Journal of Economic Perspectives—Volume 10, Number 2—Spring 1996—Pages 105–122

The Economics of Convention

H. Peyton Young

Conventions regulate much of economic and social life, yet they have received surprisingly little attention from economists.[1] By a *convention,* we mean a pattern of behavior that is customary, expected and self-enforcing. Everyone conforms, everyone expects others to conform, and everyone has good reason to conform because conforming is in each person's best interest when everyone else plans to conform (Lewis, 1969). Familiar examples include following rules of the road, adhering to conventional codes of dress and using words with their conventional meanings. Conventions with direct economic implications include species of money and credit, industrial and technological standards, accounting rules and forms of economic contracts. Indeed, it would scarcely be an exaggeration to say that almost all economic and social institutions are governed to some extent by convention.

The main feature of a convention is that, out of a host of conceivable choices, only one is actually used. This fact also explains why conventions are needed: they resolve problems of indeterminacy in interactions that have multiple equilibria. Indeed, from a formal point of view, we may *define* a convention as an equilibrium that everyone expects in interactions that have more than one equilibrium.

The economic significance of conventions is that they reduce transaction costs. Imagine the inconvenience if, whenever two vehicles approached one another, the

[1] Important exceptions are Sugden (1986) and Ullman-Margalit (1977). The interested reader might also refer to the symposium on social norms that appeared in the Fall 1989 issue of this journal, with articles by Sugden (1989) and Elster (1989).

■ *H. Peyton Young is Professor of Economics, Johns Hopkins University, Baltimore, Maryland.*

drivers had to get out and negotiate which side of the road to take. Or consider the cost of having to switch freight from one type of railroad car to another whenever a journey involves both a wide-gauge and a narrow-gauge railroad line. This was a common circumstance in the nineteenth century and not unknown in the late twentieth: until recently, Australia had different rail gauges in the states of New South Wales and Victoria, forcing a mechanical switch for all trains bound between Sydney and Melbourne.

Conventions are also a notable feature of legal contracts. People rely on standard leases, wills, purchasing agreements, construction contracts and the like, because it is less costly to fill in the blanks of a standard contract than to create one from scratch. Even more important, such agreements are backed up by legal precedent, so the signatories have greater confidence that their terms are enforceable.

We may discern two ways in which conventions become established. One is by central authority. Following the French Revolution, for example, it was decreed that horse-drawn carriages in Paris should keep to the right. The previous custom had been for carriages to keep left and for pedestrians to keep right, facing the oncoming traffic. Changing the custom was symbolic of the new order: going on the left had become politically incorrect because it was identified with the privileged classes; going on the right was the habit of the common man and therefore more "democratic."[2]

In Britain, by contrast, there seems to have been no single defining event that gave rise to the dominant convention of left-hand driving. Rather, it grew up by local custom, spreading from one region to another. This is the second mechanism by which conventions become established: the gradual accretion of precedent.

The two mechanisms are not mutually exclusive, of course. Society often converges on a convention first by an informal process of accretion; later it is codified into law. In many countries, rules of the road were not legislated until the nineteenth century, but by this time the law was merely reiterating what had already become established custom.

This paper examines the process by which conventions emerge through the second mechanism—the accumulation of precedent. The intuitive idea is that, as an interaction is repeated over and over by many different individuals, one particular way of resolving the interaction gains an edge through chance. This lends it greater prominence, which means that more people hear about it, which leads to more people using it, and so forth. A positive feedback loop is created. Eventually one way of doing things drives out the others, not because it is inherently better, but because historical circumstances gave it an early lead that allowed it to pull ahead of the rest.[3]

[2] See Hopper (1982), Hamer (1986) and Lay (1992, pp. 199–200).

[3] A similar process has been proposed for the diffusion of technologies that exhibit networking externalities (Katz and Shapiro, 1985; David, 1985; Arthur, Ermoliev and Kaniovski, 1987; Arthur, 1989). The more people who choose IBM-compatible personal computers over Macintoshes, for example, the more

In this paper I propose to analyze the process of forming conventions by an evolutionary "bottom up" model. The paper begins with a game-theoretic description of how individuals make choices at the micro level. Unlike traditional game theory, however, we will not assume that people have unlimited powers of calculation and foresight. Instead we shall be dealing with people who are boundedly rational and take the world around them as given. Over time their dispersed and uncoordinated decisions interact in ways that they do not necessarily foresee. The question is whether these interactions eventually converge to a convention. If so, how long does it take? Can similar societies converge to different conventions? Are some conventions more likely to be observed than others? We shall sketch out a model that can be used to answer these questions and then offer some applications.

Right or Left

Let's begin with a simple example. It is the middle of the eighteenth century, the place is the English countryside, and two horse-drawn carriages are rapidly approaching one other from opposite directions. (Since the roads are bad, both are probably hogging the middle.) The choices for the coachmen are to pass on the left or on the right. Assume for the moment that there is no established convention or law to help them decide what to do. This interaction can be viewed as a one-shot game with payoffs of the following form:

	L	R
L	1,1	0,0
R	0,0	1,1

This game has three equilibria: both go left, both go right or both randomize with 50–50 probability between left and right.

Classical game theory does not give a coherent account of how people would play a game like this. The problem is that there is nothing in the structure of the game itself that allows the players—even purely rational players—to deduce what they ought to do. How then do people (rational or otherwise) actually play such games? The answer is that they rely on contextual cues that lie outside the game per se that allow them to coordinate on a particular equilibrium. Schelling (1960)

likely it is that software will be developed for IBMs as compared to Macintoshes, and thus the more worthwhile it becomes to own an IBM. A similar story can be told for the layout of typewriter keyboards, VHS versus Beta for videocassette recorders, cassettes versus compact discs for prerecorded music and so forth. When the diffusion of such technologies occurs through many individual, decentralized decisions, the process can be modelled in much the same way as the adoption of conventions. When diffusion is governed by the strategic choices of a few large firms, however, the process must be analyzed differently. See the symposium on network externalities that appeared in the Spring 1994 issue of this journal, with articles by Katz and Shapiro (1994), Besen and Farrell (1994) and Liebowitz and Margolis (1994).

called these cues *focal points.* The relevant focal points cannot be defined a priori; they depend on the coordination problem at hand and on the culture in which the players are embedded. Schelling observed that people are usually remarkably good at knowing what the salient focal points are even though they may be quite idiosyncratic to a given situation. However, he did not explain how they come into being. This is the problem that this paper seeks to address.

In the case of our two coachmen, it is reasonable to suppose that the relevant focal points are provided by precedent. They know from personal experience (and from talking with others) whether coaches tend to pass on the left or on the right. If their information is that coaches usually pass on the left, then it makes sense to go left; if they usually pass on the right, it makes sense to go right. Note that this reasoning does not require them to be especially rational. Each operates on the assumption that the world around them is going to continue today much as it did yesterday, which makes information about precedents relevant. Nor do they need to conjecture *why* people are doing what they are doing. They simply assume that others are going to follow previous patterns. Note also that they are typically making their decisions on the basis of very incomplete information. A coachman cannot know in practice what happened in all previous encounters between two coaches; he knows about a few cases, which lead him to develop expectations about what the oncoming coach is going to do.

An Evolutionary Model of Convention Formation

Let us now step back from this example and consider the problem of convention formation more generally. A number of individuals face a one-shot game that is played repeatedly by different players drawn from a large population. The payoffs may depend on who is playing, but the rules and strategies of the game are fixed. We assume that the players are boundedly rational and only partially informed. They do not have perfect foresight, they do not know in detail the structure of the process they are engaged in, and they do not know why other players are acting the way they are. Instead, they use simple rules of thumb to adjust their behavior based on their information about what other players are currently doing or have done in previous periods. Furthermore, there are unexplained variations in their behavior that play a role analogous to mutations in biological evolution. These three elements—local interaction between individuals, boundedly rational responses to the perceived environment, and random perturbations—define an "evolutionary game dynamic."

To be concrete, we shall describe the evolutionary process in terms of a two-person coordination game, though in fact the model extends easily to general n-person games (Young, 1993a). In each time period, one pair of players is drawn at random from the population to play the game. We shall assume for simplicity that all pairs are equally likely to be drawn, though this is not essential. (Later we shall consider variations in which players only meet their neighbors.) The expectations of the players are governed by the history of the process up to that time, where the "history" is simply a list of all pairs of agents who have played so far and the actions that they took.

Players obtain information about the history in two ways—from their personal experiences and from information they pick up from others. For example, a coachman, in talking with his friends at the local inn, hears about encounters they had with oncoming carriages in the recent past. In addition, he may have had some encounters himself, but let's assume for the moment that he has not been driving for awhile or is new on the job. He also may hear about incidents that occurred years ago, but disregards these as being outdated.

Among the encounters he knows about, suppose that more than half the carriages attempted to take the right side of the road. Our coachman then predicts that, when he next meets a carriage on the road, the probability is better than 50–50 that it will go right. Given this expectation, it is best for him to go right also (assuming that the payoffs are symmetric between left and right).

This behavioral rule can be generalized as follows. Let m (memory) be the maximum number of time periods that an agent looks back into the past. Assume for simplicity that in each period there is exactly one encounter and that each person hears about a certain number of prior encounters through word of mouth. A convenient way to model this way of acquiring information is to suppose that the agent draws a random sample from the last m encounters. The sample size s reflects the size of the agent's informational network and is an inherent property of the agent, not a result of an optimal search. Based on this information, he calculates the observed frequency distribution of left and right, and uses this to predict the probability that the next carriage he meets will go left or right. He then chooses a best reply. We shall call this behavioral rule "best reply to recent sample evidence."

While this rule seems reasonable enough, it does not seem likely that agents would always follow it (or any other) rule exactly. There will always be unexplained variation in their behavior—ways in which our predictions fail to capture what agents actually do. We model this unexplained behavior by a random variable. Let's say that with some small probability ε an agent chooses randomly between left and right, while with probability $1 - \varepsilon$ he adopts a best reply to recent sample evidence. In reality, of course, there will almost always be a specific reason why an agent acts the way he does. But this is like saying there is always a reason why a coin lands heads instead of tails—air currents, initial spin, height of throw and so forth. Since these effects cannot be known or predicted in advance, it is appropriate to model their aggregate impact as a random variable.

To illustrate concretely how such an evolutionary process works, let people's potential memory reach back 10 periods ($m = 10$), and suppose that each agent knows about only three of the past 10 encounters ($s = 3$). Finally, let's assume that there is one chance in a hundred that an agent chooses randomly, so that $\varepsilon = 1/100$. Suppose that people managed to coordinate on left for each of the last 10 periods. Then we have the following situation, where decisions of one player are on the top row and decisions of the other are on the bottom row, and time is moving from left to right:

LLLLLLLLL

LLLLLLLLL

This pattern represents a convention because it perpetuates itself when agents always choose best replies given their information. The reason is that, no matter what agents arrive on the scene next period, and no matter what particular precedents they know about from the above history, left is necessarily a best reply. It is an "absorbing state" in the sense that, once reached, the process never leaves it.

When the possibility of mistakes is introduced, however, the process has no absorbing states; there is always a chance that someone will take an unconventional action. For example, if society is in the left convention, several agents may nevertheless try to go right (with possibly disastrous results). If enough instances of right accumulate through random variation, the process will tip into the basin of attraction of the right convention, which then gains ground purely by best-reply dynamics. In this particular example, it suffices that three agents choose right in order to tip the process into the right convention. For example, we might obtain the following configuration by chance:

LLLLLLLRR

LLLLLLLLR

Now suppose a newcomer enters the scene in the next period and samples three of the previous encounters at random. By chance these might include the two most recent encounters and one other. Thus, out of the six carriages he has heard about, three went to the right. This leads him to conclude that left and right driving are equally popular. (As a newcomer, he has no experience of his own to modify this conclusion.) This leaves him indifferent between going left and right, so we can say that there is a 50–50 probability he will choose right the next time he meets an oncoming coach. However, the same reasoning holds for the agent who is driving toward him: there is a positive probability that he will also sample the last two encounters and choose right in response. Note that neither is making a mistake in doing so; both are acting rationally based on their information. Thus with positive probability we obtain the following situation in the next period:

LLLLLLRRR

LLLLLLLRR

Observe that the left-most pair L, L has been deleted because each record disappears after it is more than 10 periods old. At this point it should be clear that, within eight more periods, there is a positive probability of reaching a situation where for the preceding 10 periods that lie within memory, both players have always chosen the right side. This is the right-hand convention. In other words, three idiosyncratic choices ("mistakes") were sufficient to create a situation in which rational choices after that point could, with positive probability, tip the process out of the left convention and into the right convention. The crucial point is that, when the probability ε of making a mistake is small, and mistakes are made independently, the probability of tipping out of one convention and into the other is on the order of ε^3. This is because all the subsequent events needed to reach the opposite convention have much higher probability than the original three mistakes when ε is sufficiently small.

Simplified as this account may be compared to the decision situations that people actually face, it does capture three crucial elements of such decisions: they are based on limited information, the decision makers are boundedly rational, and there is unexplained variation in their behavior.

Moreover, the model yields several important qualitative predictions. It says that even if we know the initial state of society, we cannot predict what the prevailing convention will be at each future date. To put the matter differently, if two societies start off under similar initial conditions, there is a positive probability that at any given future time they will be operating different conventions. Such a process is said to be *path dependent.*

Unlike some of the path-dependent processes hypothesized in the literature on technological adoption, however, this model does not exhibit lock-in. It is path dependent in the short run but not in the long run. The reason is that, in our model, past decisions are eventually forgotten. This reduces the inertia to change and means that every so often society "flips" from one convention to another.[4] These incidents are triggered by repeated stochastic shocks to the system that we have modeled as unexplained behaviors. Over a long period of time, these shocks generate an ergodic process; in other words, there is statistical regularity in the frequency with which different states are observed *independent of the initial conditions.*

It can be shown that, if all agents have a positive probability of interacting, if they have sufficiently incomplete information and if random deviations have sufficiently low probability, then most of the time most of the population will be using the same convention (Young, 1993a; Kandori, Mailath and Rob, 1993).[5] We call

[4] In some models of technological adoption, it is assumed that products never wear out and adoption decisions are once and for all (Arthur, Ermoliev and Kaniovski, 1987; Arthur, 1989). Under this assumption, perturbations in the process ultimately cannot overcome the weight of the past (which grows indefinitely over time), and it locks into an outcome in which one product takes all of the market with probability one. In our case, finite memory corresponds to the idea that products have a finite life. Under this assumption, persistent stochastic shocks prevent lock-in over the long run.

[5] The Kandori, Mailath and Rob (1993) model is similar to the one presented here except that agents in their model do not have partial information or a long memory. They choose a best reply to the

this the *local conformity effect.* It does not say that the world is in equilibrium *all* of the time, but within a community of interacting agents it is close to equilibrium *most* of the time.

Now compare several communities that are alike in every way except that they do not interact with one another (for example they might be located on separate islands). If we run the process in each starting from similar initial conditions, then at any sufficiently distant future time there is a positive probability that they will be using different conventions. This is the *global diversity effect.*

While a convention tends to remain in place for a long period of time once it is established, it will eventually be dislodged by a series of random shocks. Society then careens toward a new convention, which also tends to remain in force for a long time. Similar dynamics are observed in biological evolution where random mutations occasionally give rise to new species that displace the old ones in a rush. This is known as the *punctuated equilibrium effect.*

Is there any evidence that conventions actually do follow a dynamic process characterized by these three effects? In the remainder of the paper we shall consider evidence from two quite different sources—the evolution of rules of the road and norms in distributive bargaining.

Rules of the Road

Before about 1750, there was relatively little vehicular traffic on the roadways of Europe, and what traffic there was tended to keep to the center to avoid falling into the ditch.[6] To the extent that left-right conventions existed at all, they appear to have been highly localized. The earliest mention of definite rules is for major bridges. In the year 1300, for example, Pope Boniface VIII issued an edict that pilgrims to Rome should keep to the left when crossing the Ponte Sant'Angelo. In 1736, Saxony decreed that traffic should keep to the right on a new bridge over the River Elbe. In 1756, the English Parliament passed a law requiring vehicles on London Bridge to keep left.

Whether seeded by such local decrees, or arrived at by a process of individual trial and error, these local rules gradually diffused more widely, until they were established at the regional and then national level. At this point we can think of nations, rather than individuals, as the principal actors in the convention game. If a nation is surrounded by countries who are following a different convention, it may have an incentive to switch to a rule that conforms with its neighbors.

---

frequency distribution of all actions in the preceding period with high probability and choose randomly among actions with low probability. The qualitative dynamics of this process are very similar to those described in the text.

[6] This account is based on Lay (1992). See also Hopper (1982), Kincaid (1986) and Hamer (1986). I am indebted to Jean-Marcel Goger and Arnulf Gruebler for historical information on France and Austria, respectively.

The surprising fact is, however, that until the end of the eighteenth century, the dominant convention in much of western Europe was for horse-drawn carriages to keep *left.* Great Britain, France, Sweden, Portugal, Austria, Hungary, Bohemia and parts of Italy and Germany drove left, and some continued to do so until well into this century. How then did continental Europe evolve to mostly right-hand driving today?

The answer is through a chain of historical accidents that gradually tipped the balance in favor of right. As we have already mentioned, France switched to right-hand driving after the Revolution for symbolic reasons. Napoleon then adopted this convention for his armies and imposed it in some of the countries he occupied. This set a train of events in motion that eventually caused other countries to switch. Portugal, which shares a border with right-driving Spain, converted after the First World War. The western-most provinces of Austria (Vorarlberg and Tyrol, which are sandwiched in between Germany, Switzerland and Italy) adopted the right-hand rule in the nineteenth century. The rest of Austria followed suit province by province in the 1920s, though the changeover was not completed until the *Anschluss* with Germany in 1938. (Hungary and Czechoslovakia also switched to the right-hand rule under German occupation.) Historically, Italy was a jumble of conflicting customs, with traffic in some cities keeping to the left, and traffic in much of the countryside keeping to the right. This situation was not straightened out until the 1930s. By the 1960s, only Sweden (among continental European countries) was driving on the left, and it switched over in 1967. While each of these cases has its particular explanation, the general impetus for change is clear: the more countries who adopt the same convention, the more likely it is that the others will follow suit.

One way to think about this process is to introduce a neighborhood structure into our earlier model.[7] The players in the game are now countries, which have specific geographical locations. Let us represent each country by a point and draw an edge between two countries if they share a common border (see Figure 1). On each edge there is cross-border traffic, and a cost is incurred by the two countries on either side of the border if they have conflicting rules of the road. Suppose that in each time period one country is drawn at random, and it asks itself whether it should switch. Let's assume that it does switch provided that its total costs, summed over all the edges adjacent to it, would go down. (We ignore the cost of switching due to the need for new infrastructure; this changes the details of the following discussion but not its substance.)

What are the equilibrium configurations of such an adjustment process? One is the situation in which all countries are using the same convention, as in Figure 1. But this is not the only possibility. For example, one could have an equilibrium in which two different conventions hold sway in regions that are weakly

[7] The first spatial model of an evolutionary type was proposed by Schelling (1971) to describe neighborhood segregation patterns. For recent work on spatial evolutionary game dynamics see Blume (1993), Ellison (1993), Anderlini and Ianni (1993), Berninghaus and Schwalbe (1992), Goyal and Jenssen (1993) and An and Kiefer (1993).

Figure 1
A Schematic Map of Countries (Nodes) and Border Crossings (Edges)

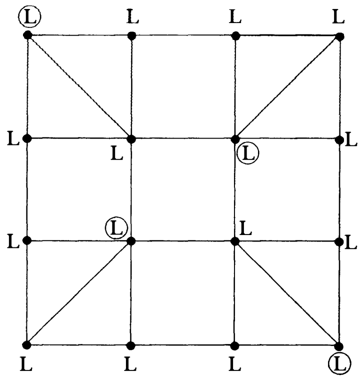

Figure 2
A Bottleneck Separating Two Regions Using Different Conventions

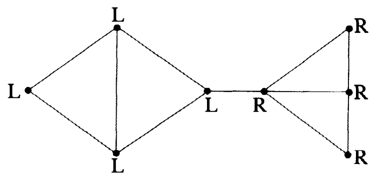

connected, as shown in Figure 2. It is also possible for two conventions to coexist even if the interconnections are quite dense. This situation is shown in Figure 3, where the countries form two blocks—one driving left, the other going right. If we assume there is an equal amount of traffic on each edge, there is no incentive for any one country to switch, because each is adjacent to more countries using the same convention than to countries using the opposite one. This shows how conflicting conventions can exist side by side.

The story changes, though, if we suppose that countries occasionally switch convention at random, that is, for reasons outside the model. (The French Revolution is an example.) In this case the process exhibits punctuated equilibrium behavior: an established configuration can give way to another when enough random changes accumulate to push the process out of the basin of attraction of one equilibrium and into the basin of attraction of another. In Figure 3, for example, it takes at least two switches to destabilize the equilibrium: if the circled countries shift to right-hand driving, then it is rational for the countries above and below it to switch to right-hand driving also, and the right convention eventually spreads across the board.

Figure 3
The Left Convention Changes to Right When the Circled Countries Switch

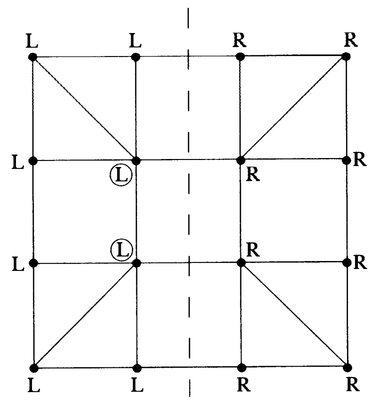

Compare this to the situation in which, at first, *all* countries are using the left convention and then a few switch to the right. The reader may check that it takes at least four switches (for example, the circled ones in Figure 1) to drive the left convention off the board.

This argument suggests the general proposition that it is easier to destabilize side-by-side conventions than to destabilize a uniform convention; hence a uniform convention will be more stable over the long run. With some additional hypotheses, this assertion can be proved rigorously using techniques borrowed from statistical mechanics (Blume, 1993). Suppose that the probability of a country switching to a new convention is proportional to ε^p, where p equals the difference in cost between the new convention and what the country is doing now, and "cost" is proportional to the amount of traffic crossing into neighboring countries with conflicting conventions. A country is therefore less likely to switch if the cost of doing so is high, but there is some probability it will switch anyway for idiosyncratic reasons. Finally, let us assume that the graph is connected, that is, one can drive from any country to any other. It can be shown that when ε is small, the probability that all countries are following the same convention is substantially higher over the long run than the probability of a side-by-side equilibrium in which some use one convention and some the other. Moreover, this result holds for any two competing conventions (not just left-right driving) when the agents' decisions are governed only by what their neighbors are doing (An and Kiefer, 1993). Of course this result addresses only the long-run behavior of the adjustment process; in the short to medium run, less stable transient patterns may well be observed.

While the model described above is quite simple, the evolutionary pattern of rules of the road is consistent with the qualitative effects it predicts. First, local communities and regions usually have a single, well-established convention (local conformity). Second, communities that do not interact may operate under different conventions (global diversity). Third, an established convention is not locked in forever; if we wait long enough, it will eventually be displaced by chance events that tip the process into a new regime (punctuated equilibrium). Indeed, prompted by the French Revolution, we have probably been witnessing just such a shift from left to right driving over the last couple of hundred years.

Norms in Distributive Bargaining

We turn now to distributive bargaining, a different arena in which conventions play an important role. When two people bargain over their shares of a common "pie," in what proportions do they divide it? In this case, the evolutionary model predicts quantitative as well as qualitative outcomes.

Let us first recall what classical theory has to say about the problem. In the model advanced by Nash (1950, 1953), the outcome of a distributive bargain depends principally on the alternatives of the parties and their aversion to risk. A person who is risk neutral and has a good fallback position will tend to get a better deal than someone who is risk averse and has unattractive alternatives, all else being equal.[8] To the extent that people vary in their opportunities and preferences, the theory predicts that we can expect to see considerable variation in the shares that they negotiate.

However, this is often not what we see in practice. What is missing from the above account are social norms and customs that point the bargainers toward particular solutions. By a *distributive bargaining norm* we mean shares that are customary and expected for the parties in a given bargaining situation. The effect of these norms is to lend greater uniformity and predictability to bargaining outcomes than classical theory predicts. For example, if two people are given a sum of money to divide in a laboratory setting, they almost invariably divide it 50–50 (Nydegger and Owen, 1974). But if they earn it by contributing different amounts of effort or skill, they may divide according to some other norm (for example, in proportion to their contributions).[9]

More generally, experimental evidence shows that bargainers can be conditioned by the experimental setup to expect a particular outcome, which is then reinforced as it is played over and over again. This solution becomes a norm that the subjects rationalize afterwards as being "fair," even though it may be idiosyncratic to the particular history of play experienced by a particular group (Roth, 1987; Binmore, 1991; Binmore, Swierzbinski, Hsu and Proulx, 1993).

The same holds in real-world bargaining situations: norms condition the parties to expect certain outcomes that depend on the bargaining context. In the United States, for example, a lawyer who takes a malpractice suit on a contingency basis conventionally gets one-third of the award, and the client gets

[8] For a detailed discussion of the Nash model and variations on it, see Luce and Raiffa (1957), Binmore and Dasgupta (1987) and Osborne and Rubinstein (1990).

[9] For experimental setups that elicit a variety of norms, see Yaari and Bar-Hillel (1981).

two-thirds. Real-estate agents customarily get 6 percent of the sale price of a house. Of course, people use such sharing contracts because they give agents an incentive to work diligently on behalf of their clients. But there seems to be no economic justification for the *particular* shares that we observe. If there were, we would expect the shares to fluctuate like prices under changing supply and demand conditions, but this is not the case. Rather, the shares appear to be conventions that got an early start through chance events and remain in place because they succeed in conditioning the parties' expectations. (There may also be implicit collusion in some of these markets that helps keep the norm in place.)

A particularly telling example occurs in sharecropping. A sharecropping contract specifies the fractions of the harvest for the landowner and the laborer who works his land. Such contracts are widely used in agricultural societies, both primitive and advanced, and they have been extensively studied by agricultural economists. One of the reasons why the parties favor sharecropping contracts is that they divide the risk of crop failure more evenly than a pure wage or rental contract would (Stiglitz, 1974). However, the issue that concerns us here is not the *form* of the contract but its *terms.* A striking feature of sharecropping is how little variation there is in contracts involving different bargainers, different soil qualities, different sizes of plots, and even different crops. For example, Bardhan (1984) analyzes extensive survey data on sharecropping practices in certain regions of India for paddy and wheat. He finds that, in a sample of over 300 villages, more than two-thirds had a single type of contract within the village, and over 95 percent had at most two forms of contract within the village. While some differences exist in the terms of individual contracts, he concludes (p. 115) that there is little evidence that the outcome is very sensitive to the relative bargaining power between tenants and landowners, their degrees of risk aversion or even the size of the farm.

A second significant feature of sharecropping contracts is that the terms vary substantially from one village to another, even between villages that are similar in other respects like soil type, rainfall, scale of land holdings, and so on. Thus in village *A* the rule may be one-third for the landlord and two-thirds for the laborer, while in village *B* some distance away it is half-half, and for village *C* it is two-thirds for the landlord and one-third for the tenant. Moreover, substantial differences between villages hold up when one looks at the real incomes of tenants and landlords instead of the shares themselves (Bardhan, 1984, ch. 4; Bardhan and Rudra, 1986). Of course, such disparities can persist only if labor is relatively immobile, but this appears to be the case: personal connections, local credit arrangements, and differences in dialect and culture substantially limit the mobility of laborers outside of a small area.

Perhaps the most striking feature of the data, however, is the prevalence of equal split between labor and land. In over two-thirds of the villages Bardhan surveyed, 50–50 was the dominant form of contract, and in most of these cases it was the only form of contract (Bardhan, 1984, Table 9.3). There is nothing in classical bargaining theory—or for that matter in classical economic theory—that would predict these results.[10] It is not even clear why 50–50 would be a prominent focal point when the parties are in such obviously asymmetric positions.

We can shed some light on these matters by adopting an evolutionary point of view. Imagine that the agents in each village are boundedly rational and myopic. They do the sensible thing most of the time—that is, they optimize up to a point—but their information is limited, and they have little ability to anticipate how future events are going to unfold. Assume now that when two agents come to the bargaining table their expectations are shaped by precedent, that is, by what they and others like them have received in recent bargains. If almost all landlords in the area get three-fifths, for example, then it is reasonable to expect that the landlord in this case will insist on three-fifths. Similarly, the laborer (knowing the landlord's expectations) will find it reasonable to insist on no more than two-fifths. Thus a positive feedback loop is created—precedents affect present expectations, which determine current actions, which in turn become future precedents. All of this is subject to random perturbations that arise from misperceptions, miscalculations, changing economic circumstances and other unpredictable forces that buffet the process about.

As we have already seen, this kind of dynamic process exhibits the following qualitative behavior. First, within a given village where interactions are frequent, we can expect that most people will be using the same rule most of the time (local conformity). Second, in villages that do not interact, the prevailing conventions may differ—depending on the historical path the process follows—even though they have similar resource endowments (global diversity). Third, every so often the local convention may suddenly shift in response to accumulated stochastic shocks (punctuated equilibrium). The data on India do not cover a long enough period to test this last prediction, but they certainly are consistent with the first two effects.

Moreover, the evolutionary model predicts something else. It says that when the perturbations are fairly small, one particular equilibrium (or one class of equilibria) will be observed more often than the others. These equilibria are *stochastically stable*, that is, they are robust under small, persistent random shocks (Foster and Young, 1990).[11] In discussing rules of the road, for instance, we saw that a globally uniform convention is more stable over the long run than two opposing conventions side by side, even though both are equilibria in a static sense. In the bargaining problem, it turns out that an even more striking result holds, namely, one division of the pie is more stable than the others. To understand why, let us focus on a particular village and suppose that in each period one landlord and one

[10] As Stiglitz (1989, p. 22) puts it: "[T]he range of contract forms [in sharecropping] seems far more restricted than theory would suggest. . . . Although there have been several attempts to explain this uniformity, none has gained general acceptance."

[11] This concept was introduced by Foster and Young as an alternative to the notion of evolutionary stable strategy (ESS) due to Maynard-Smith and Price (1973). The drawback of an ESS is that it is only stable against a small one-time shock, whereas most evolutionary processes are subjected to *persistent* shocks. For applications of this concept not mentioned in the text, see Young and Foster (1991), Nöldeke and Samuelson (1993), Samuelson (1994) and Kandori and Rob (1995).

tenant must negotiate a contract. We model the one-shot bargaining game as follows. The landlord demands x, and the tenant demands y. If their demands are mutually compatible ($x + y \leq 1$), each side gets what it asked for. If the demands are incompatible ($x + y > 1$), the negotiation breaks down, and they get nothing. This is known as the Nash demand game (Nash, 1953). Notice that for every value of x strictly between zero and one, the outcome in which the landlord demands x and the tenant demands $1 - x$ is a strict Nash equilibrium of this game.

To simplify the argument, assume that each side has only three possible demands—high (3/4), medium (1/2) and low (1/4). Assume further that the payoff to each bargainer is proportional to the share of the crop he gets, that is, the parties are risk neutral. Then the payoff matrix is as follows:

		Tenants		
		H	M	L
	H	0, 0	0, 0	75, 25
Landlords	M	0, 0	50, 50	50, 25
	L	25, 75	25, 50	25, 25

There are three possible conventions: 1/4 : 3/4, 3/4 : 1/4, and 1/2 : 1/2, where we list the landlord's share first and the tenant's second.

To determine the stochastically stable convention, we must compute the probability that enough stochastic shocks accumulate to push the community out of one convention and into another, and we must do this for every pair of conventions. As before, we assume that agents make their choices based on information acquired by word of mouth and personal experience. Occasionally they make idiosyncratic, unexplained choices that function like mutations. Assume for simplicity that each bargainer knows the demands that were made in s out of the previous m bargains (whether these demands were compatible or not). He uses this information to predict how much the other side is likely to demand and chooses a best reply with probability $1 - \varepsilon$. However, with probability ε he demands 3/4, 1/2, or 1/4 at random. (When an agent chooses randomly, we may assume that all actions are chosen with equal probability, though the analysis does not depend on this assumption.)

To illustrate, suppose that the sample size is $s = 12$, $\varepsilon = .01$, and the current convention is 3/4 : 1/4. No matter which of the precedents appear in an agent's sample, the best reply is to choose 3/4 if the agent is a landlord and 1/4 if he is a tenant. To destabilize this convention requires that some agents choose something unconventional. Suppose that four tenants in succession demand 1/2 instead of 1/4. If the next landlord samples the most recent data, he may be led to believe that tenants demand 1/2 one-third of the time (because his sample has four out of 12 instances of this), and that they demand 1/4 two-thirds of the time. A best reply to this situation is 1/2. The same decision could be made by the next landlord, and the next after that and so forth. After enough landlords have demanded 1/2, tenants who sample this

Figure 4
Transition Probabilities Between Conventions

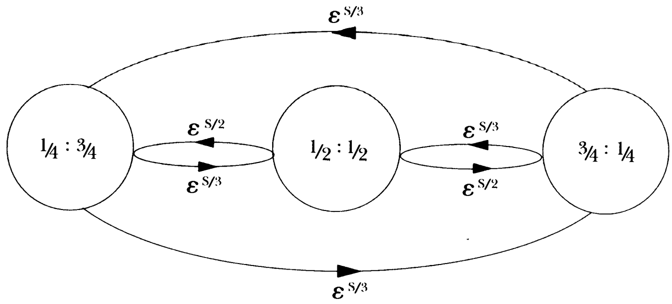

information will rationally demand 1/2. Thus the original four mistakes (whose probability is ε^4) can tip the process from the 3/4 : 1/4 regime to the 1/2 : 1/2 regime.

It can be checked, however, that it takes at least six mistakes to go the other way, that is, to tip out of the 1/2 : 1/2 convention into either of the others. Figure 4 shows the transition probabilities as a function of ε and s between every two conventions (when $s/2$ or $s/3$ are fractional, they are rounded to the next higher integer). When ϵ is small the probability of getting out of 1/2 : 1/2 is substantially lower than the probability of getting in. It follows that, when ε is sufficiently small (and s is sufficiently large), the process is much more likely to be in the 1/2 : 1/2 convention than in either of the others.[12]

This result generalizes as follows. Suppose that i) all agents are risk neutral or risk averse; ii) for every landlord with a given utility function there is a positive probability that the laborer against whom he is matched has the same utility function and vice versa; iii) all agents have the same amount of information s, which is sufficiently large and sufficiently incomplete ($s/m \leq .5$); and iv) the agents can demand any decimal fraction of the pie, where the demands are rounded to some fixed number of decimal places to keep the state space finite. Then a 1/2 : 1/2 split is the unique stochastically stable convention (Young, 1993b).

Concretely, this means the following. Suppose the evolutionary process is run in each of 100 villages starting from arbitrary initial conditions and with no communication between the villages. Then after a sufficiently long period of time, 50–50 will be the dominant mode of division in most (though not necessarily all) of the villages. This is quite consistent with the pattern we see in the data.

Of course other explanations for the prevalence of 50–50 are possible. The parties may simply consider 50–50 to be the "fairest" solution. Or they may select it because it is "focal" (Schelling, 1960). On reflection, however, it is not clear why anyone would

[12] The argument uses techniques from the theory of perturbed dynamical systems (Freidlin and Wentzell, 1984; Foster and Young, 1990; Young, 1993a; Kandori, Mailath and Rob, 1993).

consider equal division to be either fair or focal when the bargainers contribute such different inputs to the production process. Nor would one expect focal points to play such an important role when the stakes are so high. Yet it appears that they do.

We can make sense of these facts by viewing focal points as the outcome of an evolutionary process of expectations formation. Something comes to be focal; it is not necessarily focal in itself. Using a relatively simple model, we have shown that some norms are more likely to be observed than others—not because they are more prominent, or more natural, or more ethical *a priori,* but simply because they are more *stable* in the long run. We may even conjecture that long-run stable solutions, like equal division, acquire some of their ethical force *because* of their inherent stability.

These results do not hinge on unrealistically strong assumptions about rationality and common knowledge; nor do they presuppose that individuals can foresee all the results of their (and others') actions. At the same time, we do not dispense with rationality altogether. Instead, we think of society as groping its way through the uncoordinated actions of individuals doing the best they can with limited information. In spite of this rather untidy state of affairs (and in part because of it), the resulting evolutionary dynamic has both qualitative and quantitative features that seem to be broadly in keeping with empirical evidence.

■ *I am indebted to Alan Auerbach, George Mailath, Carl Shapiro and Timothy Taylor for many helpful comments and suggestions.*

References

An, Mark Y., and Nicholas M. Kiefer, "Evolution and Equilibrium Selection in Repeated Lattice Games," preprint, Cornell University, 1993.

Anderlini, Luca, and Antonella Ianni, "Local Learning on a Torus," preprint, Cambridge University, 1993.

Arthur, W. Brian, "Competing Technologies, Increasing Returns, and Lock-in by Historical Events," *Economic Journal,* March 1989, *99,* 116–31.

Arthur, W. Brian, Yuri Ermoliev, and Yuri Kaniovski, "Strong Laws for a Class of Path-Dependent Stochastic Processes." In Arkin, V., A. Shiryaev, and R. Wets, eds., *Proceedings of the International Conference on Stochastic Optimization.* Berlin: Springer Lecture Notes in Control and Information Sciences, 1986, *81,* 287–300.

Arthur, W. Brian, Yuri Ermoliev, and Yuri Kaniovski, "Path-Dependent Processes and the Emergence of Macrostructure," *European Journal of Operational Research,* 1987, *30,* 294–303.

Bardhan, Pranab, *Land, Labor, and Rural Poverty.* New York: Columbia University Press, 1984.

Bardhan, Pranab, and Ashok Rudra, "Labour Mobility and the Boundaries of the Village Moral Economy," *Journal of Peasant Studies,* 1986, *13,* 90–115.

Berninghaus, Siegfried K., and Ulrich Schwalbe, "Learning and Adaptation Processes in Games with Local Interaction Structures," preprint, University of Mannheim, 1992.

Besen, Stanley M., and Joseph Farrell, "Choosing How to Compete: Strategies and Tactics in Standardization," *Journal of Economic Perspectives,* Spring 1994, *8,* 116–31.

Binmore, Ken, "Do People Exploit their Bargaining Power: An Experimental Study," *Games and Economic Behavior,* 1991, *3,* 295–322.

Binmore, Ken, and Partha Dasgupta, eds., *The Economics of Bargaining.* Oxford: Blackwell, 1987.

Binmore, Ken, J. Swierzbinski, S. Hsu, and C. Proulx, "Focal Points and Bargaining," *International Journal of Game Theory,* 1993, *22,* 381–409.

Blume, Larry, "The Statistical Mechanics of Strategic Interaction," *Games and Economic Behavior,* 1993, *4,* 387–424.

David, Paul, "Clio and the Economics of QWERTY," *American Economic Review,* May 1985, *75,* 332–37.

Ellison, Glenn, "Learning, Local Interaction, and Coordination," *Econometrica,* September 1993, *61,* 1047–71.

Elster, Jon, "Social Norms and Economics," *Journal of Economic Perspectives,* Fall 1989, *3,* 99–117.

Foster, Dean, and H. Peyton Young, "Stochastic Evolutionary Game Dynamics," *Theoretical Population Biology,* 1990, *38,* 219–32.

Freidlin, Mark, and Alexander Wentzell, *Random Perturbations of Dynamical Systems.* Berlin: Springer Verlag, 1984.

Hamer, Mick, "Left is Right on the Road: The History of Road Traffic Regulations," *New Scientist,* December 25, 1986, *112.*

Hopper, R. H., "Left-Right: Why Driving Rules Differ," *Transportation Quarterly,* 1982, *36,* 541–48.

Goyal, Sanjeev, and Maarten Janssen, "Interaction Structure and the Stability of Conventions," preprint, Erasmus University, 1993.

Kandori, Michihiro, and Rafael Rob, "Evolution of Equilibria in the Long Run: A General Theory and Applications," *Journal of Economic Theory,* 1995, *65,* 383–415.

Kandori, Michihiro, George Mailath, and Rafael Rob, "Learning, Mutation, and Long-Run Equilibria in Games," *Econometrica,* January 1993, *61,* 29–56.

Kaniovski, Yuri, and H. Peyton Young, "Learning Dynamics in Games with Stochastic Perturbations," *Games and Economic Behavior,* 1995, *11*:2, 330–63.

Katz, Michael, and Carl Shapiro, "Network Externalities, Competition, and Compatibility," *American Economic Review,* June 1985, *75,* 424–40.

Katz, Michael, and Carl Shapiro, "Technology Adoption in the Presence of Network Externalities," *Journal of Political Economy,* August 1986, *94,* 822–41.

Katz, Michael, and Carl Shapiro, "Systems Competition and Network Effects," *Journal of Economic Perspectives,* Spring 1994, *8,* 93–115.

Kincaid, P., *The Rule of the Road.* New York: Greenwood, 1986.

Lay, Maxwell G., *Ways of the World.* New Brunswick N.J.: Rutgers University Press, 1992.

Lewis, David, *Convention: A Philosophical Study.* Cambridge, Mass.: Harvard University Press, 1969.

Liebowitz, S. J., and Stephen E. Margolis, "Network Externality: An Uncommon Tragedy," *Journal of Economic Perspectives,* Spring 1994, *8,* 133–50.

Luce, R. Duncan, and Howard Raiffa, *Games and Decisions.* New York: Wiley, 1957.

Maynard-Smith, John, and George R. Price, "The Logic of Animal Conflict," *Nature,* 1973, *246,* 15–18.

Nash, John, "The Bargaining Problem," *Econometrica,* April 1950, *18,* 155–62.

Nash, John, "Two-Person Cooperative Games," *Econometrica,* January 1953, *21,* 128–40.

Nöldeke, Georg, and Larry Samuelson, "An Evolutionary Analysis of Backward and Forward Induction," *Games and Economic Behavior,* 1993, *5,* 425–54.

Nydegger, R. V., and Guillermo Owen, "Two-Person Bargaining: An Experimental Test of the Nash Axioms," *International Journal of Game Theory,* 1974, *3,* 239–49.

Osborne, Martin, and Ariel Rubinstein, *Bargaining and Markets.* New York: Academic Press, 1990.

Roth, Alvin, ed., *Laboratory Experimentation in Economics.* Cambridge, U.K.: Cambridge University Press, 1987.

Rudra, Ashok, and Pranab Bardhan, *Agrarian Relations in West Bengal: Results of Two Surveys.* Bombay: Somaiya, 1986.

Samuelson, Larry, "Stochastic Stability in Games with Alternative Best Replies," *Journal of Economic Theory,* 1994, *64,* 35–65.

Schelling, Thomas C., *The Strategy of Conflict.* Cambridge, Mass.: Harvard University Press, 1960.

Schelling, Thomas C., "Dynamic Models of Segregation," *Journal of Mathematical Sociology,* 1971, *1,* 143–86.

Stiglitz, Joseph E., "Incentives and Risk Sharing in Sharecropping," *Review of Economic Studies,* April 1974, *41,* 219–55.

Stiglitz, Joseph E., "Rational Peasants, Efficient Institutions, and a Theory of Rural Organization: Methodological Remarks for Development Economics." In Bardhan, P., ed., *The Economic Theory of Agrarian Institutions.* Oxford: Clarendon Press, 1989, pp. 18–29.

Sugden, Robert, *The Evolution of Rights, Cooperation, and Welfare.* New York: Basil Blackwell, 1986.

Sugden, Robert, "Spontaneous Order," *Journal of Economic Perspectives,* Fall 1989, *3,* 85–97.

Ullman-Margalit, Edna, *The Emergence of Norms.* Oxford: Oxford University Press, 1977.

Yaari, Menachem, and Maya Bar-Hillel, "On Dividing Justly," *Social Choice and Welfare,* 1981, *1,* 1–24.

Young, H. Peyton, "The Evolution of Conventions," *Econometrica,* January 1993a, *61,* 57–84.

Young, H. Peyton, "An Evolutionary Model of Bargaining," *Journal of Economic Theory,* February 1993b, *59,* 145–68.

Young, H. Peyton, and Dean Foster, "Cooperation in the Short and in the Long Run," *Games and Economic Behavior,* 1991, *3,* 145–156.