Online Appendix to "Stochastic Imitative Game Dynamics with Committed Agents"

William H. Sandholm*

January 16, 2012

O.1. Imitative protocols, mean dynamics, and equilibrium selection

In this section, we consider stochastic evolution under two well-known imitative protocols: imitation of success and imitation driven by dissatisfaction.[1]

Example O.1.1. Imitation of success. Suppose that when an agent receives a revision opportunity, he randomly samples an opponent, switching to the opponent's strategy with probability proportional to the difference between the opponent's payoff and some fixed baseline $m \in \mathbf{R}$. This protocol is described by the conditional imitation probabilities

$$r_{ij}(\pi) = \lambda(\pi_j - m) \quad \text{for } j \neq i, \tag{O.1}$$

where $\lambda > 0$ and m are chosen to ensure that conditions (B) and (1) both hold. If there are no mutations, then substituting into equation (M_0) and simplifying shows that finite-horizon behavior under this protocol is described by the mean dynamic

$$\dot{\chi} = \lambda \chi(1 - \chi)(F_1(\chi) - F_0(\chi)). \tag{O.2}$$

This is the *replicator dynamic* of Taylor and Jonker [4], with the constant λ representing a uniform change in speed. ◆

*Department of Economics, University of Wisconsin, 1180 Observatory Drive, Madison, WI 53706, USA. e-mail: whs@ssc.wisc.edu; website: http://www.ssc.wisc.edu/~whs.

[1] Early analyses of these protocols include Björnerstedt and Weibull [1], Weibull [5], and Hofbauer [2].

Example O.1.2. Imitation driven by dissatisfaction. Suppose that a revising agent compares his own payoff to some aspiration level $M \in \mathbf{R}$. He then opts to switch strategies with probability proportional to the amount by which his payoff is deficient, switching to the strategy of a randomly chosen opponent in this event. This rule is captured by the conditional imitation probabilities

$$r_{ij}(\pi) = \lambda(M - \pi_i) \text{ for } j \neq i, \tag{O.3}$$

where $\lambda > 0$ and M are such that (B) and (1) both hold. If there are no mutations, substituting into equation (M_0) and simplifying shows that this protocol too generates the replicator dynamic (O.2) as its mean dynamic. ◆

Since protocols (O.1) and (O.3) generate the same mean dynamic, the deterministic approximation theorem described in Section 3 implies that in large populations, the finite-horizon aggregate behavior trajectories generated by the two protocols are essentially indistinguishable. Despite this strong agreement, we now use our results from Section 4 to show that infinite-horizon behavior under protocols (O.1) and (O.3) can be quite different: the two protocols can generate different stochastically stable states, even in very simple games.

Example O.1.3. Consider a population of agents matched to play the symmetric normal form game

$$A = \begin{pmatrix} h & h \\ b & B \end{pmatrix}$$

with $b < h < B$, which we call *Boar Hunt*. The strategies in this coordination game are hunting for Hare (strategy 0) and hunting for Boar (strategy 1). Hare yields a certain payoff of h, while coordinating on Boar yields the highest possible payoff of B. If a population of agents are matched to play Boar Hunt, the resulting population game has payoffs $F_0(\chi) = h$ and $F_1(\chi) = b(1 - \chi) + B\chi$. In addition to the all-Hare ($\chi = 0$) and all-Boar ($\chi = 1$) equilibria, there is a mixed equilibrium in which fraction $\chi^* = \frac{h-b}{B-b}$ of the population chooses Boar. Note that either strategy in Boar Hunt can be risk dominant.

If committed agents are present, Theorems 4.2 and 4.3 tell us that the state that is stochastically stable in the large population limit is the one that maximizes the ordinal potential function J. It is easy to check that since F is a coordination game, the ordinal potentials induced by protocols (O.1) and (O.3) are strictly convex. Thus, if $J(1) < J(0) \equiv 0$, $\chi = 0$ is stochastically stable, while if $J(1) > 0$, state $\chi = 1$ is stochastically stable.

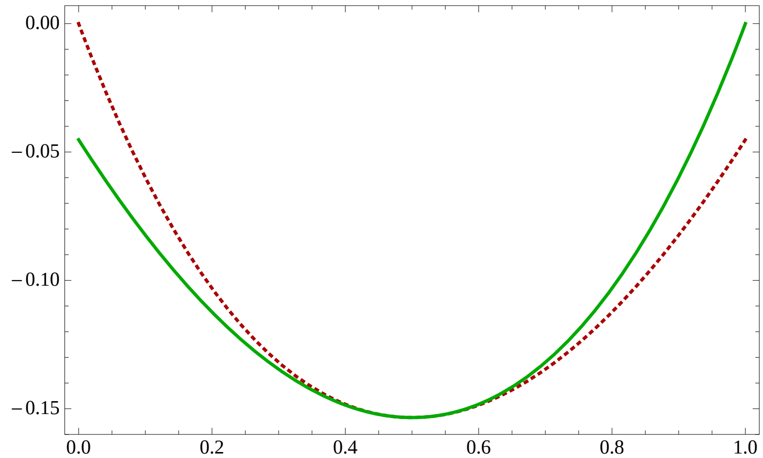

Figure 1: The ordinal potentials ΔJ_S (dashed) and ΔJ_D (solid) when $m = 0$, $b = 1$, $h = 2$, $B = 3$, and $M = 4$.

This criterion is easy to evaluate. Under imitation of success (O.1), the value of the ordinal potential function at $\chi = 1$ is

$$J_S(1) = \int_0^1 \log \frac{F_1(y) - m}{F_0(y) - m}\,\mathrm{d}y = \log \frac{B - m}{h - m} + \frac{b - m}{B - b} \log \frac{B - m}{b - m} - 1.$$

Under imitation driven by dissatisfaction (O.3), this value is

$$J_D(1) = \int_0^1 \log \frac{M - F_0(y)}{M - F_1(y)}\,\mathrm{d}y = \log \frac{M - h}{M - B} + \frac{M - b}{B - b} \log \frac{M - B}{M - b} + 1.$$

To make this example more concrete, let us fix $m = 0$, $b = 1$, $h = 2$, $B = 3$, and $M = 4$. Then the mixed Nash equilibrium of F is $\chi^* = \frac{h-b}{B-b} = \frac{1}{2}$, implying that neither strategy is strictly risk dominant; obviously, the all-Boar equilibrium $\chi = 1$ is payoff dominant. The functions ΔJ_S and ΔJ_D are graphed in Figure 1. Since $J_S(1) = \frac{3}{2}\log 3 - \log 2 - 1 \approx -.0452$, imitation of success leads to the selection of the payoff dominated all-Hare equilibrium, $\chi = 0$. But since and $J_D(1) = -J_S(1) \approx .0452$, imitation driven by dissatisfaction leads to the selection of the payoff dominant all-Boar equilibrium, $\chi = 1$. Clearly, the same qualitiative results obtain in any game with similar payoff values, implying that either protocol can select or fail to select a strictly risk dominant equilibrium.

How do the differences between the protocols lead to the different equilibrium selections? In general, the identity of the stochastically stable state is determined by the relative unlikelihoods of excursions between equilibria: here, from all-Hare to all-Boar, and from all-Boar to all-Hare. The mean dynamic (M_0), computed as the *differences* between the probability $p_\chi \approx \chi(1-\chi)r_{01}(F(\chi))$ of a switch from 0 to 1 and the probability

$q_\chi \approx \chi(1-\chi)r_{10}(F(\chi))$ of a switch from 1 to 0, provides some information about the relative unlikelihoods of excursions. But the stationary distribution (8), being the limiting distribution of the Markov process, is determined by *products of ratios* of the one-step transition probabilities. Ultimately, the stationary distribution weights depend on the ratios of the imitation probabilities $r_{01}(\cdot)$ and $r_{10}(\cdot)$, or, equivalently, on the differences between the logarithms of $r_{01}(\cdot)$ and $r_{10}(\cdot)$. This is evident in the definition (11) of the ordinal potential function J, which allows us to write $J(1)$ as

$$J(1) = \int_0^1 \log r_{01}(F(y))\,\mathrm{d}y \; - \int_0^1 \log r_{10}(F(y))\,\mathrm{d}y. \tag{O.4}$$

Equation (O.4) says that the stochastically stable state is determined by comparing the averages over population states χ of $\log r_{01}(F(\chi))$ and $\log r_{10}(F(\chi))$. Since we are averaging logarithms of imitation probabilities, the manner in which the levels of dispersion of these probabilities depend on the direction of imitation—from Hare to Boar or from Boar to Hare—plays a basic role in explaining infinite-horizon behavior.

As χ varies from 0 to 1, the payoff to Hare is constant at $F_0(\chi) = 2$, while the payoff to Boar, $F_1(x) = 2\chi + 1$, is uniformly distributed between 1 and 3. Under imitation of success (O.1), the conditional imitation probabilities are given by $r_{01}(\chi) = \lambda(F_1(\chi) - m) = \lambda F_1(\chi)$ and $r_{10}(\chi) = \lambda(F_0(\chi) - m) = 2\lambda$.[2] If we view $F_1(\chi)$ as a random variable with a *uniform*$[1,3]$ distribution, then $r_{01}(\chi)$ has a *uniform*$[\lambda, 3\lambda]$ distribution. Thus, since the logarithm function is concave, Jensen's inequality implies that the first integral in (O.4) is smaller than the second, and hence that the all-Hare equilibrium is stochastically stable. Intuitively, the conditional imitation probabilities in each direction have the same mean, but only those from Hare to Boar are variable. This implies that the Hare to Boar transition is less likely than its opposite, and so that the all-Hare equilibrium is selected.

Under imitation driven by dissatisfaction (O.3), the conditional imitation rates are given by $r_{01}(\chi) = \lambda(M - F_0(\chi)) = 2\lambda$ and $r_{10}(\chi) = \lambda(4 - F_1(\chi))$. This time, the conditional imitation probabilities from Hare to Boar are fixed, while those from Boar to Hare are variable; the latter transition is therefore less likely, and the all-Boar equilibrium is stochastically stable. ◆

Maruta [3] also considers a stochastic evolutionary model in which decision rules focusing on the payoff of the current strategy favor selection of the payoff dominant equilibrium, while decision rules focusing on the payoff of the alternative strategy favor selection of the "safe" equilibrium. Maruta's [3] model differs from ours in a number of

[2] Recall that the constant $\lambda > 0$ was introduced in (O.1) to ensure that $r_{01}(\chi)$ and $r_{10}(\chi)$ are less than 1.

important respects: for example, choices in his model are not based on imitation, and his analysis concerns the small noise limit rather than the large population limit.

O.2. The proof of Theorem 4.1

When there are no committed agents and a mutation rate of ε as described in protocol (3), the stationary distribution of the stochastic evolutionary process takes the form

$$\frac{\mu_\chi^{N,\varepsilon}}{\mu_0^{N,\varepsilon}} = \prod_{j=1}^{N\chi} \frac{p_{(j-1)/N}^{N,\varepsilon}}{q_{j/N}^{N,\varepsilon}} = \prod_{j=1}^{N\chi} \frac{\frac{N-j+1}{N}\left((1-\varepsilon)\frac{j-1}{N-1}\, r_{01}(F^N(\frac{j-1}{N})) + \varepsilon\right)}{\frac{j}{N}\left((1-\varepsilon)\frac{N-j}{N-1}\, r_{10}(F^N(\frac{j}{N})) + \varepsilon\right)}. \tag{O.5}$$

To prove the first statement in part (i), we expand equation (O.5) with $\chi = 1$ to obtain

$$\begin{aligned}\frac{\mu_1^{N,\varepsilon}}{\mu_0^{N,\varepsilon}} &= \frac{N\varepsilon}{(1-\varepsilon)r_{10}(F^N(\frac{1}{N})) + \varepsilon} \cdot \prod_{j=2}^{N-1} \frac{N-j+1}{j} \cdot \frac{(1-\varepsilon)\frac{j-1}{N-1}\, r_{01}(F^N(\frac{j-1}{N})) + \varepsilon}{(1-\varepsilon)\frac{N-j}{N-1}\, r_{10}(F^N(\frac{j}{N})) + \varepsilon} \cdot \frac{(1-\varepsilon)r_{01}(F^N(\frac{N-1}{N})) + \varepsilon}{N\varepsilon} \\ &= \prod_{j=1}^{N-1} \frac{N-j}{j} \cdot \frac{(1-\varepsilon)\frac{j}{N-1}\, r_{01}(F^N(\frac{j}{N})) + \varepsilon}{(1-\varepsilon)\frac{N-j}{N-1}\, r_{10}(F^N(\frac{j}{N})) + \varepsilon}.\end{aligned}$$

Since the $N\varepsilon$ terms cancel, none of the terms in the final product converge to zero or infinity. Indeed, we have

$$\lim_{\varepsilon\to 0} \log \frac{\mu_1^{N,\varepsilon}}{\mu_0^{N,\varepsilon}} = \log\left(\prod_{j=1}^{N-1} \frac{N-j}{j} \cdot \frac{\frac{j}{N-1}\, r_{01}(F^N(\frac{j}{N}))}{\frac{N-j}{N-1}\, r_{10}(F^N(\frac{j}{N}))}\right) = \sum_{j=1}^{N-1} \log \frac{r_{01}(F^N(\frac{j}{N}))}{r_{10}(F^N(\frac{j}{N}))},$$

so bound (B) and the bounded convergence theorem imply that

$$\lim_{N\to\infty} \lim_{\varepsilon\to 0} \frac{1}{N} \log \frac{\mu_1^{N,\varepsilon}}{\mu_0^{N,\varepsilon}} = \lim_{N\to\infty} \frac{N-1}{N} \cdot \frac{1}{N-1} \sum_{j=1}^{N-1} \log \frac{r_{01}(F^N(\frac{j}{N}))}{r_{10}(F^N(\frac{j}{N}))} = J(1).$$

To prove the second statement in part (i), observe that for $\chi \in \mathcal{X}^N - \{0,1\}$,

$$\frac{\mu_\chi^{N,\varepsilon}}{\mu_0^{N,\varepsilon}} = \frac{N\varepsilon}{(1-\varepsilon)r_{10}(F^N(\frac{1}{N})) + \varepsilon} \cdot \prod_{j=2}^{N\chi} \frac{N-j+1}{j} \cdot \frac{(1-\varepsilon)\frac{j-1}{N-1}\, r_{01}(F^N(\frac{j-1}{N})) + \varepsilon}{(1-\varepsilon)\frac{N-j}{N-1}\, r_{10}(F^N(\frac{j}{N})) + \varepsilon}$$

by equation (O.5). Since all terms except the initial $N\varepsilon$ approach positive constants as ε approaches zero, $\mu_\chi^{N,\varepsilon}/\mu_0^{N,\varepsilon}$ is in $\Theta(\varepsilon)$. The analysis of $\mu_\chi^{N,\varepsilon}/\mu_1^{N,\varepsilon}$ is similar.

To prove part (ii) of the theorem, we note from equation (O.5) that

$$\frac{1}{N}\log\frac{\mu_{\chi}^{N,\varepsilon}}{\mu_{0}^{N,\varepsilon}} = \frac{1}{N}\sum_{j=1}^{N\chi}\left(\log\frac{N-j+1}{N} - \log\frac{j}{N} + \log\frac{(1-\varepsilon)\frac{j-1}{N-1}\,r_{01}(F^{N}(\frac{j-1}{N})) + \varepsilon}{(1-\varepsilon)\frac{N-j}{N-1}\,r_{10}(F^{N}(\frac{j}{N})) + \varepsilon}\right).$$

Comparing this equation to equation (19) and the subsequent arguments, it is easy to verify that the remainder of the proof is identical to that of Theorem 4.3(ii). ■

O.3. The proof of the Theorem (4.4)

By substituting the committed agents protocol (4) with $c_i^N = N\gamma_i$ into equations (6) and (7), we obtain the one-step transition probabilities

$$p_{\chi}^{N} = (1-\chi)\cdot\tfrac{N(\chi+\gamma_1)}{N(1+\gamma_T)-1}r_{01}(F^{N}(\chi)) \quad \text{and}$$
$$q_{\chi}^{N} = \chi\cdot\tfrac{N(1-\chi+\alpha_0)}{N(1+\gamma_T)-1}r_{10}(F^{N}(\chi)),$$

where $\gamma_T = \gamma_0+\gamma_1$. Inserting these expressions into equation (8), we find that the stationary distribution of the process $\{X_t^N\}$ is given by

$$\frac{\mu_{\chi}^{N}}{\mu_{0}^{N}} = \prod_{j=1}^{N\chi}\frac{p_{(j-1)/N}^{N}}{q_{j/N}^{N}} = \prod_{j=1}^{N\chi}\frac{\frac{N-j+1}{N}\cdot\frac{j+N\gamma_1-1}{N(1+\gamma_T)-1}\,r_{01}(F^{N}(\frac{j-1}{N}))}{\frac{j}{N}\cdot\frac{N-j+N\gamma_0}{N(1+\gamma_T)-1}\,r_{10}(F^{N}(\frac{j}{N}))}. \tag{O.6}$$

It follows that

$$\begin{aligned}\frac{1}{N}\log\frac{\mu_{\chi}^{N}}{\mu_{0}^{N}} &= \frac{1}{N}\sum_{j=1}^{N\chi}\left(\log\frac{r_{01}(F^{N}(\frac{j-1}{N}))}{r_{10}(F^{N}(\frac{j}{N}))} + \log\frac{\frac{N-j+1}{N}\cdot\frac{j+N\gamma_1-1}{N(1+\gamma_T)-1}}{\frac{j}{N}\cdot\frac{N-j+N\gamma_0}{N(1+\gamma_T)-1}}\right)\\ &= J^{N}(\chi) + \frac{1}{N}\sum_{j=1}^{N\chi}\left(\log\tfrac{N-j+1}{N} - \log\tfrac{j}{N} + \log\tfrac{j+N\gamma_1-1}{N(1+\gamma_T)-1} - \log\tfrac{N-j+N\gamma_0}{N(1+\gamma_T)-1}\right).\end{aligned} \tag{O.7}$$

The proof of Theorem 4.2 shows that J^N converges uniformly to J. Moreover, applying the dominated convergence theorem as in the proof of Theorem 4.3 shows that the second term of (O.7) converges uniformly to

$$\begin{aligned}&\int_0^{\chi}\left(\log(1-y) - \log y + \log\tfrac{y+\gamma_1}{1+\gamma_T} - \log\tfrac{1-y+\gamma_0}{1+\gamma_T}\right)\mathrm{d}y\\ &\quad= -\chi\log\chi - (1-\chi)\log(1-\chi) + (\chi+\gamma_1)\log(\chi+\gamma_1) - \gamma_1\log\gamma_1\\ &\qquad+ (1-\chi+\gamma_0)\log(1-\chi+\gamma_0) - (1+\gamma_0)\log(1+\gamma_0)\end{aligned}$$

$$= L^{\gamma}(\chi).$$

Thus since $J^{\gamma}(\chi) = J(\chi) + L^{\gamma}(\chi)$, we find that

$$\lim_{N\to\infty} \max_{\chi \in \mathcal{X}^N} \left| \frac{1}{N} \log \frac{\mu_{\chi}^N}{\mu_0^N} - J^{\gamma}(\chi) \right| = 0.$$

The proof is completed by combining this conclusion with a slight variation on the last part of the proof of Theorem 4.2. ■

References

[1] J. Björnerstedt, J. W. Weibull, Nash equilibrium and evolution by imitation, in: K. J. Arrow, et al. (eds.), The Rational Foundations of Economic Behavior, St. Martin's Press, New York, 1996, pp. 155–181.

[2] J. Hofbauer, Imitation dynamics for games, unpublished manuscript, University of Vienna (1995).

[3] T. Maruta, Binary games with state dependent stochastic choice, Journal of Economic Theory 103 (2002) 351–376.

[4] P. D. Taylor, L. Jonker, Evolutionarily stable strategies and game dynamics, Mathematical Biosciences 40 (1978) 145–156.

[5] J. W. Weibull, Evolutionary Game Theory, MIT Press, Cambridge, 1995.