

5 Evolution of Social Conventions

On Saturday mornings in Madison, many people buy produce at the farmers' market on the Capitol square. Given the placement of farm stands on the sidewalk around the square, it is difficult for pedestrian traffic to move in both directions. Thus, over the years, Madisonians have adopted the *social convention* of walking counterclockwise around the square. This "rule" seems somewhat arbitrary. Why not walk clockwise? However, it is obviously helpful to have some rule to coordinate pedestrians. Further, now that the convention has become established, it is largely self-enforcing (since walking clockwise would be difficult).

Once you begin looking for them, social conventions seem ubiquitous. To give another example, consider the meaning of words within a language. I'm currently sitting on an object called a "chair." Again, this "rule" seems arbitrary. Why not use the word "elephant" in place of "chair"? (We might then reserve the word "chair" for those large gray animals with floppy ears.) Nevertheless, coordination of human activity is facilitated by the adoption of some convention which ultimately becomes self-enforcing. We'll consider some further examples below.

Economists view social conventions as the outcome of a *coordination game*. However, as discussed in the next section, standard game-theoretic analysis merely predicts coordination on some outcome (e.g., clockwise or counterclockwise) without addressing the process by which the convention emerges. Recent work in *evolutionary game theory* attempts to fill this gap. Section 5.2 sketches an evolutionary approach to social conventions developed by Young (1998). This approach leads to the Markov chain model considered in section 5.3. This chain can be either absorbing or regular, depending whether individuals sometimes make "mistakes." In section 5.4, we'll then address a more complicated example using simulation methods.

5.1 Coordination games

Everyone drives on the right-hand side of the road in some countries (e.g., the US), while everyone drives on the left-hand side in other countries (e.g., the UK). Obviously, these social conventions became entrenched long ago (and were eventually reenforced by law).¹ But consider a simple two-person society that has not yet developed a convention. Adopting a game-theoretic perspective, each individual ("player") can choose L or R. If the players coordinate (both choosing R or both choosing L), then they obtain good outcomes (say 1 point each). But if the players fail to coordinate (making different choices), then they obtain bad outcomes (say 0

¹See Young (1998) for a brief history of driving conventions.

points each). This simple game – economists would call it a *coordination game* – is summarized by the payoff matrix below.²

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

Our choice of payoff levels was somewhat arbitrary. In particular, we could multiply all payoffs by a positive constant (or add a constant to all payoffs) without altering the game in any essential way. Nevertheless, this payoff structure does reflect both the incentives of the individuals to coordinate on *some* convention and their indifference over *which* convention is adopted.

To try to predict the outcome of this game, we might first consider the standard game-theoretic analysis. The first step is to determine each player’s *best response* to each action that could be taken by the opponent. To begin, consider the row player. If her opponent chooses L, then she can either obtain 1 point by choosing L or else obtain 0 points by choosing R. Thus, L is her best response to L. Similarly, if her opponent chooses R, then she can obtain 0 points by choosing L or else obtain 1 point by choosing R. Thus, R is her best response to R. The column player’s best responses could be determined in a similar fashion. The second step is to identify any *Nash equilibria* – any pair of actions such that (i) the row player’s action is a best response to the column player’s action and (ii) the column player’s action is a best response to the row player’s action. For our coordination game, we find that both (L,L) and (R,R) are Nash equilibria.³

On one hand, this game-theoretic analysis is appealing because it clearly captures the “self-enforcing” aspect of social conventions. Viewing a convention as a Nash equilibrium of a coordination game, individuals clearly have no incentive to deviate unilaterally from the convention (by definition of Nash equilibrium). On the other hand, standard game-theoretic analysis offers no specification of the *process* by which players reach a Nash equilibrium. Relatedly, given the existence of multiple Nash equilibria, it offers no clear prediction about *which* equilibrium is reached.

²Some further explanation of the payoff matrix may be helpful for readers unfamiliar with game theory. For each choice i that could be made by the row player (player 1) and each choice j that could be made by the column player (player 2), cell (i, j) of the payoff matrix reports the pair of payoffs (u_1, u_2) received by the players. Note that this indexing of the row and column players is itself a social convention within game theory. Of course, given the symmetry of the payoff matrix in the present example, the precise indexing of the players is irrelevant here.

³We will see in Chapter xx that there is also a third *mixed-strategy* Nash equilibrium. But this is irrelevant for our present purposes.

5.2 An evolutionary approach

Recent work in *evolutionary game theory* attempts to address these issues. We'll focus here on the approach developed by Young (1998).⁴ It begins by assuming a sequence of coordination games that continues indefinitely. In any period t , two players are drawn randomly from a large population of potential players. These players do not know the entire history of actions taken in the past, but merely know what happened in some recent periods. More precisely, given the “social memory” composed of the actions taken in the past m periods (i.e., periods $t-m$ through $t-1$), each player draws a random sample of s periods (hence $2s$ actions) from the social memory. Each player then determines her best response based upon her sample. Let p denote the proportion of Rs in the sample. Intuitively, R is a best response when $p \geq 1/2$, L is a best response when $p \leq 1/2$, and both actions are best responses when $p = 1/2$.⁵ Most of the time, each player simply chooses her best response. However, to capture the possibility that players sometimes make “mistakes,” we assume that each player randomizes over available actions with probability ϵ (where ϵ is small). After actions are chosen by the players, the social memory is updated (so that it now contains the actions taken in periods $t-m+1$ through t), two new players are drawn for period $t+1$, and the process continues.

From even this cursory description of the process, we can already foresee its long-run dynamics. If the social memory initially contains enough Ls and Rs, each player could potentially draw a sample containing more Ls than Rs (so that L is the best response) or more Rs than Ls (so that R is the best response) or perhaps equal numbers of Ls and Rs (so that both are best responses). But eventually, due simply to “random drift” induced by sampling variation, the social memory will become weighted more heavily toward one of the actions. If the number of Rs falls below s , then L is the best response for any sample that could possibly be drawn. In the absence of mistakes (i.e., $\epsilon = 0$), the number of Rs would continue to fall until the social memory holds no Rs. Conversely, if the number of Rs rises above $2m-s$, then R is the best response for any sample. In the absence of mistakes, the number of Rs would continue to rise until the social memory holds all Rs. Thus, if mistakes never occur, the society eventually adopts a social convention from which it never departs.

In contrast, if mistakes are possible (i.e., $\epsilon > 0$), social conventions are no longer permanent. Even if the social memory currently contains no Rs, one could be added if a player makes a mistake. Assuming s large and ϵ small, L remains the best response for any possible sample, and an isolated mistake will likely fade from the

⁴Young's approach might be labeled “stochastic” for reasons that will soon become apparent. In Chapter xx, we'll discuss an alternative “deterministic” approach to evolutionary game theory.

⁵More formally, best responses are determined by comparing the expected values of the actions. For the present game, the expected value of playing L is $EV(L) = 1 \times (1-p) + 0 \times p = 1-p$, while the expected value of playing R is $EV(R) = 1 \times p + 0 \times (1-p) = p$. R is a best response when $EV(R) \geq EV(L)$ which implies $p \geq 1/2$. Conversely, L is a best response when $EV(L) \geq EV(R)$ which implies $p \leq 1/2$.

social memory (after m periods elapse). However, given a sufficiently long *sequence* of mistakes, the number of Rs in the social memory could rise to s . Once this happens, R becomes a best response for some samples, and (due to random drift) the number of Rs in the social memory may continue to rise even if there are no further mistakes. Potentially, the number of Rs may rise beyond $m - 2s$ so that L is no longer a best response. Of course, given large s and small ϵ , it may take a long time before a sequence of mistakes pushes the system away from one convention toward another. But in finite stochastic processes (where sampling variation matters), anything that is possible will happen eventually.

Before turning to the formalization, let me emphasize three implications of this approach. First, the model predicts “local conformity” within societies. In most periods, the social memory is predominated by a single action which is the unique best response to all possible samples. Consequently, any mistakes are likely to fade quickly from the social memory. Second, the model predicts “global diversity” across societies. Even if two societies start from the same initial condition (i.e., the same initial configuration of the social memory), they may eventually adopt different conventions (as illustrated by the US and UK). Third, the model predicts “punctuated equilibrium.” Within a society, a shift from one convention to another occurs only rarely. But given a long enough time horizon, a series of mistakes will eventually push the society toward an alternative convention.

5.3 A Markov chain model

Young’s model can be formalized as a Markov chain process. However, the states of the process are more complicated than you might initially anticipate: each state corresponds to a configuration of the social memory. For each period within the social memory, the number of Rs is either 0 or 1 or 2, so there are 3 possible outcomes.⁶ Given m periods in the social memory, the process thus has 3^m states. Because the number of states rises rapidly in m , specification of a transition matrix is impractical unless m is very small. Setting $m = 2$, we obtain the 9 states listed below.⁷

00, 01, 02, 10, 11, 12, 20, 21, 22

States are indexed so that, if the process is in state ij in period t , the social memory contains i Rs for period $t - 2$ and j Rs for period $t - 1$. (Recall that we used a similar indexing scheme in Chapter 2 when we allowed the respondent’s social mobility to depend upon both his grandfather’s and father’s social class.) Consequently, for state ij , the proportion of Rs in the social memory is equal to $(i + j)/4$.

⁶For our present purposes, there is no need record which player (row or column) took which action (L or R). Otherwise, there would be 4 possible outcomes: (L,L), (L,R), (R,L) and (R,R).


⁷Obviously, our choice of m is greatly limited by tractability considerations. In the next section, we consider another example using a more flexible simulation method which permits choice of any m and any $s \leq m$.

Having enumerated the states of the process, we can now consider the transition probabilities. To simplify our task, we'll set $s = m = 2$ so that each player's "sample" is the entire social memory. Best response(s) are thus completely determined by the proportion of Rs in the social memory of the current state ij , so that $p = (i + j)/4$. Each player chooses her best response with probability $1 - \epsilon$, and randomizes over actions otherwise (choosing L with probability $\epsilon/2$ and R with probability $\epsilon/2$). Thus, for states **00** or **01** or **10** in which $p < 1/2$, each player chooses L with probability $1 - \epsilon/2$, and R with probability $\epsilon/2$. Conversely, for states **12** or **21** or **22** in which $p > 1/2$, each player chooses L with probability $\epsilon/2$, and R with probability $1 - \epsilon/2$. Finally, for states **02** or **11** or **20** in which $p = 1/2$, we assume that each player chooses L with probability $1/2$, and R with probability $1/2$. Further assuming that any randomization is independent across players, and recognizing that a chain in state ij must transition to some state jk , we thus obtain the transition matrix below.

	00	01	02	10	11	12	20	21	22
00	$(1 - \frac{\epsilon}{2})^2$	$2(\frac{\epsilon}{2})(1 - \frac{\epsilon}{2})$	$(\frac{\epsilon}{2})^2$	0	0	0	0	0	0
01	0	0	0	$(1 - \frac{\epsilon}{2})^2$	$2(\frac{\epsilon}{2})(1 - \frac{\epsilon}{2})$	$(\frac{\epsilon}{2})^2$	0	0	0
02	0	0	0	0	0	0	1/4	1/2	1/4
10	$(1 - \frac{\epsilon}{2})^2$	$2(\frac{\epsilon}{2})(1 - \frac{\epsilon}{2})$	$(\frac{\epsilon}{2})^2$	0	0	0	0	0	0
11	0	0	0	1/4	1/2	1/4	0	0	0
12	0	0	0	0	0	0	$(\frac{\epsilon}{2})^2$	$2(\frac{\epsilon}{2})(1 - \frac{\epsilon}{2})$	$(1 - \frac{\epsilon}{2})^2$
20	1/4	1/2	1/4	0	0	0	0	0	0
21	0	0	0	$(\frac{\epsilon}{2})^2$	$2(\frac{\epsilon}{2})(1 - \frac{\epsilon}{2})$	$(1 - \frac{\epsilon}{2})^2$	0	0	0
22	0	0	0	0	0	0	$(\frac{\epsilon}{2})^2$	$2(\frac{\epsilon}{2})(1 - \frac{\epsilon}{2})$	$(1 - \frac{\epsilon}{2})^2$

5.3.1 Case 1: absorbing chain

We can see from this matrix that, if mistakes never occur ($\epsilon = 0$), states **00** and **22** are absorbing. From the transition diagram below, it is also apparent that the process can transition (perhaps in multiple steps) from every non-absorbing state to one (or both) of the absorbing states. Thus, the chain itself is absorbing.



To facilitate numerical experiments, I've written a Matlab function m-file placed in Appendix 5.6.1. Given the input $\epsilon = 0$, this function generates a transition matrix which matches the diagram above. Raising this matrix to a very high power, we find the long-run probability distribution for each initial state.

```
>> help convention
```

```
function P = convention(e)
based on Peyton Smith's evolutionary analyses of conventions (JEP 1998)
assumes coordination game with actions (L,R) and payoffs [(1,1),(0,0); (0,0),(1,1)];
and size of social memory (m) = size of individual's sample (s) = 2
thus, state space is (00, 01, 02, 10, 11, 12, 20, 21, 22)
where state ij indicates that R was chosen i times in period t-2,
and R was chosen j times in period t-1
output P is the probability transition matrix
input e is the probability that each individual randomizes over actions
(instead of playing best response)
```

```
>> P = convention(0) % players never randomize (epsilon = 0)
```

```
P =
1.0000    0    0    0    0    0    0    0    0
0    0    0    1.0000    0    0    0    0    0
0    0    0    0    0    0    0.2500    0.5000    0.2500
1.0000    0    0    0    0    0    0    0    0
0    0    0    0.2500    0.5000    0.2500    0    0    0
0    0    0    0    0    0    0    0    1.0000
0.2500    0.5000    0.2500    0    0    0    0    0    0
0    0    0    0    0    0    1.0000    0    0
```

```

0      0      0      0      0      0      0      0      1.0000
>> P^10000 % long-run outcome for each initial state

ans =
1.0000      0      0      0      0      0      0      0      0
1.0000      0      0      0      0      0      0      0      0
0.2000      0      0      0      0      0      0      0      0.8000
1.0000      0      0      0      0      0      0      0      0
0.5000      0      0      0      0      0      0      0      0.5000
0          0      0      0      0      0      0      0      1.0000
0.8000      0      0      0      0      0      0      0      0.2000
0          0      0      0      0      0      0      0      1.0000
0          0      0      0      0      0      0      0      1.0000

```

Obviously, the process is ultimately absorbed in either state **00** or **22**. From some initial conditions, the eventual outcome is predetermined. For instance, if the process begins in state **01**, it is eventually absorbed in state **00** with probability 1. For other initial conditions, either social convention is possible. For instance, if the process begins in state **02**, there is a 20% chance it will eventually be absorbed in state **00** and an 80% chance it will be absorbed in state **22**. From the transition diagram, it is also clear that the expected time to absorption is low.⁸

5.3.2 Case 2: regular chain

Let's now consider the more interesting case where mistakes are possible (i.e., $\epsilon > 0$). In this case, the transition matrix P is primitive, and thus every row of P^t will converge to the limiting distribution as t becomes large. To illustrate, we'll suppose that players randomize 10% of the time (so that $\epsilon = 0.1$).

```

>> P = convention(.1) % players sometimes randomize (epsilon = 0.1)

P =
0.9025    0.0950    0.0025         0         0         0         0         0         0
0          0          0        0.9025    0.0950    0.0025         0         0         0
0          0          0         0         0         0        0.2500    0.5000    0.2500
0.9025    0.0950    0.0025         0         0         0         0         0         0
0          0          0        0.2500    0.5000    0.2500         0         0         0
0          0          0         0         0         0        0.0025    0.0950    0.9025
0.2500    0.5000    0.2500         0         0         0         0         0         0
0          0          0        0.0025    0.0950    0.9025         0         0         0
0          0          0         0         0         0        0.0025    0.0950    0.9025

```

```

>> P^1000 % long-run outcome

```

⁸Because we have not specified the transition matrix in canonical form, the Q matrix is given by rows and columns 2 through 8 of the P matrix. Once you've obtained the Q matrix, it is straightforward to compute the fundamental matrix and expected time to absorption. I'll leave this as an exercise for the reader.

```
ans =
  0.4039    0.0432    0.0015    0.0432    0.0164    0.0432    0.0015    0.0432    0.4039
  0.4039    0.0432    0.0015    0.0432    0.0164    0.0432    0.0015    0.0432    0.4039
  0.4039    0.0432    0.0015    0.0432    0.0164    0.0432    0.0015    0.0432    0.4039
  0.4039    0.0432    0.0015    0.0432    0.0164    0.0432    0.0015    0.0432    0.4039
  0.4039    0.0432    0.0015    0.0432    0.0164    0.0432    0.0015    0.0432    0.4039
  0.4039    0.0432    0.0015    0.0432    0.0164    0.0432    0.0015    0.0432    0.4039
  0.4039    0.0432    0.0015    0.0432    0.0164    0.0432    0.0015    0.0432    0.4039
  0.4039    0.0432    0.0015    0.0432    0.0164    0.0432    0.0015    0.0432    0.4039
  0.4039    0.0432    0.0015    0.0432    0.0164    0.0432    0.0015    0.0432    0.4039
```

In the long run, the chain thus spends 40.39% of its time in state **00**, another 40.39% of its time in state **22**, and the remaining time distributed across the other states.


The existence of a limiting distribution might initially seem difficult to reconcile with our earlier discussion of “punctuated equilibrium.” But upon reflection, there is no contradiction. Recall from Chapter 1 that we can adopt both “micro” and “macro” perspectives on Markov chain processes. In the context of social mobility, the micro perspective focused on the history of a single family line (a single chain), while the macro perspective focused on a large population of family lines (many independent chains). From the macro perspective, the limiting distribution can be interpreted as the proportion of family lines occupying each social class in the long run. But while there is stability at the population level, any particular family line will continue perpetually to move between classes in the manner dictated by the transition matrix.

Returning now to the social convention model, it is crucial to recognize that we have implicitly adopted a “micro” perspective by focusing on the history of single society (a single chain). Thus, “punctuated equilibrium” is a description of the pattern of transitions that would be observed within a single chain. At least in principle, we could also adopt a “macro” perspective by assuming a large number of societies (many independent chains). Given our computations above, we might then interpret the limiting distribution as the proportion of societies in each state, with 40.39% of societies in state **00**, another 40.39% of societies in state **22**, and so on. This is the formal basis for our claim that the model generates “global diversity” across societies. On the other hand, because the actual number of societies is small (at least relative to the number of individuals within societies), it seems specious to assume a “large population” of societies. Thus, the “macro” perspective seems less relevant for the present application.

To illustrate the (micro-level) dynamics of the model, we can use the `markovchain` `m-file` function (— from Chapter 1 —). Given the transition matrix above (with $\epsilon = 0.1$), we can generate a chain of length 1000, reflecting the history of our hypothetical society for 1000 periods. We can then compute the number of Rs in the social memory in each period, and plot the result.

```
>> c = markovchain(P,1000,1);
```

```
>> y = zeros(1,1000); y(c==2|c==4) = 1; y(c==3|c==5|c==7) = 2; y(c==6|c==8) = 3; y(c==9) = 4;
>> plot(y) % proportion of Rs in social memory in each period
```



For this example, mistakes are not especially rare, occurring 5% of the time (since $\epsilon/2 = 0.05$). Nevertheless, we still observe “local conformity.” In almost every period, the social memory is predominated by a single action (with at least 3 out of 4 actions the same). At the same time, we also observe “punctuated equilibrium.” Occasionally, *pairs* of mistakes enter the social memory, and this sometimes (though not always) causes the society to “jump” to the alternative convention.

5.3.3 Stochastic stability

Let’s now return to our concerns with standard game-theoretic analysis stated at the end of section 9.1. The evolutionary approach clearly addresses one concern, providing a specification of the process by which societies reach (and sometimes depart from) social conventions. Perhaps we have also implicitly addressed the concern about multiple equilibria. In the evolutionary approach, “accidents of history” play an important role in determining which convention develops. Thus, when confronted with multiple Nash equilibria, we might simply say that no precise prediction is possible without some knowledge of past play. In a phrase: “history matters.”

Still, the evolutionary approach offers another way to address the problem of equilibrium selection. Consider how the limiting distribution would change as the probability of randomization becomes very small. Intuitively, because mistakes become extremely rare, the process would spend very close to 0% of its time in most

states. But any states which retain a positive probability in the limiting distribution might be regarded as especially “stable.” More formally, given the limiting distribution \mathbf{x} (which depends implicitly on ϵ), we say that state i is *stochastically stable* when $\mathbf{x}(i)$ remains positive (bounded away from 0) as $\epsilon \rightarrow 0$.

To illustrate, let’s set ϵ very small (say 0.0001). For this parameter value, the rows of the P^t matrix do not converge to the limiting distribution by period $t = 10^6$, but do converge by period $t = 10^{12}$.

```
>> P = convention(.0001) % have chosen epsilon very small

P =
    0.9999    0.0001    0.0000         0         0         0         0         0         0
         0         0         0    0.9999    0.0001    0.0000         0         0         0
         0         0         0         0         0         0    0.2500    0.5000    0.2500
    0.9999    0.0001    0.0000         0         0         0         0         0         0
         0         0         0    0.2500    0.5000    0.2500         0         0         0
         0         0         0         0         0         0    0.0000    0.0001    0.9999
    0.2500    0.5000    0.2500         0         0         0         0         0         0
         0         0         0    0.0000    0.0001    0.9999         0         0         0
         0         0         0         0         0         0    0.0000    0.0001    0.9999

>> P^1000000 % after a million periods

ans =
    0.9929    0.0001    0.0000    0.0001    0.0000    0.0000    0.0000    0.0000    0.0069
    0.9928    0.0001    0.0000    0.0001    0.0000    0.0000    0.0000    0.0000    0.0070
    0.2041    0.0000    0.0000    0.0000    0.0000    0.0001    0.0000    0.0001    0.7957
    0.9929    0.0001    0.0000    0.0001    0.0000    0.0000    0.0000    0.0000    0.0069
    0.4999    0.0000    0.0000    0.0000    0.0000    0.0000    0.0000    0.0000    0.4999
    0.0069    0.0000    0.0000    0.0000    0.0000    0.0001    0.0000    0.0001    0.9929
    0.7957    0.0001    0.0000    0.0001    0.0000    0.0000    0.0000    0.0000    0.2041
    0.0070    0.0000    0.0000    0.0000    0.0000    0.0001    0.0000    0.0001    0.9928
    0.0069    0.0000    0.0000    0.0000    0.0000    0.0001    0.0000    0.0001    0.9929

>> ans^1000000 % after a million million periods

ans =

    0.4999    0.0000    0.0000    0.0000    0.0000    0.0000    0.0000    0.0000    0.4999
    0.4999    0.0000    0.0000    0.0000    0.0000    0.0000    0.0000    0.0000    0.4999
    0.4999    0.0000    0.0000    0.0000    0.0000    0.0000    0.0000    0.0000    0.4999
    0.4999    0.0000    0.0000    0.0000    0.0000    0.0000    0.0000    0.0000    0.4999
    0.4999    0.0000    0.0000    0.0000    0.0000    0.0000    0.0000    0.0000    0.4999
    0.4999    0.0000    0.0000    0.0000    0.0000    0.0000    0.0000    0.0000    0.4999
    0.4999    0.0000    0.0000    0.0000    0.0000    0.0000    0.0000    0.0000    0.4999
    0.4999    0.0000    0.0000    0.0000    0.0000    0.0000    0.0000    0.0000    0.4999
    0.4999    0.0000    0.0000    0.0000    0.0000    0.0000    0.0000    0.0000    0.4999
```

For this example, we thus find that both states **00** and **22** are stochastically stable. However, in other coordination games where players would prefer to coordinate on

one outcome rather than another, only the most preferred (“Pareto dominant”) Nash equilibrium is stochastically stable. (—This claim will be explored further in Exercise 5.x.—)

5.4 Bargaining

Beyond the simple left-right coordination game, other sorts of social conventions might also be addressed by the evolutionary approach. In particular, consider a negotiation between a seller and a buyer. The seller has an incentive to “hold out” for a high price, while the buyer might similarly hold out for a low price. However, if neither side is willing to compromise, the negotiation may break down entirely, with both sides losing potential gains from trade. Perhaps for this reason, many societies develop social conventions about “fair” bargaining outcomes.

This situation is captured (in an admittedly stylized way) by the following two-player game. To interpret this game, suppose the players are attempting to divide a “pie” worth 100 points. Each player can request a “low” share of the pie (25 points) or a “medium” share (50 points) or a “high” share (75 points). If the requests are compatible – if the sum of the requests does not exceed the size of the pie – then each player receives the amount requested. Otherwise, both players receive nothing.⁹

		<i>buyer</i>		
		L	M	H
<i>seller</i>	L	25,25	25,50	25,75
	M	50,25	50,50	0,0
	H	75,25	0,0	0,0

Standard game-theoretic analysis yields three Nash equilibria – (H,L), (M,M), and (L,H) – and hence three potential social conventions. Sellers receive most of the gains from trade in the first convention, while buyers receive most of the gains from trade in the third. In the second convention, gains from trade are split equally between the seller and buyer.

⁹We could modify the payoff structure so that nothing is “left over” if the players reach agreement. But our present specification of this game (called the *Nash bargaining game*) is standard in economics.

To specify the process by which some convention is reached, and perhaps say more about *which* convention is reached, we again adopt the evolutionary approach. In contrast to our previous example, we now assume *two* populations of players. In each period, the current seller is drawn randomly from a large population of potential sellers. After she draws a sample from the social memory, this seller uses actions of past *buyers* to determine her best response. Similarly, the current buyer is drawn randomly from a large population of potential buyers, and uses actions of past *sellers* to determine her best response. This two-population assumption is necessary to give either of the asymmetric conventions (H,L) or (L,H) a chance to evolve. However, it further increases the number of states of the chain. Because there are now 9 possible outcomes within each period, there are 9^m states of the chain. In place of the analytic strategy taken in the preceding section (setting m small and constructing a transition matrix), we will thus proceed via simulation analysis. Although this has some drawbacks (e.g., we can no longer compute the precise limiting distribution), it also has some obvious benefits (e.g., we are now free to set any m and any $s \leq m$ without worrying about tractability).

Simulation analysis of the model can undertaken using the following Matlab function m-file.

```
>> help bargaining

function [C,M] = bargaining(m,s,e,T,M)
bargaining between sellers and buyers (Nash bargaining game)
based on Young, JEP, 1998
input m = length of social memory
      s = sample size (where s <= m)
      e = probability that players randomize (instead of choosing best response)
      T = length of chain
      M = initial social memory
      if M = [] then initial social memory is generated randomly
output C is one realization of Markov chain (2xT matrix)
      C(1,t) gives action taken by player 1 (seller) in period t
      C(2,t) gives action taken by player 2 (buyer) in period t
      actions are indexed so that low (L) is 1, medium (M) is 2, high (H) is 3
      M is the final social memory (2xm matrix)
```

Readers interested in the details of the program are encouraged to study the code placed in the Appendix 5.6.2. But to proceed here, you need simply to understand the inputs and outputs of this function. In particular, note that the output C is a $2 \times T$ matrix reporting the entire history of play (including the initial social memory). For convenience, the actions of players are denoted by numerical index rather than alphabetical label (so that L is 1, M is 2, and H is 3). To illustrate, consider the two trials below.

```
>> C = bargaining(5,3,.1,15,[ ]) % m = 5, s = 3, epsilon = .1, T = 15, initial social memory random
C =
```

```

    1    2    2    2    1    2    2    2    2    2    2    2    2    2
    2    2    3    2    2    2    2    2    2    2    2    2    2    2
>> C = bargaining(5,3,.1,15,[ ]) % second trial (with same input parameters)
C =
    2    3    3    1    1    3    3    1    1    1    1    1    1    3
    2    3    1    3    3    3    3    3    3    3    3    3    3    3

```

Given the input parameters, the first 5 periods reflect the initial (randomly chosen) social memory, while the next 10 periods report the actions taken by the seller (top row) and buyer (bottom row) in each period.

5.4.1 Case 1: absorbing chain

Following our analytical strategy from the preceding section, we'll first explore the case where players never make mistakes (i.e., $\epsilon = 0$). Although we have not attempted to specify a transition matrix (nor drawn the associated transition diagram), it is apparent that the chain is absorbing. The three absorbing states correspond to the three Nash equilibria discussed above (although each absorbing state is an entire configuration of the social memory rather than a pair of actions). Thus, any chain will eventually be absorbed into one of those 3 states.

To see which social convention is more likely to emerge, we can run many trials (many independent chains) and then compute the proportion of chains reaching each absorbing state. The chains need to be long enough that every chain is absorbed by the final period. Given the large number of initial conditions (there are $9^5 = 59,049$ states), we will not choose a particular initial state, but instead randomly draw a new initial state for each trial.¹⁰ Note my use of the Matlab `unique` command to verify that every chain was absorbed, and to count the number of chains absorbed in each state.

```

>> [C,M] = bargaining(5,3,0,100,[ ]); % single trial with T = 100
>> M % final social memory for this trial
M =
    1    1    1    1    1
    3    3    3    3    3
>> reshape(M,1,10) % reshape memory as a 1 x 10 vector
ans =
    1    3    1    3    1    3    1    3    1    3
>> % now run 1000 trials, save final state of each chain as a row of the X matrix

```

¹⁰To undertake a more thorough analysis, we could conduct a grid search over initial conditions, obtaining a probability distribution over absorbing states for each initial condition.

```

>> X = []; for i = 1:1000; [C,M] = bargaining(5,3,0,100,[]); X = [X; reshape(M,1,10)]; end
>> [a,b,c] = unique(X,'rows'); a % unique rows of X matrix

a =
     1     3     1     3     1     3     1     3     1     3
     2     2     2     2     2     2     2     2     2     2
     3     1     3     1     3     1     3     1     3     1

>> % thus, all 1000 trials have reached one of the 3 absorbing states

>> sum([c==1 c==2 c==3])/1000 % probability distribution over absorbing states

ans =
    0.1920    0.6350    0.1730

```

We thus find that 63.50% of chains were absorbed in the state where the players choose (M,M). Of course, we might run even more trials to obtain a more precise estimate, and we have not considered the effect of the parameters m , s , and e on this distribution. Nevertheless, this result does begin to suggest that the social convention in which both parties choose M is more likely to emerge than the other two conventions.

5.4.2 Case 2: regular chain


We now consider the case where mistakes are possible (i.e., $\epsilon > 0$). Because the chain is now regular, we know there is a unique limiting distribution across states. However, given the number of states, it becomes impractical even to estimate this limiting distribution. Nevertheless, simulation analysis can provide some insight into this case. To start, we can again illustrate the concept of punctuated equilibrium by looking at the sequence of play within one trial (with $m = 5$, $s = 3$, and $\epsilon = 0.1$). To illustrate, I've plotted the seller's actions in the upper panel and the buyer's actions in the lower panel.

```


>> C = bargaining(5,3,.1,500,[]); % one trial with T = 500 periods

>> plot(C(1,:)) % seller's actions (L is 1, M is 2, H is 3)

```



```
>> plot(C(2,:)) % buyer's actions (L is 1, M is 2, H is 3)
```



For this particular trial (with the initial condition chosen randomly), we see that the actions initially “hover around” the Nash equilibrium (H,L) for most of the first 200 periods, but then hover around the Nash equilibrium (M,M) for most of the remaining periods.

Obviously, we should not place much weight on the outcome of any single trial. To get a better sense of the long-run dynamics of the system, we can again run many trials (many independent chains). Focusing on the last period of each trial, we can then estimate the probability distribution over the possible pairs of actions.¹¹ Some (unreported) preliminary analysis indicates that the chains need to be quite long in order to achieve independence from initial conditions.¹² For the trials below, each chain is thus $T = 1000$ periods long.

```
>> C = bargaining(5,3,.1,1000,[]); C(:,end)' % final pair of actions for one trial
```

¹¹Again, given the number of states, it is impractical to estimate the probability distribution over the (59,049) states of the chain, and we instead focus simply on the probability distribution over the (9) pairs of actions.

¹²Simulation analysts refer more colloquially to the “burn-in time” necessary for observed chains to reflect long-run dynamics.

```

ans =
     1     3

>> % now conduct 1000 trials, save last pair of actions as a row of the X matrix
>> X = []; for i = 1:1000; C = bargaining(5,3,.1,1000,[]); X = [X; C(:,end)']; end
>> [a,b,c] = unique(X,'rows'); a % unique rows of the X matrix

a =
     1     1
     1     2
     1     3
     2     1
     2     2
     2     3
     3     1
     3     2
     3     3

>> sum([c==1 c==2 c==3 c==4 c==5 c==6 c==7 c==8 c==9])/1000 % prob distn over pairs of actions

ans =
     0.0040     0.0410     0.0470     0.0340     0.7630     0.0340     0.0380     0.0290     0.0100

```

Thus, we find that 76.30% of trials ended with the pair of actions (M,M), while very few trials ended with players choosing the other Nash equilibria (H,L) or (L,H). This preliminary analysis is far from exhaustive. A more complete analysis might involve many more trials, consider the effect of the parameters m and s and ϵ , and also include a grid search over initial conditions. Nevertheless, our results again suggest that the “fair” bargaining convention is more likely to emerge than either of the asymmetric conventions.

5.4.3 Further analysis of best responses

Some further analysis helps explain our simulation results. Given the player’s sample of the social memory, let p denote the proportion of Hs, and let q denote the proportion of Ls. The proportion of Ms is thus $1 - p - q$. The expected value of each action is given below.

$$\begin{aligned}
 EV(L) &= p \times 25 + (1 - p - q) \times 25 + q \times 25 = 25 \\
 EV(M) &= p \times 50 + (1 - p - q) \times 50 + q \times 0 = 50(1 - q) \\
 EV(H) &= p \times 75 + (1 - p - q) \times 0 + q \times 0 = 75p
 \end{aligned}$$

Consequently, M is a better response than L when

$$EV(M) > EV(L) \quad \text{which implies} \quad p < 1/2,$$


H is a better response than L when

$$EV(H) > EV(L) \quad \text{which implies} \quad q < 1/3,$$

and H is a better response than M when

$$EV(H) > EV(M) \quad \text{which implies} \quad q > (2/3)(1 - p).$$

The best response for each sample can thus be summarized by the following diagram, called a *simplex*. On this diagram, any sample that might be drawn from the social memory is associated with some point $\{p, q\}$, and the best response for each region is indicated by the large letters.



It is apparent that the M region is larger than the other regions, indicating that M is a best response for a larger range of samples. In particular, M is the best response for “well-mixed” samples ($p = q = 1/3$) corresponding to the large dot in the center of the simplex. Thus, in the absence of mistakes, the (M,M) convention emerges for a larger range of initial conditions.

Graphically, each vertex of the simplex corresponds to one of the (mistake-free) conventions. For instance, if the society was in the (M,M) convention, all actions in the social memory would be Ms, and thus every sample would contain all Ms. Thus, the distance from one vertex (say the all-M vertex) to the other regions of the simplex (the L and H regions) reflects the number of mistakes that would need to enter the social memory in order for a new convention to emerge. From the simplex

diagram, we see that the distance from the all-M vertex to the other regions is greater than the distance from the other vertices to the relevant regions.¹³ Consequently, the (M,M) convention is more robust to mistakes.

5.5 Further reading

This chapter draws heavily on Young (J Econ Persp 1998). That paper provides another review of the evolutionary approach, along with further references to relevant economics literature. Gintis (*Game Theory Evolving*, Ch 10) provides a similar treatment of the Markov chain model in section 5.3.

If you're ever in Madison, you should visit the Dane County Farmers' Market. See www.dcfm.org for more information. But please respect our local convention.

5.6 Appendix

5.6.1 convention m-file

```
function P = convention(e)
% function P = convention(e)
% based on Peyton Smith's evolutionary analyses of conventions (JEP 1998)
% assumes coordination game with actions (L,R) and payoffs [(1,1),(0,0); (0,0),(1,1)];
% and size of social memory (m) = size of individual's sample (s) = 2
% thus, state space is (00, 01, 02, 10, 11, 12, 20, 21, 22)
% where state ij indicates that R was chosen i times in period t-2,
% and R was chosen j times in period t-1
% output P is the probability transition matrix
% input e is the probability that each individual randomizes over actions
% (instead of playing best response)

a = (1-e/2)^2;
b = 2*(e/2)*(1-e/2);
c = (e/2)^2;

P = [a b c 0 0 0 0 0 0;
     0 0 0 a b c 0 0 0;
     0 0 0 0 0 0 .25 .5 .25;
     a b c 0 0 0 0 0 0;
     0 0 0 .25 .5 .25 0 0 0;
     0 0 0 0 0 0 c b a;
     .25 .5 .25 0 0 0 0 0 0;
     0 0 0 c b a 0 0 0;
     0 0 0 0 0 0 c b a];
```

¹³In particular, the distance from the all-M vertex to the L region is 1/2, the distance from the all-L vertex to the H region is 1/3, and the distance from the all-H vertex to the M region is 1/3.

5.6.2 bargaining m-file

```
function [C,M] = bargaining(m,s,e,T,M)
% function [C,M] = bargaining(m,s,e,T,M)
% bargaining between sellers and buyers (Nash bargaining game)
% based on Young, JEP, 1998
% input m = length of social memory
%       s = sample size (where s <= m)
%       e = probability that players randomize (instead of choosing best response)
%       T = length of chain
%       M = initial social memory
%       if M = [] then initial social memory is generated randomly
% output C is one realization of Markov chain (2xT matrix)
%       C(1,t) gives action taken by player 1 (seller) in period t
%       C(2,t) gives action taken by player 2 (buyer) in period t
%       actions are indexed so that low (L) is 1, medium (M) is 2, high (H) is 3
%       M is the final social memory (2xm matrix)

V = [25 25 25; 50 50 0; 75 0 0]; % payoff matrix for Nash bargaining game

if isempty(M)
    M = ceil(rand(2,m)*3);
end
C = M; % chain starts with social memory

for t = 1:T-m

    % determine seller's action

    r = randperm(m); S = M(2, r(1:s)); % draw seller's sample
    p = [sum(S==1) sum(S==2) sum(S==3)]'/s; % probabilities of buyers' actions from sample
    br = (V*p == max(V*p)); br = br/sum(br); % seller's best response
    q = (1-e)*br + e*ones(3,1)/3; % seller's probabilities over actions
    a1 = sum(cumsum(q) < rand) + 1; % seller's action

    % determine buyer's action

    r = randperm(m); S = M(1, r(1:s)); % draw buyer's sample
    p = [sum(S==1) sum(S==2) sum(S==3)]'/s; % probabilities of sellers' actions from sample
    br = (V*p == max(V*p)); br = br/sum(br); % buyer's best response
    q = (1-e)*br + e*ones(3,1)/3; % buyer's probabilities over actions
    a2 = sum(cumsum(q) < rand) + 1; % buyer's action

    % update chain and social memory

    C = [C, [a1;a2]];
    M(:,1) = []; M = [M, [a1;a2]];

end
```