–60

Guiso L, Sapienza P, Zingales L. 2009. *Long term persistence*. NBER Work. Pap. 14278

Hammond R, Ornstein J. 2014. A model of social influence on body weight. *Ann. N. Y. Acad. Sci.* 1331:34–42

Hechter M, Opp K-D, eds. 2001. *Social Norms*. New York: Russell Sage Found.

Hofbauer J, Sigmund K. 1998. *Evolutionary Games and Population Dynamics*. Cambridge, UK: Cambridge Univ. Press

Homans GC. 1950. *The Human Group*. New York: Harcourt Brace

Hopton R. 2007. *Pistols at Dawn: A History of Duelling*. London: Portrait

Hume D. 1888 (1739). *A Treatise of Human Nature: Being an Attempt to Introduce the Experimental Method of Reasoning into Moral Subjects*, ed. LA Selby-Bigge. New York: Oxford Univ. Press

Hurkens S. 1995. Learning by forgetful players. *Games Econ. Behav.* 11:304–29

Illinois Cooperative Extension Service. 1995. *Cooperative Extension Service Farm Leasing Survey*. Dep. Agric. Consum. Econ., Coop. Ext. Serv., Univ. Ill. Champaign-Urbana

Jindani S. 2014. *A game-theoretic approach to dueling and other social norms*. MPhilos. Thesis, Dep. Econ., Oxford Univ.

Kahneman D, Tversky A, Slovic P, eds. 1982. *Judgment Under Uncertainty: Heuristics and Biases*. Cambridge, UK: Cambridge Univ. Press

Kanazawa S, Still MC. 2001. The emergence of marriage norms: an evolutionary psychological perspective. See Hechter & Opp 2001, pp. 274–304

Kandori M. 1992. Social norms and community enforcement. *Rev. Econ. Stud.* 59:63–80

Kandori M, Mailath G, Rob R. 1993. Learning, mutation, and long-run equilibria in games. *Econometrica* 61:29–56

Kohler H-P. 1997. Learning in social networks and contraceptive choice. *Demography* 34:369–83

Kohler H-P. 2000. Fertility decline as a coordination problem. *J. Dev. Econ.* 63:231–63

Kohler H-P, Bereman JR, Watkins SC. 2001. The density of social networks and fertility decisions: evidence from South Nyanza District, Kenya. *Demography* 38:43–58

Kreindler GE, Young HP. 2013. Fast convergence in evolutionary equilibrium selection. *Games Econ. Behav.* 80:39–67

Kreindler GE, Young HP. 2014. Rapid innovation diffusion in social networks. *PNAS* 111:10881–88

Kritzer HM. 2004. *Risks, Reputations, and Rewards: Contingency Fee Legal Practice in the United States*. Stanford, CA: Stanford Univ. Press

Lewis D. 1969. *Convention: A Philosophical Study*. Cambridge, MA: Harvard Univ. Press

Liefbroer AC, Billari FC. 2010. Bringing norms back in: a theoretical and empirical discussion of their importance for understanding demographic behavior. *Popul. Space Place* 16:287–305

Lindbeck A, Nyberg S, Weibull J. 1999. Social norms and economic incentives in the welfare state. *Q. J. Econ.* 114:1–35

Mackie G. 1996. Ending footbinding and infibulation: a convention account. *Am. Sociol. Rev.* 61:999–1017

Mackie G, LeJeune J. 2009. *Social dynamics of abandonment of harmful practices: a new look at the theory*. Work. Pap. 9009-06, Innocenti Res. Cent., UNICEF, Florence

Mallaby S. 2010. *More Money Than God*. London: Bloomsbury

Manski C. 1993. Identification of endogenous social effects: the reflection problem. *Rev. Econ. Stud.* 60:531–42

McFadden D. 1974. Conditional logit analysis of qualitative choice behavior. In *Frontiers in Econometrics*, ed. P Zarembka, pp. 105–42. New York: Academic

Mill JS. 1848. *Principles of Political Economy*. London: Longmans Green

Moffitt RA. 2001. Policy interventions, low-level equilibria, and social interactions. See Durlauf & Young 2001, pp. 45–82

Monderer D, Shapley LS. 1996. Potential games. *Games Econ. Behav.* 14:124–43

Montgomery MR, Casterline JB. 1996. Social influence, social learning, and new models of fertility. *Popul. Dev. Rev.* 22:151–75

Munshi K, Myaux J. 2006. Social norms and the fertility transition. *J. Dev. Econ.* 80:1–38

Myers RJ. 1973. Bismarck and the retirement age. *The Actuary*, April

Myerson R, Weibull J. 2015. Tenable strategy blocks and settled equilibria. *Econometrica*. In press

Nash J. 1950. The bargaining problem. *Econometrica* 18:155–62

Nicolaides P. 2014. *Tax compliance, social norms and institutional quality: an evolutionary theory of public good provision*. Tax. Pap. 46-2014, Eur. Comm.

North DC. 1990. *Institutions, Institutional Change, and Economic Performance*. Cambridge, UK: Cambridge Univ. Press

Nye RA. 1993. *Masculinity and Codes of Male Honor in Modern France*. New York: Oxford Univ. Press

Odell RT, Oschwald WR. 1970. *Productivity of Illinois soils*. Circ. 1016, Coll. Agric. Coop. Ext. Serv., Univ. Ill., Urbana-Champaign

Phelps CE, Mooney C. 1993. Variations in medical practice use: causes and consequences. In *Competitive Approaches to Health Care Reform*, ed. RJ Arnould, RF Rich, W White, pp. 139–78. Washington, DC: Urban Inst.

Platteau J-P. 2000. *Institutions, Social Norms, and Economic Development*. New York: Routledge

Posner EA. 2000. *Law and Social Norms*. Cambridge, MA: Harvard Univ. Press

Putnam RD. 1993. *Making Democracy Work: Civic Traditions in Modern Italy*. Princeton, NJ: Princeton Univ. Press

Robertson AF. 1987. *The Dynamics of Productive Relationships: African Share Contracts in Comparative Perspective*. Cambridge, UK: Cambridge Univ. Press

Saez-Marti M, Weibull JW. 1999. Clever agents in Young's bargaining model. *J. Econ. Theory* 86:268–79

Samuelson L. 1997. *Evolutionary Games and Equilibrium Selection*. Cambridge, MA: MIT Press

Sandholm W. 2010. *Population Games and Evolutionary Dynamics*. Cambridge, MA: MIT Press

Schelling TC. 1960. *The Strategy of Conflict*. Cambridge, MA: Harvard Univ. Press

Schelling TC. 1978. *Micromotives and Macrobehavior*. New York: Norton

Skyrms B. 1996. *Evolution of the Social Contract*. Cambridge, UK: Cambridge Univ. Press

Skyrms B. 2004. *The Stag Hunt and the Evolution of Social Structure*. Cambridge, UK: Cambridge Univ. Press

Smith A. 1776. *An Inquiry into the Nature and Causes of the Wealth of Nations*. London: Strahan & Cadell

Soumerai SB, McLaughlin TJ, Gurwitz JH, Guadagnoli E, Hauptman PJ, et al. 1998. Effect of local medical opinion leaders on quality of care for acute myocardial infarction: a randomized controlled trial. *JAMA* 279:1358–63

Stiglitz JE. 1974. Incentives and risk sharing in sharecropping. *Rev. Econ. Stud.* 41:219–55

Sugden R. 1986. *The Economics of Rights, Cooperation and Welfare*. Oxford: Basil Blackwell

Sunstein CR. 1996. Social norms and social roles. *Columbia Law Rev.* 96:903–68

Tabellini G. 2008. Institutions and culture. *J. Eur. Econ. Assoc.* 6:255–94

Ullmann-Margalit E. 1977. *The Emergence of Norms*. New York: Oxford Univ. Press

Vega-Redondo F. 1996. *Evolution, Games, and Economic Behavior*. New York: Oxford Univ. Press

Wennberg J, Gittelsohn A. 1973. Small area variations in health care delivery. *Science* 182:1102–8

Wennberg J, Gittelsohn A. 1982. Variations in medical care among small areas. *Sci. Am.* 182:1102–8

Wenzel M. 2004. An analysis of norm processes in tax compliance. *J. Econ. Psychol.* 25:213–28

Wenzel M. 2005. Motivation or rationalization? Causal relations between ethics, norms, and tax compliance. *J. Econ. Psychol.* 26:491–508

Young HP. 1993a. The evolution of conventions. *Econometrica* 61:57–84

Young HP. 1993b. An evolutionary model of bargaining. *J. Econ. Theory* 59:145–68

Young HP. 1995. The economics of convention. *J. Econ. Perspect.* 10(2):105–22

Young HP. 1998a. Conventional contracts. *Rev. Econ. Stud.* 65:773–92

Young HP. 1998b. *Individual Strategy and Social Structure: An Evolutionary Theory of Institutions*. Princeton, NJ: Princeton Univ. Press

Young HP. 2009. Innovation diffusion in heterogeneous populations: contagion, social influence, and social learning. *Am. Econ. Rev.* 99:1899–924

Young HP. 2011. The dynamics of social innovation. *PNAS* 108:21285–91

Young HP, Burke MA. 2001. Competition and custom in economic contracts: a case study of Illinois agriculture. *Am. Econ. Rev.* 91:559–73