Forthcoming in the *New Palgrave Dictionary of Economics, Second Edition*, edited by Steven N. Durlauf and Lawrence E. Blume. London: Macmillan.

stochastic adaptive dynamics

Economic systems often involve large numbers of agents whose behaviour, and patterns of interaction, have stochastic components. The dynamic properties of such systems can be analysed using stochastic dynamical systems theory. A key feature of these processes is that their long-run behaviour often differs substantially from the behaviour of the deterministic process obtained by taking expectations of the random variables. Furthermore, unlike the deterministic dynamics, the theory yields sharp predictions about the probability of being in different equilibria independently of the initial conditions.

Stochastic adaptive dynamics require analytical methods and solution concepts that differ in important ways from those used to study deterministic processes. Consider, for example, the notion of asymptotic stability: in a deterministic dynamical system, a state is locally asymptotically stable if any sufficiently small deviation from the original state is self-correcting. We can think of this as a first step toward analysing the effect of stochastic shocks; that is, a state is locally asymptotically stable if, after the impact of a small, one-time shock, the process evolves back to its original state.

This idea is not entirely satisfactory, however, because it treats shocks as if they were isolated events. Economic systems are usually composed of large numbers of interacting agents whose behaviour is *constantly* being buffeted by perturbations from various sources. These *persistent* shocks have substantially different effects from *one-time* shocks; in particular, persistent shocks can accumulate and tip the process out of the basin of attraction of an asymptotically stable state. Thus, in a stochastic setting, conventional notions of dynamic stability – including *evolutionarily stable strategies* – are inadequate to characterize the long-run behaviour of the process. Here we shall outline an alternative approach that is based on the theory of large deviations in Markov processes (Freidlin and Wentzell, 1984; Foster and Young, 1990; Young, 1993a).

Types of stochastic perturbations

Before introducing formal definitions, let us consider the various kinds of stochastic shocks to which a system of interacting agents may be exposed. First, there is the interaction process itself whereby agents randomly encounter other agents in the population. Second, the agents' behaviour will be *intentionally* stochastic if they are employing mixed strategies. Third, their behaviour may be *unintentionally* stochastic if their payoffs are subject to unobserved utility shocks. Fourth, mutation processes may cause one type of agent to change spontaneously into another type. Fifth, in- and out-migration can introduce new behaviours into the population or extinguish existing ones. Sixth, the system may be hit by aggregate shocks that change the *distribution* of behaviours. This list is by no means exhaustive, but it does convey some sense of the range of stochastic influences that arise quite naturally in economic (and biological) contexts.

Stochastic stability

The early literature on evolutionary game dynamics tended to sidestep stochastic issues by appealing to the law of large numbers. The reasoning is that, when a population is large, random influences at the individual level will tend to average out, and the aggregate state variables will evolve according to the *expected* (hence deterministic) direction of motion. While this approximation may be reasonable in the short and medium run, however, it can be quite misleading when extrapolated over longer periods of time. The difficulty is that, even when the stochastic shocks have very small probability, their accumulation can have dramatic long-run effects that push the process far away from its deterministic trajectory.

The key to analysing such processes is to observe that, when the aggregate stochastic effects are 'small' and the resulting process is ergodic, the long-run distribution will often be concentrated on a very small subset of states – possibly, in fact, on a *single* state. This leads to the idea of *stochastic stability*, a solution concept first proposed for general stochastic dynamical systems by Foster and Young (1990, p. 221): 'the stochastically stable set (SSS) is the set of states S such that, in the long run, it is nearly certain that the system lies within every open set containing S as the noise tends slowly to zero.' The analytical technique for computing these states relies on the theory of large deviations first developed for continuous-time processes by Freidlin and Wentzell (1984), and subsequently extended to general finite-state Markov chains by Young (1993a). It is in the latter form that the theory is usually applied in economic contexts.

An illustrative example

The following simple model illustrates the basic ideas. Consider a population of n agents who are playing the 'Stag Hunt' game:

	A	B
A	10, 10	0, 7
B	7, 0	7, 7

The *state* of the process at time t is the current number of agents playing A, which we shall denote by $z_t \in Z = \{0, 1, 2, \ldots, n\}$. Time is discrete. At the start of period $t+1$, one agent is chosen at random. Strategy A is a best response if $z_t \leq .7n$ and B is a best response if $z_t \geq .7n$. (We assume that the player includes herself in assessing the current distribution, which simplifies the computations.) With high probability, say $1-\varepsilon$, the agent chooses a best response to the current distribution of strategies; while with probability ε she chooses A or B at random (each with probability $\varepsilon/2$).

We can interpret such a departure from best response behaviour in various ways: it might be a form of experimentation, it might be a behavioural 'mutation', or it might simply be a form of ignorance – the agent may not know the current state. Whatever the explanation, the result is a *perturbed best response process* in which individuals choose (myopic) best responses to the current state with high probability and depart from best response behaviour with low probability.

This process is particularly easy to visualize because it is one-dimensional: the states can be viewed as points on a line, and in each period the process moves to the left by one step, to the right by one step, or stays put. Figure 1 illustrates the situation when the population consists of ten players.

The transitions indicated by solid arrows have high probability and represent the direction of best response, that is, the main flow of the process. The dashed arrows go against the flow and have low probability, which is the

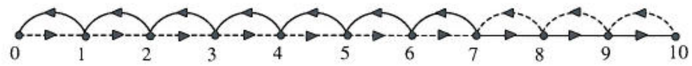

Figure 1 Transitions for the perturbed best response process in the Stag Hunt game and a population of ten agents. *Note*: Each vertex represents the number of agents playing action A at a given time. Solid arrows are transitions with high-probability, dotted arrows are transitions with low (order-ε) probability

same order of magnitude as ε. (The process can also loop by staying in a given state with positive probability; these loops are omitted from the figure to avoid clutter.)

In this example the transition probabilities are easy to compute. Consider any state z to the left of the critical value $z^* = 7$. The process moves right if and only if one more agent plays A. This occurs if and only if an agent currently playing B is drawn (an event with probability $1 - z/10$) and this agent mistakenly chooses A (an event with probability $\varepsilon/2$). In other words, if $z < 7$ the probability of moving right is $R_z = (1 - z/10)(\varepsilon/2)$. Similarly, the probability of moving left is $L_z = (z/10)(1 - \varepsilon/2)$. The key point is that the right transitions have much smaller probability than the left transitions when ε is small. Exactly the reverse is true for those states $z > 7$. In this case the probability of moving right is $R_z = (1 - z/10)(1 - \varepsilon/2)$, whereas the probability of moving left is $L_z = (z/10)(\varepsilon/2)$. (At $z = 7$ the process moves left with probability .15, moves right with probability .35, and stays put with probability .50.)

Computing the long-run distribution

Since this finite-state Markov chain is irreducible (each state is reachable from every other via a finite number of transitions), the process has a unique long-run distribution. That is, with probability 1, the relative frequency of being in any given state z equals some number μ_z *independently of the initial state*. Since the process is one-dimensional, the equations defining μ are particularly transparent, namely, it can be shown that for every $z < n$, $\mu_z R_z = \mu_{z+1} L_{z+1}$. This is known as the *detailed balance condition*. It has a simple interpretation: in the long run, the process transits from $z+1$ to z as often as it transits from z to $z+1$.

The solution in this case is very simple. Given any state z, consider the directed tree T_z consisting of all right transitions from states to the left of z and all left transitions from states to the right of z. This is called a *z-tree* (see Figure 2).

An elementary result in Markov chain theory says that, for one-dimensional chains, the long-run probability of being in state z is proportional to the product of the probabilities on the edges of T_z:

$$\mu_z \propto \prod_{y<z} \mathrm{R}_y \prod_{y>z} L_y. \tag{1}$$

This is a special case of the *Markov chain tree theorem*, which expresses the stationary distribution of any finite chain in terms of the probabilities of its z-trees. (Versions of this result go back at least to Kirchhoff's work in the 1840s; see Haken, 1978, s. 4.8. Freidlin and Wentzell, 1984, use it to study large deviations in continuous-time Wiener processes.)

Formula (1) allows us to compute the order-of-magnitude probability of each state without worrying about its exact magnitude. Figure 2 shows, for example, that μ_3, the long-run probability of state $z = 3$, must be proportional

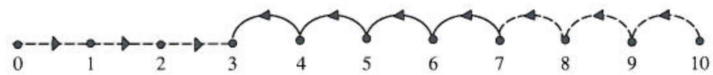

Figure 2 The unique 3-tree

to ε^6, because the 3-tree has six dotted arrows, each of which has probability of order ε. Using this method we can easily compute the relative probabilities of each state.

Stochastic stability and equilibrium selection

This example illustrates a general property of adaptive processes with small persistent shocks. That is, the persistent shocks act as a *selection mechanism*, and the *selection strength increases the less likely the shocks are*. The reason is that the long-run distribution depends on the probability of escaping from various states, and the critical escape probabilities are *exponential* in ε. Figure 1 shows, for example, that the probability of all-B (the left endpoint) is larger by a factor of $1/\varepsilon$ than the probability of any other state, and it is larger by a factor of $1/\varepsilon^4$ than the probability of all-A (the right endpoint). It follows that, as ε approaches zero, the long-run distribution of the process is concentrated entirely on the all-B state. It is the unique stochastically stable state.

While stochastic stability is defined in terms of the limit as the perturbation probabilities go to zero, sharp selection can in fact occur when the probabilities are quite large. To illustrate, suppose that we take $\varepsilon = .20$ in the above example. This defines a very noisy adjustment process, but in fact the long-run distribution is still strongly biased in favour of the all-B state. It can be shown, in fact, that the all-B state is nearly 50 times as probable as the all-A state. (See Young, 1998b, ch. 4, for a general analysis of stochastic selection bias in one-dimensional evolutionary models.)

A noteworthy feature of this example is that the stochastically stable state (all-B) does not correspond to the Pareto optimal equilibrium of the game, but rather to the risk dominant equilibrium (Harsanyi and Selten, 1988). The connection between stochastic stability and risk dominance was first pointed out by Kandori, Mailath and Rob (1993). Essentially their result says that, in any symmetric 2×2 game with a uniform mutation process, the risk dominant equilibrium is stochastically stable provided the population is sufficiently large. The logic of this connection can be seen in the above example. In the pure best response process ($\varepsilon = 0$) there are two absorbing states: all-B and all-A. The basin of attraction of all-B is the set of states to the left of the critical point, while the basin of attraction of the all-A is the set of states to the right of the critical point. The left basin is bigger than the right basin. To go from the left endpoint into the opposite basin therefore requires more 'uphill' motion than to go the other way around. In any symmetric 2×2 coordination game the risk dominant equilibrium is the one with the widest basin, hence it is stochastically stable under uniform stochastic shocks of the above type.

How general is this result? It depends in part on the nature of the shocks. On the one hand, if we change the probabilities of left and right transitions in an arbitrary way, then we can force any given state – including non-equilibrium states – to have the highest long-run probability; indeed this follows readily from formula (1). (See Bergin and Lipman, 1996.) On the other hand, there are many natural perturbations that do lead to the risk dominant equilibrium in 2×2 games. Consider the following class of perturbed best response dynamics. In state z, let $\Delta(z)$ be the expected payoff from playing A against the population minus the payoff from playing B against the population. Assume that in state z the probability of choosing A divided by the probability of choosing B is well-approximated by a function of form $e^{h(\Delta(z))/\beta}$ where $h(\Delta)$ is non-decreasing in Δ, strictly increasing at $\Delta = 0$, and skew-symmetric ($h(\Delta) = - h(-\Delta)$). The positive scalar β is a measure of the noise level. In this set-up, a state is *stochastically stable* if its long-run probability is bounded away from zero as $\beta \to 0$. Subject to some minor additional regularity assumptions, it can be shown that, in any symmetric 2×2 coordination game, if the population is large enough, the unique stochastically stable state is the one in which everyone plays the risk-dominant equilibrium (Blume, 2003).

Unfortunately, the connection between risk dominance and stochastic stability breaks down – even for uniform mutation rates – in games with more than two strategies per player (Young, 1993a). The difficulty stems from the fact that comparing 'basin sizes' works only in special situations. To determine the stochastically stable states in more general settings requires finding the path of least resistance – the path of greatest probability – from every absorbing set to every other absorbing set, and then constructing a rooted tree from these critical paths (Young, 1993a). (An absorbing set is a minimal set of states from which the *unperturbed* process cannot escape.) What makes the one-dimensional situation so special is that there are only two absorbing sets – the left endpoint and the right endpoint – and there is a unique directed path going from left to right and another unique path going from right to left. (For other situations in which the analysis can be simplified, see Ellison, 2000; Kandori and Rob, 1995.)

There are many games of economic importance in which this theory has powerful implications for equilibrium selection. In the non-cooperative Nash bargaining model, for example, the Nash bargaining solution is essentially the unique stochastically stable outcome (Young, 1993b). Different assumptions about the one-shot bargaining process lead instead to the selection of the Kalai–Smorodinsky solution (Young, 1998a; for further variations see Binmore, Samuelson and Young, 2003). In a standard oligopoly framework, marginal cost pricing turns out to be the stochastically stable solution (Vega-Redondo, 1997).

Speed of adjustment

One criticism that has been levelled at this approach is that it may take an exceedingly long time for the evolutionary process to reach the stochastically stable states when it starts from somewhere else. The difficulty is that, when the shocks have very small probability, it takes a long time (in expectation) before enough of them accumulate to tip the process into the stochastically stable state(s). While this is correct in principle, the waiting time can be very sensitive to various modelling details. First, it depends on the size and probability of the shocks themselves. As we have already noted, the shocks need not be small for sharp selection to occur, in which case the waiting time need not be long either. (In the above example we found that an error rate of 20 per cent still selects the all-B state with high probability.) Second, the expected waiting time depends crucially on the *topology* of interaction. In the above example we assumed that each agent reacts to the distribution of actions in the whole population. If instead we suppose that people respond only to actions of those in their immediate geographic (or social) neighbourhood, the time to reach the stochastically stable state is greatly reduced

(Ellison, 1993; Young, 1998b, ch. 6). Third, the waiting time is reduced if the stochastic perturbations are not independent, either because the agents act in a coordinated fashion, or because the utility shocks among agents are statistically correlated (Young, 1998b, ch. 9; Bowles, 2004).

Path dependence

The results discussed above rely on the assumption that the adaptive process is ergodic, that is, its long-run behaviour is almost surely independent of the initial state. Ergodicity holds if, for example, the number of states is finite, the transition probabilities are time-homogeneous, and there is a positive probability of transiting from any state to any other state within a finite number of periods. One way in which these conditions may fail is that the weight of history grows indefinitely. Consider, for example, a two-person game G together with a population of potential row players and another population of potential column players. Assume that an initial history of plays is given. In each period, one row player and one column player are drawn at random, and each of them chooses an ε-trembled best reply to the opposite population's previous actions (alternatively, to a random sample of fixed size drawn from the opponent's previous actions). This is a stochastic form of fictitious play (Fudenberg and Kreps, 1993; Kaniovski and Young, 1995). The *proportion* of agents playing each action evolves according to a stochastic difference equation in which the magnitude of the stochastic term decreases over time; in particular it decreases at the rate $1/t$.

This type of process is not ergodic. It can be shown, in fact, that the long-run proportions converge almost surely either to a neighbourhood of all-A or to a neighbourhood of all-B, where the relative probabilities of these two events depend on the initial state (Kaniovski and Young, 1995). Processes of this type require substantially different techniques of analysis from the ergodic processes discussed earlier; see in particular Arthur, Ermoliev and Kaniovski (1987), Benaim and Hirsch (1999) and Hofbauer and Sandholm (2002).

Summary

The introduction of persistent random shocks into models with large numbers of interacting agents can be handled using methods from stochastic dynamical systems theory; moreover, there is virtually no limit on the dimensionality of the systems that can be analysed using these techniques. Such processes can exhibit path dependence if the weight of history is allowed to grow indefinitely. If instead past actions fade away or are forgotten, the presence of persistent random shocks makes the process ergodic and its long-run behaviour is often easier to analyse. An important feature of such ergodic models is that some equilibrium states are much more likely to occur in the long run than others, and this holds independently of the initial state. The length of time that it takes to reach such states from out-of-equilibrium conditions depends on key structural properties of the model, including the size and frequency of the stochastic shocks, the extent to which they are correlated among agents, and the network topology governing agents' interactions with one another.

H. Peyton Young

See also

agent-based models;
evolutionary economics.

Bibliography

Arthur, W.B., Ermoliev, Y. and Kaniovski, Y. 1987. Adaptive growth processes modelled by urn schemes. *Kybernetika* 6, 49–57; English trans in *Cybernetics* 23, 779–89.

Benaim, M. and Hirsch, M.W. 1999. Mixed equilibria and dynamical systems arising from fictitious play in perturbed games. *Games and Economic Behavior* 29, 36–72.

Bergin, J. and Lipman, B.L. 1996. Evolution with state-dependent mutations. *Econometrica* 64, 943–56.

Binmore, K., Samuelson, L. and Young, H.P. 2003. Equilibrium selection in bargaining models. *Games and Economic Behavior* 45, 296–328.

Blume, L.E. 1995. The statistical mechanics of best-response strategy revision. *Games and Economic Behavior* 11, 111–45.

Blume, L.E. 2003. How noise matters. *Games and Economic Behavior* 44, 251–71.

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

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

Ellison, G. 2000. Basins of attraction, long run equilibria, and the speed of step-by-step evolution. *Review of Economic Studies* 67, 17–45.

Foster, D. and Young, H.P. 1990. Stochastic evolutionary game dynamics. *Theoretical Population Biology* 38, 219–32.

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

Fudenberg, D. and Kreps, D. 1993. Learning mixed equilibria. *Games and Economic Behavior* 5, 320–67.

Haken, H. 1978. *Synergetics*. Berlin: Springer-Verlag.

Harsanyi, J. and Selten, R. 1988. *A General Theory of Equilibrium Selection in Games*. Cambridge MA: MIT Press.

Hofbauer, J. and Sandholm, W. 2002. On the global convergence of stochastic fictitious play. *Econometrica* 70, 2265–94.

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

Kandori, M. and Rob, R. 1995. Evolution of equilibria in the long run: a general theory and applications. *Journal of Economic Theory* 65, 383–414.

Kaniovski, Y. and Young, H.P. 1995. Learning dynamics in games with stochastic perturbations. *Games and Economic Behavior* 11, 330–63.

Karlin, S. and Taylor, H.M. 1975. *A First Course in Stochastic Processes*. New York: Academic Press.

Vega-Redondo, F. 1997. The evolution of Walrasian behaviour. *Econometrica* 65, 375–84.

Young, H.P. 1993a. The evolution of conventions. *Econometrica* 61, 57–84.

Young, H.P. 1993b. An evolutionary model of bargaining. *Journal of Economic Theory* 59, 145–68.

Young, H.P. 1998a. Conventional contracts. *Review of Economic Studies* 65, 776–92.

Young, H.P. 1998b. *Individual Strategy and Social Structure*. Princeton. NJ: Princeton University Press.