Imitator-Experimenter Dynamics: A Complete Characterization*

Douglas Gale
Department of Economics
New York University
269 Mercer Street
New York, NY 10003

Robert W. Rosenthal
Department of Economics
Boston University
270 Bay State Road
Boston, MA 02215

January 2000

1 Introduction

Gale and Rosenthal (1999) [hereafter GR] analyzes a model in which a population of agents follows adaptive behavioral rules. These behaviors are referred to as experimentation and imitation. One of the agents is an *experimenter*; the rest are *imitators*. The experimenter randomly selects a new action from a neighborhood of his previous action and observes whether the payoff from the new action is higher than the payoff from his previous action, given the current actions of his opponents. If the payoff is higher, the experimenter adopts the new action; if not, he returns to his previous action. An imitator adjusts her action toward the average action adopted by her opponents.

In the basic model, agents play a one-shot game repeatedly. An agent's action in the one-shot game is a real number. Each agent has a common payoff function $\pi(x, y) = -(x - By)^2$, where x is the agent's action and y is the average of his opponent's actions. When $B > 0$, the players' actions are

*We thank Mike Harrison for technical advice and the National Science Foundation for financial assistance.

strategic complements; when $B < 0$, their actions are strategic substitutes. The game has a unique symmetric equilibrium, in which every agent chooses 0.

Note that the imitators' behavior does not depend on payoffs, so the specification of payoff functions for imitators does not play a role in the analysis. However, imitative behavior is implicitly motivated by the assumption that the imitators' payoff functions are similar to that of the experimenter. Also, the two notions of "equilibrium"–Nash equilibrium of a game and rest point of a mathematical dynamic–conveniently coincide in our model when the imitators' payoff functions are the same as the experimenter's and the dynamics are as described.

The game is played at a sequence of dates $t = 1, 2, ...$ The *state* of the game at date t is described by an ordered pair $(x_t, y_t) \in \mathbf{R} \times \mathbf{R}$, where x_t is the experimenter's action and y_t is the average of the imitators' actions. (The equilibrium corresponds to the origin of the state space.) The initial state is denoted by (x_0, y_0). With the specified dynamics, the (stochastic) evolution of play is described by a Markov chain $\Phi = \{(X_t, Y_t)\}$. GR analyzes the stability properties of the Markov chain Φ and proves the following results for the basic model:[1]

- When $B > 1$, the model is, in an intuitive sense, "explosive."
- When $B < 1$, the model is "stable in the large" in the sense that, for some compact set K containing the origin and any initial condition (x_0, y_0), the chain reaches K with probability one in finite time. Once the state K is reached, Φ remains there forever.
- When $0 < B < 1$, the model is "stable in the small" in the sense that, from every initial condition (x_0, y_0), the chain converges almost surely to the symmetric equilibrium $(0, 0)$.[2]
- When $B < 0$ and $|B|$ is sufficiently large, the model is "unstable in the small" in the sense that, for any sufficiently small neighborhood N of the origin, and any initial condition $(x_0, y_0) \in N$ with $x_0 \neq 0$, the chain Φ leaves the neighborhood N in finite time with probability

[1] A key assumption that distinguishes the results for the basic model of GR from those for some of its variants is that the size of the experimenter's search window is constant.

[2] When $B = 1$, the process neither explodes nor converges to the equilibrium (in any sense).

one; however, the model is "not too unstable" in the sense that, for any initial condition (x_0, y_0), the chain converges in probability to the origin.[3]

Except for the case where strategic complements are "too strong" ($B > 1$), the model is stable in the large. If strategic complements are not too strong ($0 < B < 1$), the model is stable in the small as well. The most interesting case is the one in which agents' actions are strategic substitutes. Here stability in the large co-exists with instability in the small, but the instability is of a strange kind because on average the chain appears to spend most of its time near the origin.

The characterization of the chain's stability properties in the presence of strategic substitutes is incomplete, however, because both results, "unstable in the small" and "not too unstable", are only proved for $B < 0$ but $|B|$ sufficiently large. The purpose of this note is to complete the characterization by filling that gap. Specifically, we show that there exists a constant $B^* < 0$ such that:

- If $B^* < B \leq 0$, the model is "stable in the small"
- If $B < B^*$, the model is "unstable in the small" but "not too unstable".

The second result actually follows directly from results in GR. Whenever $B < B^*$, a direct calculation generates the conclusion of Lemma 4 of GR, and the proofs of Theorems 2 (unstable) and 6 (not too unstable) apply without change. This note is therefore devoted to establishing almost-sure convergence of Φ when $B^* < B \leq 0$.

The method of proof used to establish almost-sure convergence in this paper is quite different from the one GR gives for the case of strategic complements. When actions are strategic complements, it can be shown that each sample path of the chain converges monotonically (in a sense) to the origin. When actions are strategic substitutes, the dynamics are much more complicated. In fact, the convergence proof that we give here has a stronger similarity to the proof of instability given in GR. Both depend on exploiting the "drift" of the chain. The argument presented here also exploits a version of the Strong Law of Large Numbers for dependent variables applied to the

[3]When $(x_0, y_0) = (0, 0)$, there is, of course, no movement. When $x_0 = 0$ but $y_0 \neq 0$, the chain leaves N with probability strictly between 0 and 1.

logarithm of the ratio of the absolute value of successive actions of the experimenter. GR, by contrast, uses the Strong Law for independent random variables and a stochastic dominance argument to prove instability in the small.

Section 2 presents the model, a formal statement of the result, and some intuition for the proof. Section 3 contains the proof. We refer the reader to GR for motivation and results on related models.

2 Model and Convergence Result

We retain the notation of GR. The Markov chain $\Phi = \{(X_t, Y_t)\}_{t=0}^{\infty}$ is defined as follows. Let

$$(X_0, Y_0) = (x_0, y_0)$$

where (x_0, y_0) is a given initial state. Let $\{\tilde{\omega}_t\}_{t=1}^{\infty}$ be an i.i.d. sequence of random variables, where each random variable $\tilde{\omega}_t$ is distributed uniformly on the interval $[-1, 1]$. The experimenter's action follows the (stochastic) law of motion:

$$X_t = \begin{cases} X_{t-1} + \tilde{\omega}_t & \text{if } X_{t-1} + \tilde{\omega}_t \in B(X_{t-1}, Y_t) \\ X_{t-1} & \text{otherwise} \end{cases},$$

where

$$B(x, y) \equiv \{x' | \ (x' - By)^2 \leq (x - By)^2\}$$

is the set of "better responses". The average imitator's action follows the "partial-adjustment" law of motion:

$$Y_t = \lambda X_{t-1} + (1 - \lambda) Y_{t-1},$$

for some constant $0 < \lambda < 1$. We can think of the random variables X_t and Y_t as defined on a probability space $(\Omega, \mathcal{F}, \mathbf{P})$, where Ω consists of realizations $\omega = (\omega_1, \omega_2, ...)$ of the random sequence $\{\tilde{\omega}_t\}$.

Let B^* be the unique solution to the equation

$$2B - 2 + (1 - 2B)\ln(1 - 2B) = 0$$

The number B^* is approximately -1.2955. The main result of this paper is the following theorem.

Theorem 1 *If $B^* < B \leq 0$, then for any initial condition (x_0, y_0) the Markov chain Φ converges to $(0,0)$ with probability one.*

Intuition: The chain Φ visits any neighborhood of the origin infinitely often; so it is sufficient to show that Φ converges to the origin with probability one when the initial state (x_0, y_0) is close to the origin. For states (x, y) very close to the origin, the experimenter's search window $[x-1, x+1]$ is much larger than $B(x, y)$, so the probability of finding a "better response" in any trial is very small. For this reason, there are likely to be long strings of failures between occasional successes. During a string of failures the experimenter stands still but the imitators continue to close the gap between x and y. As a result, the chain is very likely to be close to the diagonal $\{x = y\}$ when a success eventually occurs.

Suppose the experimenter has a success when the current state is $(X_t, Y_t) = (x, x)$ and x is near 0. From such a state, the distribution of $|X_{t+1}| / |X_t|$, conditional on success, is independent of x. Furthermore, direct calculation shows that in this case the conditional mean of the random variable $\ln(|X_{t+1}| / |X_t|)$ is negative if and only if $B > B^*$. In this sense, the stochastic process $\{|X_t|\}$ has negative "drift".

In GR we were able to show through an approximation argument that when the mean of $\ln(|X_{t+1}| / |X_t|)$ is positive (conditional on success and starting from the diagonal near the origin), the drift of the process $\{|X_t|\}$ is positive and the process $\{(X_t, Y_t)\}$ therefore "drifts" away from the origin.[4] This is the essence of the instability argument. Here we would like to show that the drift of the process $\{(X_t, Y_t)\}$ is toward the origin whenever $B^* < B \leq 0$ and use this fact to show that the process converges to the origin with probability one.

There are two difficulties to overcome in carrying out this plan. First, the approximation argument fails whenever $|y|/|x|$ is sufficiently large, no matter how close (x, y) is to the origin, and such states can arise with positive probability whenever $(x_0, y_0) \neq (0,0)$. Second, the process can escape from any sufficiently small neighborhood of the origin with positive probability, and this again invalidates the approximation argument. Our way around these difficulties is to analyze properties of an embedded subprocess $\{(X_{\tau_n}, Y_{\tau_n})\}$ of $\{(X_t, Y_t)\}$ observed after successes when $|X_{\tau_n}|$ is close to zero and when $|Y_{\tau_n}|/|X_{\tau_n}|$ is below some constant. We first use an approximation argument

[4]The approximation argument is valid as long as x is not exactly 0, but states having values $x = 0$ occur with probability zero as long as $x_0 \neq 0$.

to show that this embedded process converges almost surely to the origin and then use that result to show that the original process also converges to the origin.

3 Proof of the Theorem

It is convenient to define the state space of the chain $\Phi = \{(X_t, Y_t)\}$ to be $(\mathbf{R}\backslash\{0\}) \times \mathbf{R}$. This entails no essential loss of generality because, for any initial condition (x_0, y_0) such that $x_0 \neq 0$, the probability of $X_t = 0$ for any t is zero; and when $x_0 = 0$ and $X_t = 0$ for all t, the a.s. convergence of $\{(X_t, Y_t)\}$ is obvious.

Without loss of generality, we can assume that the initial condition belongs to some compact set K containing the origin (GR, Theorem 1). A set $S \subset K$ is *uniformly accessible* if there exists a constant $\delta > 0$ such that, for any initial condition $(x_0, y_0) \in K$, the probability that Φ reaches S in finite time is at least δ. If the set S is uniformly accessible, then Φ visits S infinitely often with probability one (Meyn and Tweedie, 1993, Theorem 9.3.1). This fact will be used to prove the following result.

For any $(\varepsilon, a) \gg (0, 0)$, let

$$N(\varepsilon) \equiv \{(x, y) \in K : 0 < |x| \leq \varepsilon\};$$

$$C(a) \equiv \{(x, y) : |y| \leq a|x|\};$$

and

$$\Delta(\varepsilon, a) = N(\varepsilon) \cap C(a).$$

Lemma 2 *For any $\varepsilon > 0$ and any $a > 1$, $(X_t, Y_t) \in \Delta(\varepsilon, a)$ infinitely often with probability one.*

Proof. First, define

$$N'(\varepsilon) \equiv \left\{(x, y) \in N(\varepsilon) : \frac{\varepsilon}{2} \leq |x| \leq \varepsilon \text{ whenever } xy > 0\right\}.$$

So $N(\varepsilon)\backslash N'(\varepsilon)$ consists of a strip in each of the first and third quadrants of $\mathbf{R}^2$.

Second, for some fixed k and any initial state in $N'(\varepsilon)$, a string of k successive failures takes the chain into $\Delta(\varepsilon, a)$, and this event has probability at least $1/2^k$. Third, it is easy to see that for some $\delta' > 0$ sufficiently small, from any initial state in $K \backslash N(\varepsilon)$ an event consisting of a string of successive successes can be constructed having probability at least δ' that takes the chain into $N'(\varepsilon)$ in finite time. Finally, from any initial state in $N(\varepsilon) \backslash N'(\varepsilon)$, since the probability of a success in finite time is one and the probability is at least one-half that the first success takes the chain across the vertical axis and hence into $N'(\varepsilon)$, we can set $\delta \leq \left(1/2^k\right) \min\{1/2, \delta'\}$ and uniform accessibility is established. ■

Suppose the current state is (x, x), so that the chain remains in the same state after a failed search, and suppose that $|x| > 0$ is small enough that the diameter of the better response set $B(x, x)$ is less than one. Then direct calculation shows that if X' denotes the realized outcome of the next successful search, the cumulative distribution function of $Q = \ln\left(|X'|/|x|\right)$ is

$$H(z) = \begin{cases} \frac{2e^z}{2-2B} & \text{if } z < 0 \\ \frac{1+e^z}{2-2B} & \text{if } 0 \leq z < \ln(1-2B) \\ 1 & \text{if } z \geq \ln(1-2B) \end{cases},$$

independently of x. Direct calculation also shows that the expectation of the random variable Q is strictly negative if $B^* < B \leq 0$, as we assume from now on.

Lemma 3 *Choose $\delta > 0$ small enough so that $E[Q] < -2\delta$ and choose $a > 0$ and $b \leq \ln(1-2B)$. Define $\tilde{Q}(\omega)$ by putting*

$$\widetilde{Q}(\omega) \equiv \begin{cases} Q(\omega) & \textit{whenever } -\ln(a) < Q(\omega) \leq b \\ R(\omega) & \textit{otherwise} \end{cases}$$

where $R(\omega) \in (-\ln(a), b]$ almost surely. Then there exists a number $\bar{a} > 0$ such that

$$E[\widetilde{Q}] < -2\delta,$$

for all $a \in (\bar{a}, \infty)$.

Proof. Let $A = \{\omega : Q \in (-\ln(a), b]\}$ and $A^c = \{\omega : Q \notin (-\ln(a), b]\}$. From the definition of $\tilde{Q}$,

$$\begin{aligned} E[\tilde{Q}] &= E\left[Q\,|A\right]\Pr\left[A\right] + E\left[R\,|A^c\right]\Pr\left[A^c\right] \\ &\leq E\left[Q\,|A\right]\Pr\left[A\right] + b\Pr\left[A^c\right] \\ &\leq E\left[Q\,|-\ln(a) < Q\right]\Pr\left[-\ln(a) < Q\right] + b\Pr\left[Q < -\ln(a)\right]. \end{aligned}$$

For any a sufficiently large, $b\Pr[Q < -\ln(a)]$ is arbitrarily close to 0 and

$$E[Q \,|{-\ln(a)} < Q]\Pr[-\ln(a) < Q] \approx E[Q] < -2\delta.$$

■

Lemma 4 *For any* $\widetilde{Q}$ *satisfying the hypotheses of Lemma 3,* $var\left[\widetilde{Q}\right] \leq \ln^2(a^{-1})$.

Proof. Direct calculation. ■

Let δ and $\bar{a}$ satisfy the conditions of Lemma 3 and fix $a \in (\max\{1, \bar{a}\}, \infty)$. Let $\bar{x} > 0$ be a fixed but arbitrary number and define Markov stopping times τ' and τ by putting $\tau'(\omega)$ equal to the time of the first successful experiment and

$$\tau(\omega) = \inf\{t \geq \tau'(\omega) : (X_t, Y_t)(\omega) \in \Delta(\overline{x}, a)\}.$$

In words, we want to "observe" the Markov chain the first time τ it is in $\Delta(\bar{x}, a)$ after a success.

Lemma 5 *There exists a constant* $\overline{x} > 0$ *such that, for any initial condition* $(x_0, y_0) \in \Delta(\overline{x}, a)$,

$$\Pr[\tau < \infty] = 1$$

and

$$E\left[\ln\left(\max\left\{\frac{1}{a}, \frac{|X_\tau|}{|x_0|}\right\}\right)\right] < -\delta.$$

Proof. With probability one, a success occurs in finite time. For any $\bar{x} > 0$, it follows from Lemma 2 that $\Delta(\overline{x}, a)$ is reached in finite time with probability one. This shows that $\tau < \infty$ with probability one.

Then if $\overline{x}$ is small enough, the mean of the distribution of $\ln(|X_{\tau'}|/|x_0|)$ is approximated by that of Q, and so

$$E\left[\ln\left(\max\left\{\frac{1}{a}, \frac{|X_{\tau'}|}{|x_0|}\right\}\right)\right] < -\delta.$$

Now $\tau \geq \tau'$, and if there are no additional successes between these two times, then $X_\tau = X_{\tau'}$ and there is nothing more to show. On the other hand, if there is an additional success between the two times, either $|X_{\tau'}| > \overline{x}$ or $|Y_{\tau'}| > a|X_{\tau'}|$. The probability of the latter event is small enough that even if all its probability were reassigned to the event $|X_\tau| = \overline{x}$, the sign of the conditional expectation in question would not be changed, as in Lemma 3. To handle the former event, set $b = \ln(\overline{x}/|x_0|)$ and apply Lemma 3 again. ■

In the light of Lemma 5, we can define an embedded chain $\overline{\Phi}$ by observing the Markov chain Φ the first time it is in $\Delta(\bar{x}, a)$ after a success. Define the stopping time τ_1 by putting

$$\tau_1(\omega) = \inf\{t : (X_t, Y_t)(\omega) \in \Delta(\bar{x}, a)\}.$$

Then define the stopping times τ'_{n+1} and τ_{n+1} recursively as follows:

$$\tau'_{n+1}(\omega) = \inf\{t > \tau_n(\omega) : X_t(\omega) \neq X_{\tau_n(\omega)}(\omega)\},$$

$$\tau_{n+1}(\omega) = \inf\left\{t \geq \tau'_{n+1}(\omega) : (X_t, Y_t)(\omega) \in \Delta(\bar{x}, a)\right\}.$$

Then with a slight abuse of notation we can define the chain $\overline{\Phi} = \{\overline{\Phi}_n\}$ by putting

$$\overline{\Phi}_n(\omega) = (X_n, Y_n)(\omega) = (X_{\tau_n(\omega)}, Y_{\tau_n(\omega)})(\omega), n = 1, 2, ...$$

Using Lemma 2 and the observation that an infinite number of successes occurs with probability one, the embedded chain is well defined. For this chain, define a sequence of random variables $\{Z_n\}$ by putting

$$Z_{n+1} \equiv \ln\left(\max\left\{\frac{1}{a}, \frac{|X_{n+1}|}{|X_n|}\right\}\right).$$

Then

$$E(Z_k|Z_1, ..., Z_{k-1}) = E\left[\ln\left(\max\left\{\frac{1}{a}, \frac{|X_k|}{|X_{k-1}|}\right\}\right)\middle| |X_{k-1}|\right] < -\delta$$

from Lemma 5.

So far, we have shown that the log-ratio of the embedded chain has negative "drift." To make use of this result, we use a version of the Law of Large Numbers for dependent random variables.

Strong Law of Large Numbers for Dependent Random Variables: [Loeve (1978), p. 125, Theorem 32.E] Suppose that $\{Z_n\}$ is a sequence of random variables satisfying

$$\sum_n \frac{\text{var}[Z_n]}{n^2} < \infty.$$

Then

$$\frac{1}{n}\sum_{k=1}^{n}(Z_k - E(Z_k|Z_1, ..., Z_{k-1})) \to_{a.s.} 0.$$

Lemma 6 *The sequence* $\{var[Z_n]\}$ *is uniformly bounded.*

Proof. Immediate from Lemma 4. ■

Lemma 7 $\{X_n\} \to_{a.s.} 0$.

Proof. $>$From Lemmas 5 and 6, and the Strong Law of Large Numbers for dependent random variables, there exists a number $\gamma > 0$ such that

$$\frac{1}{n}\sum_{k=1}^{n}\ln\left(\max\left\{\frac{1}{a}, \frac{|X_k|}{|X_{k-1}|}\right\}\right) \to_{a.s.} -\gamma \leq -\delta < 0.$$

Hence

$$\begin{aligned}\ln(|X_n|) &= \sum_{k=1}^{n}\ln\left(\frac{|X_k|}{|X_{k-1}|}\right) + \ln(|X_1|)\\ &\leq \sum_{k=1}^{n}\ln\left(\max\left\{\frac{1}{a}, \frac{|X_k|}{|X_{k-1}|}\right\}\right) + \ln(|X_1|) \to_{a.s.} -\infty.\end{aligned}$$

Hence $\{X_n\} \to_{a.s.} 0$. ■

To complete the proof of Theorem 1, we need to show that $(X_t, Y_t) \to 0$ almost surely. By Lemma 7, $\{X_n\} \to_{a.s.} 0$. The only way that $\{X_t\} \to_{a.s.} 0$ can fail is if there is positive probability of an infinite sequence of successes at nonembedded times t and the corresponding values of X_t do not converge to 0. To rule this out, take any converging sample path $\{X_n(\omega)\}$ of the chain $\overline{\Phi}$; so for any $\varepsilon > 0$ and some m sufficiently large, $|X_n(\omega)| < \varepsilon$ for all $n > m$. Since $(X_n, Y_n)(\omega) \in \Delta(\bar{x}, a)$, we have $|Y_n(\omega)| < a\varepsilon$ and $|BY_n(\omega)| < a\varepsilon|B|$. Thus, if t is the time of the next success after $(X_n, Y_n)(\omega)$ it must be that

$$|X_t(\omega)| < 2(|X_n(\omega)| + |BY_n(\omega)|) < 2(\varepsilon + a\varepsilon|B|).$$

For $\varepsilon > 0$ sufficiently small, $|X_t(\omega)| < \bar{x}$. But since t is not an embedded time, it must be the case that $(X_t, Y_t)(\omega) \notin C(a)$. Furthermore, since $a > 1$,

$|X_s(\omega)| \leq |X_n(\omega)|$ for as long as the chain remains in $N(\bar{x})\backslash C(a)$. Finally, when the chain leaves $N(\bar{x})\backslash C(a)$ it can only do so by re-entering $\Delta(\bar{x}, a)$, producing the embedded time τ_{n+1}. Thus,

$$|X_t(\omega)| \leq |X_n(\omega)| \; \forall t \in (\tau_n(\omega), \tau_{n+1}(\omega)), \; \forall n$$

with probability one. Thus, $\{X_t\} \to 0$ almost surely. The almost-sure convergence of $\{Y_t\}$ follows immediately. ■

References

Gale, D., and R. W. Rosenthal, "Experimentation, Imitation, and Stochastic Stability," *Journal of Economic Theory* **84** (1999), 1-40.

Loeve, M., *Probability Theory II*, 4th Edition, Springer-Verlag, New York, 1978.

Meyn, S., and R. Tweedie, *Markov Chains and Stochastic Stability*, Springer-Verlag, London, 1993.