Computational Economics **19:** 25–49, 2002.

Evolutionary Models of Bargaining: Comparing Agent-based Computational and Analytical Approaches to Understanding Convention Evolution

JEFFREY P. CARPENTER
Department of Economics, Middlebury College, Middlebury, VT 05753, U.S.A.
E-mail: jpc@middlebury.edu

Abstract. This paper compares two methodologies that have been used to understand the evolution of bargaining conventions. The first is the analytical approach that employs a standard learning dynamic and computes equilibria numerically. The second approach simulates an environment with a finite population of interacting agents. We compare these two approaches within the context of three variations on a common model. In one variation agents randomly experiment with different demands. A second variation posits assortative interactions, and the third allows for sophistication in agent strategies. The simulation results suggest that the agent-based approach performs well in selecting equilibria in most instances, but exact predicted population distributions often vary from those calculated numerically. Classification Numbers: C63, C73, C78.

Key words: bargaining, convention, simulation, Agent-Based, fairness

1. Introduction

As a result of advances in computer technology, analytical models which were previously discarded because closed form solutions could not be calculated, are now being solved computationally by simulation. This achievement has spawned new interest in complicating simple models to improve their descriptive ability. However, researchers make choices when they develop computational tools to compute equilibria. One important subject, that has been given little attention, has to do with the extent to which various numerical methods provide similar results when controlling for a given model's basic institutional environment. This paper explores this question by comparing three evolutionary bargaining models using two forms of simulation. The first method iterates the discrete analogue of the differential equations used to model the evolution of different bargaining strategies. The second method is the more decentralized agent-based method which develops a bargaining environment and then populates it with agents who adopt strategies according to a simple learning rule.

Another related issue arises when only one of the two methods mentioned above is an effective method to pursue. More particularly, it is often the case that interesting questions cannot be answered adequately by simply writing down a system of dynamical equations and therefore the method of numerical iteration cannot be helpful. Alternatively, the underlying dynamics may be so complicated that simulating them in difference form may be either impractical or impossible (e.g. stochastic dynamical systems)[1]. In these situations, it would be interesting to know that the agent-based approach can be shown to generate the same behavior as the numerical iteration approach in a simpler environment.

Lastly, there is one other issue that is of particular importance for the evolutionary models to be discussed in more detail below. The replicator dynamic used in many evolutionary models is based on the assumption that populations are infinite. Therefore, by employing the law of large numbers, expected payoffs can be treated as actual payoffs in fitness calculations. However, it is obvious that the agent-based approach must be developed with finitely many agents. If the agent-based approach is an appropriate way to numericaly calculate equilibria in evolutionary models, then it must be shown that finite populations can approximate the behavior of an infinite population. The comparison of the agent-based approach and the numerical iteration approach using a common theoretical model as a basis allows us to explore all three questions.

The three models, that have been simulated, all have been developed in greater depth elsewhere to analyze the evolution of the equal split as a bargaining convention and use a common structure, the Nash demand game. Because all three models share the structure of the same underlying normal form game and are developed to answer a common question, they provide a sound basis for making comparisons between the two simulation techniques. Additionally, using three variants of the underlying model is useful, because it allows us to stress-test each method by changing the model slightly.

Before we begin the analysis we will say a few words about the bargaining models and about the reasons for evolutionary models in particular. Although many variants of the Stahl-Rubinstein non-cooperative theory of bargaining (Stahl, 1972; Rubinstein, 1982) exploit externally given differences in negotiator preferences or use stylized descriptions of asymmetries in the bargaining environment to derive equilibria, the models developed below exclude these factors and focus on the tradeoff between coordination and payoff maximization. The most appropriate game structure on which to base models of this sort is Nash's demand game (Nash, 1953). In this game two agents make demands for a share of some resource. If the sum of their demands is less than or equal to the resource, then they are both satisfied, otherwise they both receive nothing. This formulation of the bargaining problem is interesting because, in absence of asymmetries between players, unequal outcomes can be sustained by the fact that players are concerned about coordinating to take full advantage of the resource. As a result, although the equal split is a bargaining convention that can arise in this game, it is not the only convention that can emerge[2].

Evolutionary models are particularly interesting for the comparison we wish to make because they lend themselves to being simulated in more than one way. Because non-linear dynamical equations are regularly used to describe the time paths of a population distribution, these models often must be solved numerically by iterating discrete time analogues or by developing agents who interact according to the underlying structure imposed by the dynamic employed. Additionally, the evolutionary approach is often adopted because such models do not rely on informational or rationality assumptions made in standard game-theoretic models. Agents in standard models know the preferences of their bargaining counterpart and understand that their counterparts will act in accordance with these preferences. By comparison, agents in evolutionary models need not know how their counterpart orders outcomes or how they will respond to different allocations – they are simply born to blindly play a strategy. Stategies, and the agents who play them, flourish in an environment if the strategy does better, on average, than the other strategies present. The fact that agents in evolutionary models know nothing about their bargaining partner, remember nothing about previous interactions (i.e. they have zero-intelligence), and blindly make demands (i.e. they are relatively unsophisticated) makes them easy to simulate in both the ways we are interested in.

Although the models of Skyrms (1994), Ellingsen (1997), and the one we develop below all use the Nash demand game as a structure to base an evolutionary theory of the equal split on, each is unique in the institutional assumptions that drive the result. Skyrms (1994) models assortative interactions between agents. When agents are more likely than chance would predict to interact with a like agent, interactions between agents who demand less than half the surplus always result in inefficient agreements. Interactions between agents who demand half are both efficient and avoid conflict, and interactions between agents demanding more than half the pie always result in impasse.

Hence, demanding half the pie does better, on average, than other strategies. Ellingsen (1997) relies on a static notion of stability developed by Maynard Smith and Price (1973) to describe bargaining conventions that arise in a slightly more sophisticated environment. In this model agents can be obstinate and make demands as in the original specification of the game, or they can be responsive. Responsive agents are sophisticated in that they avoid conflict by claiming whatever is left over after obstinate agents have satisfied their demands. Finally, as a third extension of the standard Nash demand game Section 2.2 formulates a model in which a small fraction of agents experiment by randomly changing their demands. In this case, the only stable convention, as the amount of experimentation becomes negligible, is the equal split.

Section 2 develops a generic evolutionary model of the Nash demand game to provide the basis for the three models and then presents summaries of the three models just mentioned, starting with the model of experimentation. For the sake of expositional continuity and to control for the underlying structure of bargaining, each of the three models will be developed as an extension of the generic model. Section 3 reports the details of how the simulations were programmed and conducted. Section 4 concludes by discussing how well the agent-based approach approximates the behavior found using numerical techniques.

2. Three Models of the Equal Split

In this section we outline three evolutionary models of bargaining conventions. We start by developing a genetic bargaining environment and by introducing the notions of stability that we will use to analyze equilibria in three alternations of the generic model. In the first case, the model is altered by allowing agents to experiment with different strategies. In the second case, interactions become assortative, in that like types meet more often than would be expected by random chance. Finally, in the third model the strategy set is altered to allow for more sophisticated behavior. Here agents are sophisticated to the extent that they defer to their counterpart's demand and therefore always avoid conflict.

2.1. THE GENERIC MODEL

The bargaining environment is characterized by an infinite population of negotiators who make demands for a share of some renewable resource that has a value of 1 unit. Further, partners in this interaction are determined randomly and their demands are taken from a finite set of all the possible fractional divisions of the unit resource. This set will be denoted as D. Interactions are structured as follows. Agents simultaneously make demands for a share of the resource. For example, say one agent demands x and the other demands y where $x,y \in D$. If both demands sum to an amount that is less than or equal to 1, then each agent gets their demanded share. Otherwise, both get nothing. If the game is played only using pure strategies, the number of Nash equilibria will be determined by the number of elements in D, where each pair of equilibrium demands sums to 1.

For the Nash demand game to be interesting, we only need three pure strategies, one where agents demand something less than one-half, one where agents demand exactly one-half, and a third where agents demand one minus the lower demand. That is, the game remains interesting where $D = \{1/3, 1/2, 2/3\}$ because it retains the important balance between payoff maximization and coordination. We will call this three-pure strategy game the *finite Nash demand game*. The normal form of this game is illustrated as Figure 1 wherein one can see the three pure strategy Nash equilibria that lie along the diagonal from the southwest corner of the matrix to the northeast corner[3].

Given this bargaining structure we can now use the standard tools of evolutionary game theory to analyze equilibria. The *fitness* of a demand is determined by

	1/3	1/2	2/3
1/3	1/3, 1/3	1/3, 1/2	1/3, 2/3
1/2	1/2, 1/3	1/2, 1/2	0, 0
2/3	2/3, 1/3	0, 0	0, 0

Figure 1. The finite Nash demand game.

the expected benefit it confers on an agent who randomly interacts with another agent. For example, the fitness of demanding 1/3 is determined both by the payoff received from meeting each of the three types in the population, as well as by the the likelihood of meeting each type. More formally, we can define the finite Nash demand game as $\mathcal{D} = \langle D, \pi(x, y)\rangle$, where D is the set of strategies defined above and $\pi(x, y)$ is the payoff of an agent demanding x who meets another agent demanding y for $x,y \in D$. Further, let a be a vector of probabilities, a_y being the probability of demanding y. Then the fitness, $F(\cdot)$ of a demand is equal the sum of the payoffs to meeting each agent type weighted by the probability of meeting said type or.

$$F(x) = \sum_{y \in D} a_y \pi(x, y) \qquad \text{where} \qquad \sum_{y \in D} a_y = 1\ . \tag{1}$$

Nash equilibria of $\mathcal{D}$ are found by calculating probability distributions that equate the payoff to playing each of the pure strategies that are used with positive probability. Using this method, all agents demanding 1/2 and agents demanding 1/3 with probability one-half and demanding 2/3 with probability one-half are Nash equilibria of the current game. Call the first equilibrium the *symmetric* Nash equilibrium and the second the *asymmetric* Nash equilibrium. There is also an *interior* Nash equilibrium where all three demands are used with positive probability. Here the probability distribution $(a_{1/3}, a_{1/2})$ that equalizes the payoff to demanding 1/3, the payoff to demanding 1/2, and the payoff to demanding 2/3 is (1/2, 1/6). The existence of three equilibria makes it difficult to predict where the population will actually end up. To deal with this problem we now introduce two notions of stability that will be used to narrow the set of equilibria.

Maynard Smith (1982) refines the concept of a Nash equilibrium for a population of agents who replicate themselves over the course of many generations in frequencies that depend on the outcome of pair-wise interactions. In this formulation the notion of fitness determines how many offspring an agent passes to the next generation. A strategy is said to be *evolutionarily stable* (*ESS*) if, when most of the population uses the strategy each agent using the strategy has higher fitness than any agent using another strategy[4].

PROPOSITION 1. *No fully supported strategy (i.e., where every pure strategy is played with positive probability) can be an ESS of the finite Nash demand game* $\mathcal{D} = \langle D, \pi(x, y) \rangle$.
Proof. See Appendix.

The implication of Proposition 1 is that the internal Nash equilibrium is not an ESS of $\mathcal{D}$. However, calculating the payoffs to playing the asymmetric Nash and the symmetric Nash against each other, against the internal equilibrium strategy, and against all the pure strategies shows that these two Nash equilibria are supported as ESSs of $\mathcal{D}$. The stability of these two equilibria can be determined in another more general way – replicator dynamics.

Replicator dynamics are often preferred because they do not require agents to calculate mixed strategies. That is, now assume that each agent is hard-wired to make a specific demand. In particular, a fraction $a_{1/3}$ of the agents demands 1/3 of the resource no matter who they are paired with. Another fraction $a_{1/2}$ always demand 1/2 and the remaining $1 - a_{1/3} - a_{1/2}$ of the agents demand 2/3. As stated in Section 1, the population is assumed to be large so that using the law of large numbers we can equate expected outcomes and actual outcomes. Now following Appendix D of Maynard Smith (1982) define *average fitness*, $\bar{F}$ as the mean fitness of the population such that

$$\bar{F} = \sum_{y \in D} a_y F(y) \ . \tag{2}$$

Because average fitness is a measure of the productivity of the population as a whole, the frequency of agent type x in the next generation is determined by how productive demanding x is in relationship to the average demand. Assuming that the differences between generations are not too large we can express the time paths of the population of agents as

$$\frac{\mathrm{d}a_x}{\mathrm{d}t} \equiv \dot{a}_x = a_x(F(x) - \bar{F}) \ . \tag{3}$$

Equation (3) is the basic replicator dynamic and can be interpreted to mean that the rate of growth of x demanders in the population is proportional to the fitness of demanding x relative to the average fitness of the current population, given the distribution of types, a.

For the three-strategy Nash demand game the growth rates of demanding 1/3, $\mathrm{d}a_{1/3}/\mathrm{d}t$ and demanding 1/2, $\mathrm{d}a_{1/2}/\mathrm{d}t$ are

$$\frac{\mathrm{d}a_{1/3}}{\mathrm{d}t} \equiv \dot{a}_{1/3} = a_{1/3}[1/3 - (a_{1/3} - a_{1/3}a_{1/2}/6 - 2a_{1/3}^2/3 + a_{1/2}^2/2)] \tag{4}$$

and

$$\frac{da_{1/2}}{dt} \equiv \dot{a}_{1/2} =$$

$$a_{1/2}[(a_{1/3} + a_{1/2})/2 - (a_{1/3} - a_{1/3}a_{1/2}/6 - 2a_{1/3}^2/3 + a_{1/2}^2/2)] \,. \qquad (5)$$

A stationary state of this dynamical system occurs when both of the rates of growth approach zero and therefore an *evolutionary equilibrium* for the finite Nash demand game arises where $\dot{a}_{1/3} = \dot{a}_{1/2} = 0$.

Recall that the criteria for an ESS in the finite Nash demand game eliminated any fully supported mixed strategy. Using the dynamic system that we have just developed, we can prove that the internal Nash equilibrium is not a stable fixed point of the replicator dynamics either[5].

PROPOSITION 2. *The asymmetric and the symmetric Nash equilibria are asymptotically stable evolutionary equilibria of the replicator dynamics described by (4) and (5).*
Proof. See Appendix.

In this section we have developed an environment in which strategies compete for a finite resource and replicate themselves according to a simple dynamical learning rule – look around and see how well you are doing with respect to the average. Additionally, we have seen that within this structure two stable equilibria (conventions) can evolve and another non-stable equilibrium is not sustainable. As such, this provides a good foundation to explore the performance of the two simulation methods. However, to add additional pitfalls for the simulations, we will first look at three different refinements that have been developed to single out the equal split as the only stable convention. We begin by exploring a stronger stability requirement.

2.2. SYSTEMICALLY STABLE BARGAINING CONVENTIONS

As a first adjustment to the basic model, we consider a stronger stability condition related to the notion of stochastic stability develop originally in Young (1993). The basic premise is as follows. The requirement used to test the stability of equilibria in the genetic model (using either the static notion of ESS or the dynamic notion of the asymptotic stability of a fixed point) boils down to a test of whether an equilibrium is what we call *invasion proof*. An invasion can be thought of as a small perturbation away from an equilibrium population distribution where a small number of agents adopt some other strategy. An equilibrium is said to be stable if the population distribution does not continue to diverge from the equilibrium after invasion (i.e. the invaders do not do better than the general population, on average). Further an equilibrium is said to be asymptotically stable if the invasion can be repelled in the sense that invaders do worse on average than the general population and therefore the system returns to the original equilibrium.

With these definitions in mind, consider an alteration of the generic scenario. Rather than a small one-time invasion, what if perturbations occur randomly and repeatedly so that the system does not necessarily recover before another perturbation? In this case perturbations are systemic. To visualize the situation, suppose that a small fraction of the population experiments with different strategies by randomly choosing demands or by making mistakes in each period. In this case, a population distribution is said to be *systemically stable* if the system is certain to stay in the neighborhood of the equilibrium population distribution for small rates of experimentation.

We can now check for systemic stability in the finite Nash demand game by altering Equations (4) and (5) in the following way. Assume, for whatever reason (experimentation or mistake), a fraction d of the population disregard the strategy they are born with and randomly choose among the strategies in D. Further, the remaining $(1 - d)$ of the agents make demands as before. In this case we have

$$\dot{a}_{1/3} = (1-d)a_{1/3}[1/3 - (a_{1/3} - a_{1/3}a_{1/2}/6- \\ 2a_{1/3}^2/3 + a_{1/2}^2/2)] + d(1/3 - a_{1/3}) \tag{6}$$

and

$$\dot{a}_{1/2} = (1-d)a_{1/2}[(a_{1/3} + a_{1/2})/2 - (a_{1/3} - a_{1/3}a_{1/2}/6- \\ 2a_{1/3}^2/3 + a_{1/2}^2/2)] + d(1/3 - a_{1/2})\ . \tag{7}$$

The last term following the plus sign in each equation indicators that a fraction, d of the population choose a strategy randomly. To see this let $d \to 1$ and find the fixed points of the system. It is easy to see that there is an asymptotically stable equilibrium, where each strategy is played by one-third of the population – exactly what would be expected from a population making demands randomly.

PROPOSITION 3. *Only the equal slit is asymptotically stable for levels of experimentation greater than 0.013*[6].
Proof. See Appendix.

In altering the generic bargaining model in a very reasonable way be allowing for systemic perturbations to equilibrium, we see that only the equal split is robust to the inertia the system is subjected to when agents experiment or make mistakes. We shall see whether this refinement poses a challenge for the agent-based simulation in Section 4.

2.3. ASSORTATIVE INTERACTIONS

Now we ask what happens to the set of equilibrium population distributions when interactions are somewhat assortative[7]. Interactions are assortative in the current model because there is some positive correlation between pairing and type. In particular, let there be gravity between like types so that agents demanding x are more likely to meet other x demanders than would be expected by random chance. Formally, let z be the probability that like types meet. This changes the dynamic, (3). Remember that the payoff to demanding x against another x demander is $\pi(x, x)$ which means that $(1 - z)$ of the x demanders will have fitness $F(x)$ calculated as in (1) and a fraction z of the x demanders will have fitness $\pi(x, x)$. Combining these facts and again assuming the population does not change too much from generation, we have

$$\frac{\mathrm{d}a_x}{\mathrm{d}t} \equiv \dot{a}_x = a_x[(1 - z)F(x) - z\pi(x, x) - \bar{F}] \,. \tag{8}$$

Notice that as the amount of assortation becomes large, the growth rate of strategies depends only on the difference between $\pi(x, x)$ and the average fitness of the population. Clearly, this difference is maximized by agents using the equal split rule. However, we are interested in a different result. That is, what is the critical value of z above which only one-half demanders survive?

Changing (4) and (5) in accordance with (8) we have

$$\dot{a}_{1/3} = a_{1/3}[1/3 - (a_{1/3} - a_{1/3}a_{1/2}/6 - 2a_{1/3}^2/3 + a_{1/2}^2/2)] \tag{9}$$

and

$$\dot{a}_{1/2} = a_{1/2}[z(1 - a_{1/3} - a_{1/2})/2 + (a_{1/3} + a_{1/2}q)2 - (a_{1/3} - a_{1/3}a_{1/2}/$$
$$6 - 2a_{1/3}^2/3 + a_{1/2}^2/2)] \,. \tag{10}$$

Notice that there is no difference between (9) and (4) because agents demanding 1/3 get one-third of the surplus no matter whom they interact with.

PROPOSITION 4. *Only the equal split is asymptotically stable under assortative interactions when* $z > 1/3$.
Proof. See Appendix.

Proposition 4 demonstrates that when it is sufficiently probable that like types interact, modest demands lead to inefficient bargaining and greedy demands too often end in conflict. The end result is a stable (invasion-proof) population of agents demanding half the surplus.

	0	⅓	½	⅔	1	r
0	1/2, 1/2	0, 1	0, 1	0, 1	0, 1	0, 1
⅓	1, 0	1/2, 1/2	2/5, 3/5	1/3, 2/3	0, 0	1/3, 2/3
½	1, 0	3/5, 2/5	1/2, 1/2	0, 0	0, 0	1/2, 1/2
⅔	1, 0	2/3, 1/3	0, 0	0, 0	0, 0	2/3, 1/3
1	1, 0	0, 0	0, 0	0, 0	0, 0	1, 0
r	1, 0	2/3, 1/3	1/2, 1/2	1/3, 2/3	0, 1	1/2, 1/2

Figure 2. The finite Nash demand game with strategy set D'.

2.4. SOPHISTICATED AGENTS

As a third variation from the genetic model, consider a larger strategy set that includes the demands of more sophisticated agents[8]. Here sophistication means that some agents do not obstinately make demands, rather they accept the demands of their counterpart and react by claiming any residual. Therefore, these *responsive* agents always avoid conflict. To stay as close as possible to the original version of this model we will also admit two more strategies, the ultra modest strategy of demanding nothing and the greedy strategy of demanding all the surplus. In addition, we use the static notions of stability (ESS and NSS) to analyze the evolution of this expanded group of negotiators.

Formally, let D' be the expanded set of demands such that $D' = \{0, 1/2, 1/2, 2/3, 1, r\}$ where $r = 1 - x$ for all $x \neq r$ and $x \in D'$. We also need to re-specify the payoff function so that it matches that of Ellingsen (1997). In this model we consider only fully efficient negotiations such that

$$\pi(x, y) = \begin{cases} \frac{x}{x+y}, & \text{if } x + y \leq 1 \\ 0, & \text{otherwise}. \end{cases} \tag{11}$$

Lastly, define $\pi(r, r) = 1/2$ so that when two responsive agents meet, they share the surplus equally. Figure 2 illustrates the normal form game that represents this model. Using the concept of iterative elimination of weakly dominated strategies, we find a unique solution where all agents are greedy (i.e. demand the whole surplus). However, it is not necessarily the case that evolution eliminates weakly dominated strategies (Samuelson and Zhang, 1992; Samuelson, 1997; Mailath, 1998). We shall return to the question of dominance when the simulation results are discussed in Section 4. For now, we are interested in whether the equal split is viable.

PROPOSITION 5. *A population distribution, a is neutrally stable if and only if* $a_{1/2} > 1/2$ *and* $a_r = 1 - a_{1/2}$.
Proof. See Appendix.

In this model agents need to be wary of greedy strategies that are viable in populations that contain responsive agents because responsive agents seek to avoid conflict at all cost. As the simulations will show, responsive agents are both a blessing and a curse. They add to the support of the equal split which is good for the simple reason that the level of conflict (efficiency) in such populations is at a minimum (maximum). However, they also contribute to the fragility of the equal plit because they free-ride on obstinate fair agents who punish greedys. This summarizes the last pitfall for which the simulations will need to account.

3. Numerical Iteration and Agent-Based Simulations

As mentioned above, we choose to simulate the finite version of the Nash demand game because it is relatively easy to do so both by iterating the discrete analogues of the replicator dynamics derived above and by simulating the bargaining environment with agents[9]. In this section we will describe in detail the implementation of these two forms of simulation.

3.1. NUMERICAL SOLUTIONS

The numerical simulations simply iterate the replicator dynamic (in difference form) from randomly drawn initial conditions[10]. By running enough sessions, one can sweep the simplex created by the three strategies relatively efficiently. A sample of the code is instructive. What follows is the heart of the code for the numerical simulation of the systemic model of Section 2.2 (here $p = a_{1/3}$ and $q = a_{1/2}$) (Scheme 1).

```
Do
    t = t + 1
    q(t) = pn
    q(t) = qn

    FitP = 1 / 3
    FitQ = 0.5 * (p(t) + q(t))
    FitAvg = FitP * p(t) + FitQ * q(t) + (2 / 3) * p(t) * (1 - p(t) - q(t))

    pn = (((1 – d) * p(t) * (FitP – FitAvg) + d * ((1 / 3) – p(t))) / FitAvg) + p(t)
    qn = (((1 – d) * q(t) * (FitQ – FitAvg) + d * ((1 / 3) – q(t)))  FitAvg) + q(t)

    dev = Abs(pn – p(t)) + Abs(qn - q(t))

Loop Until dev < 0.0000001
```

Scheme 1.

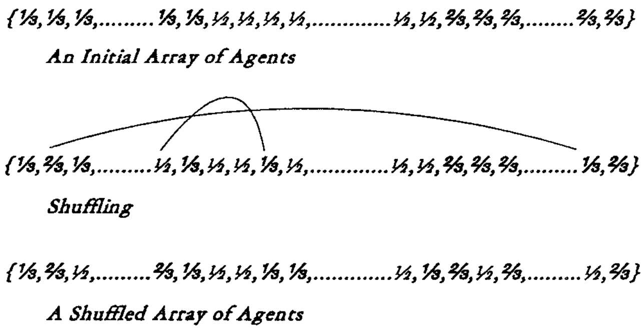

Figure 3. Creating and array of negotiators.

As can be seen from this example, the researcher inputs a value of d, the fraction of agents who experiment in each period, the program randomly picks a seed (not shown) and then the simulation loops until the total deviation in the population distribution from one generation to the next is less than 0.000001. The lack of movement from one generation to the next implies a stable population distribution. This basic code altered for the unique characteristics of each scenario is used to simulate each of the three models from a variety of initial conditions. This procedure allows one to construct phase diagrams for each model.

3.2. AGENT-BASED MODELS

The structure of the agent-based simulations is somewhat complicated and therefore warrants a more detailed discussion. At the beginning of each simulation an array of 1000 agents if created that reflects different starting distributions of the demander types (e.g the number of 1/2 demanders is $a_{1/2} \times 1000$). This array is then shuffled to randomize the order of the agents. Figure 3 illustrates the process of creating an array and shuffling it. Initially the array is populated by $1000 \times a_{1/3}$ one-third demanding agents followed by $1000 \times a_{1/2}$ half demanding agents who, in turn, are followed by $1000 \times a_{2/3}$ two-thirds demanders. This is shown at the top of Figure 3. The shuffling process proceeds by picking two agents randomly to change places in the array. This happens 1000 times so that by the end of shuffling the strategies are randomly distributed along the array.

After shuffling is completed, the agents bargain. Each even numbered agent plays the Nash demand game with their neighbor to the left and the payoffs are recorded. The distribution of types in the agent array for the next generation is determined by the success of each strategy as a whole in the generation before. More specifically, three sums are calculated, one for each demander type, which record the total earnings of each strategy. Agents in the next period are distributed according to the share of total earnings (the sum of all agents' payoffs) each strategy is responsible for. This simulates a process in which agents look around after one period at everyone's payoff and then adopt a strategy in the next period with

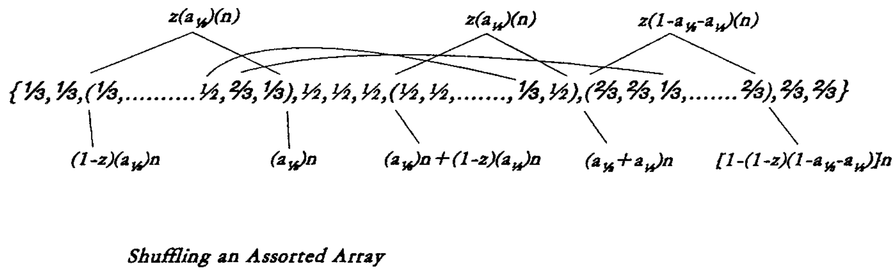

Figure 4. Creating and array of shuffled, assorted negotiators.

probability equal to the strategy's relative success (i.e. the replicator dynamic). This routine was repeated for 100 generations.

For the systemic model of Section 2.2 the process was altered by allowing for mutations. After the initial array was created, but before it was shuffled, a fraction, d of the agents were randomly chosen from the array. These agents were then 're-programmed' to play a strategy randomly and therefore with probability one-third they actually played the original strategy that they were born with. The remainder of the procedure was unchanged.

The process of assortation required that we alter the process of creating an array of negotiators. Although it is not discussed in the assortative model presented above (nor in the original model of Skyrms (1994), another assumption needed to be made in the simulations. As the population share of an unfit demand type approaches zero, two things can happen. One, to maintain the level of assortation, unfit agents can withdraw from the population and begin to only interact with each other. Or two, as the unfit become less prevalent in the population, the probability increases that they will be forced into the general population and therefore will be more likely to meet other types. We will make the second assumption. This point is better understood by illustration.

Figure 4 presents the method of creating an assorted array of negotiators. For the sake of making the simulation easier to program we reverse the notation of Section 2.3. Here, z is the fraction of each demand type that are paired randomly and $(1 - z)$ is the fraction that meet only a like demander. Now, when creating an array of negotiators we set up *home turf* where only like demanders interact. For example, where n is the population size, the left most part of the array shown in Figure 4 is the home turf of the 1/3 demanders. The size of this enclave is determined by z and by the population share of 1/3 demanders. Those agents on home turf cannot be shuffled while the remainder of the population are moved randomly to non-home territory as depicted in Figure 4. As the population share of 1/3 demanders falls, the border of the home turf moves to the left. In the limit, as $a_{1/3}$ approaches zero, the home turf vanishes and 1/2 demanders are just as likely to be shuffled as any other agent outside of their turf.

Lastly, two changes were needed for the agent-based simulations of the sophisticated model (Section 2.4). First, because the sophisticated model utilizes a richer strategy space, the number of strategies was increased from three so six. In particular the strategy set D was expanded so that $D' = \{0, 1/3, 1/2, 2/3, 1, r\}$ was used instead. Also, the payoff structure (Equation (11)) is different and therefore the revised version was also adopted for the simulations of this model.

4. Numerical Analysis vs. Agent-Based Simulations

In this section we discuss the results of the two simulation techniques. The three models we developed in Section 2 enable us to discuss how three different variations on the basic model are treated differently (or not) by the two computational methods we employ. More specifically, the systemic model allows us to see what happens when we introduce gravity towards the center of the simplex. The assortative rule allows us to compare the techniques when we change the strategy pairing rule, and lastly, the sophisticated model introduces the potential pitfalls of an expanded strategy space and weakly-dominated strategies.

We compare the two techniques based on two very general criteria – time to convergence and equilibrium selection. The first is measured in generations to population fixation. The second criteria is a comparison of how well the agent-based simulation picks the same equilibrium from a given starting state as the discrete model which we treat as picking the ‘true’ equilibrium. To quantify the potential differences we use the non-parametric Kolmogorov-Smirnov (*ks*) statistic, which tests whether the data points in the time paths are commonly distributed. Also, when we discus the results of the simulations we shall speak in terms of the evolution of the ‘demand half’ strategy in particular because it is focus of our three models. We also do this so as to not clutter the presentation.

4.1. CONVERGENCE IN THE GENERIC AND SYSTEMIC MODELS

Figure 5 plots the evolution of the equal split in the generic model ($d = 0$) from Section 2.1. As in all the graphs comparing the agent-based (AB) simulations to the numerical simulations, the AB time paths are denoted by dotted lines and the numerical paths appear as solid lines that are offset by one generation. As can be seen from Figure 5, the AB technique picks the same equilibrium as the numerical simulation in all cases and the AB simulations all seem to arrive at fixation in approximately the same number of generations as their numerical counterparts. This is in fact statistically true (with $ks = 0.0836$, $p = 0.99$ for the largest difference between the two initial state $(a_{1/3}, a_{1/2}, a_{2/3}) = (0.8, 0.1, 0.1)$) with one exception which is obvious from Figure 5. From the starting state (0.5, 0.267, 0.333) the AB model arrives at the equilibrium in 27 generations while the numerical simulation takes more than 50 generations. This difference is highly significant ($ks = 0.69$, $p = 0$). Recall that this starting state is the unstable internal Nash equilibria we

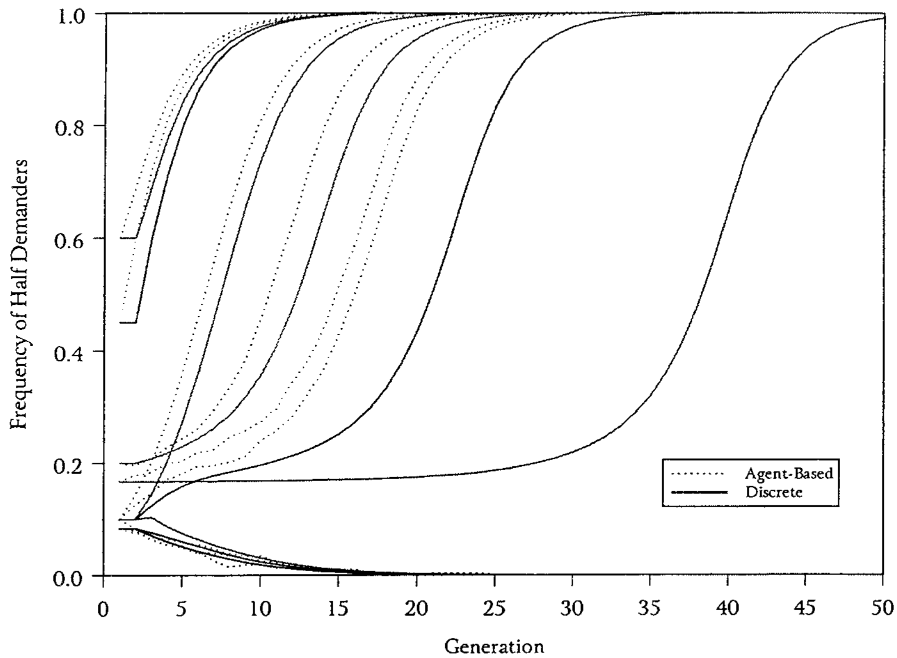

Figure 5. Comparison of agent-based and numerical simulations – generic model.

spoke of in Section 2.1, which means that the contours around the internal saddle point are steeper in the AB simulation than in the numerical. However, in general according to our two criteria, the two techniques are more or less indistinguishable.

As for the systemic model, two levels of experimentation, d where simulated using each method. The results are presented as Figures 6 and 7. As stated above (Section 2.2), the systemic model is interesting because it allows us to compare the two simulation methods when we add gravity towards the center of the strategy simplex. For the larger level of experimentation, $d = 0.05$ (Figure 6) we see that the results of the two methods are quite different. First, note that while in all cases both methods choose the 'mostly demanding 1/2' equilibrium, the AB method stubbornly fixes at a frequency of 0.97 half-demanders while the numerical method chooses the fixed point of $a_{1/2} = 0.85$[11]. Obviously, all comparisons are highly significantly different due mostly to the difference in the frequency of half-demanders selected. Also of note is the difference in time to convergence when half-demanders start out as a low fraction of the population. The contours are now much less steep in the AB simulations and hence the process converges much less quickly than the numerical simulations.

When the level of experimentation is reduced to $d = 0.01$ in Figure 7 we again find that both models pick the same equilibrium in the most general sense, but the AB simulations over-emphasize the draw of the equal split as was also seen in the $d = 0.05$ case. However the difference is not nearly as drastic here (compare 0.99

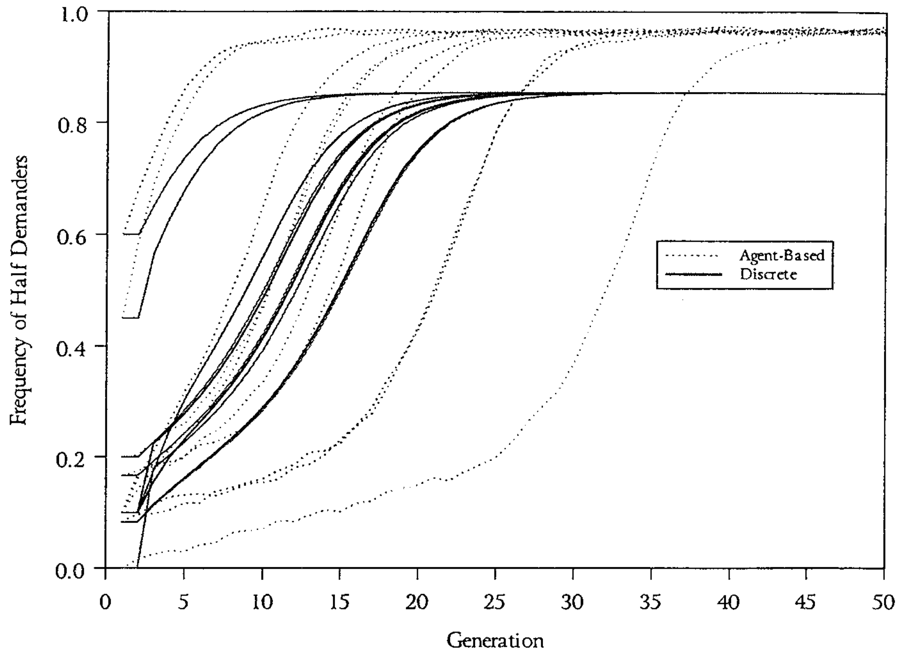

Figure 6. Comparison of agent based and numerical simulations – systemic model ($d = 0.05$).

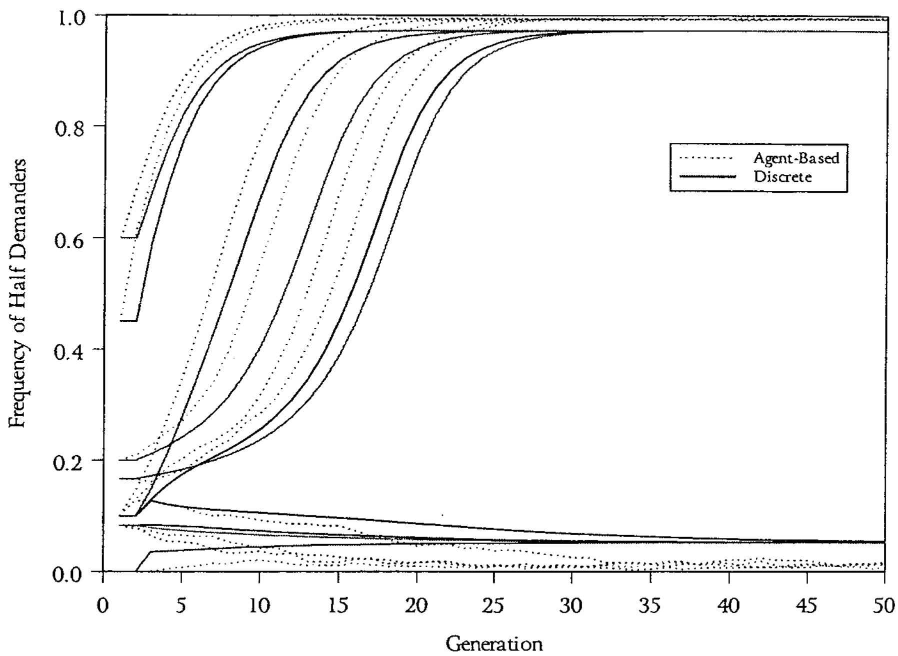

Figure 7. Comparison of agent-based and numerical simulations – systemic model ($d = 0.01$).

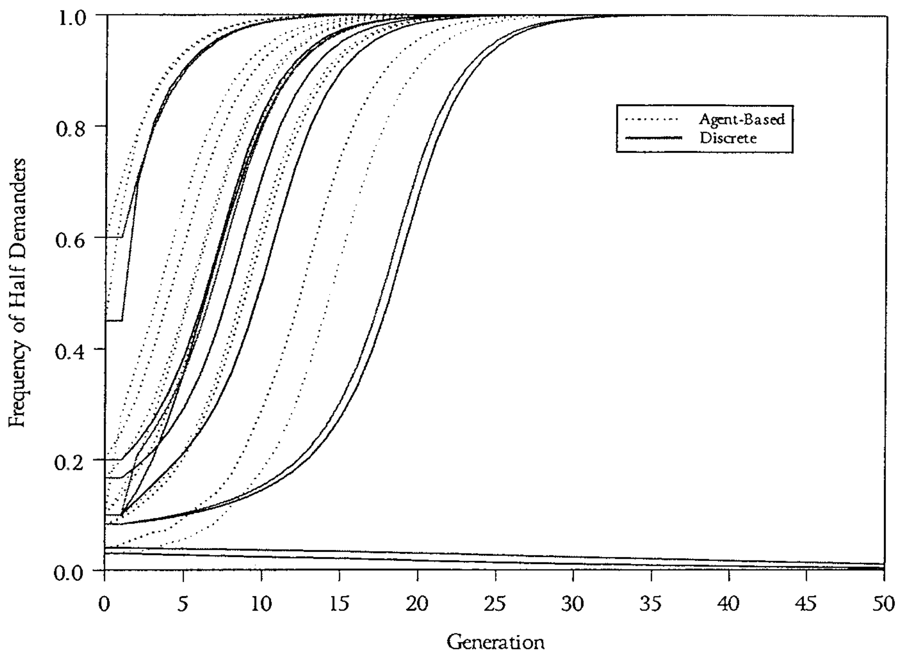

Figure 8. Comparison of agent-based and numerical simulations – assortative model ($z = 0.25$).

to 0.97). Also, the convergence times are much more similar from a common starting state than was the case with $d = 0.05$[12]. Overall, the simulations of the systemic model demonstrate that the AB approach seems to do an adequate job of selecting equilibria, but the final equilibrium state tends to differ between the two methods. It appears, too, that as the draw towards the middle of the simplex increases (i.e. $d \uparrow$), the AB method performs less well in identifying the proper population distribution. We can conjecture as to why the two methods yield different results. Obviously, the AB model underestimates the theoretical draw to the middle of the strategy simplex. One reason for this might have to do with the difference between the finite population realization of the model and the numerical results which are based on an infinite population. We return to this point in Section 4.3.

4.2. CONVERGENCE IN THE ASSORTATIVE AND SOPHISTICATED MODELS

With the assortative model we can test how the simulation methods react when we change the pairing rule of the generic model. Also, using the sophisticated model we can explore what happens as we expand the strategy set and add weakly dominated strategies. We begin by examining what happens when interactions become assortative. Figures 8 and 9 report simulation results for two different levels of assortation that have been chosen to straddle the critical value reported in Proposition 4. When 25% of interactions are between like types (Figure 8), we

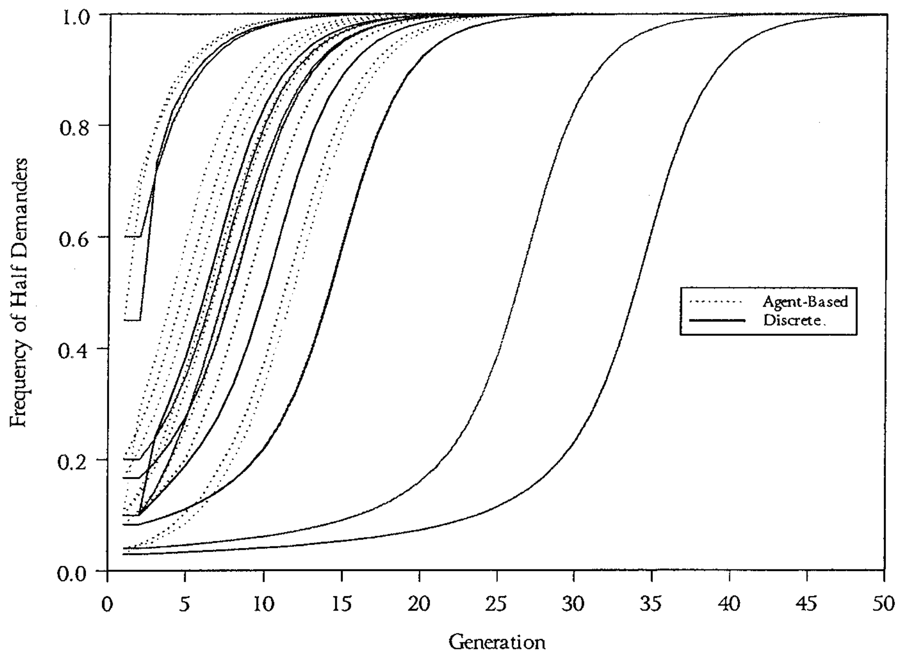

Figure 9. Comparison of agent-based and numerical simulations – assortative model ($z = 0.35$).

find two major results. First, the AB method tends to converge quicker than the numerical simulations and second, the AB model incorrectly chooses the all half-demander equilibrium for very low values of initial half-demanders[13]. When the level of assortation rises above the critical value of 0.33 to 0.35 (Figure 9) then the two methods pick the same equilibrium in each case. However, the trend of the AB simulation converging quicker is maintained as the level of assortation rises. This indicates that the AB method simulates much steeper contours than the numerical method in the assortative model. This reverses the result of the systemic model in which $d = 0.01$ and the convergence times were much quicker in the numerical simulations.

As predicted by Proposition 5, which states that any distribution composed of more than fifty percent half demanders and the rest responsive agents is an equilibrium, there are many more stable equilibria in the sophisticated model. Five equilibria are shown in Figure 10, which plots the evolution of half-demanders from various starting states. Notice first that there appears to be a critical level of half-demanders around 0.20 such that when there are more than 20 percent half demanders in the initial populations, they end up sharing an equilibrium with responsive agents, but for less than 20 percent the greedies quickly dominate the population and there is 100 percent conflict.

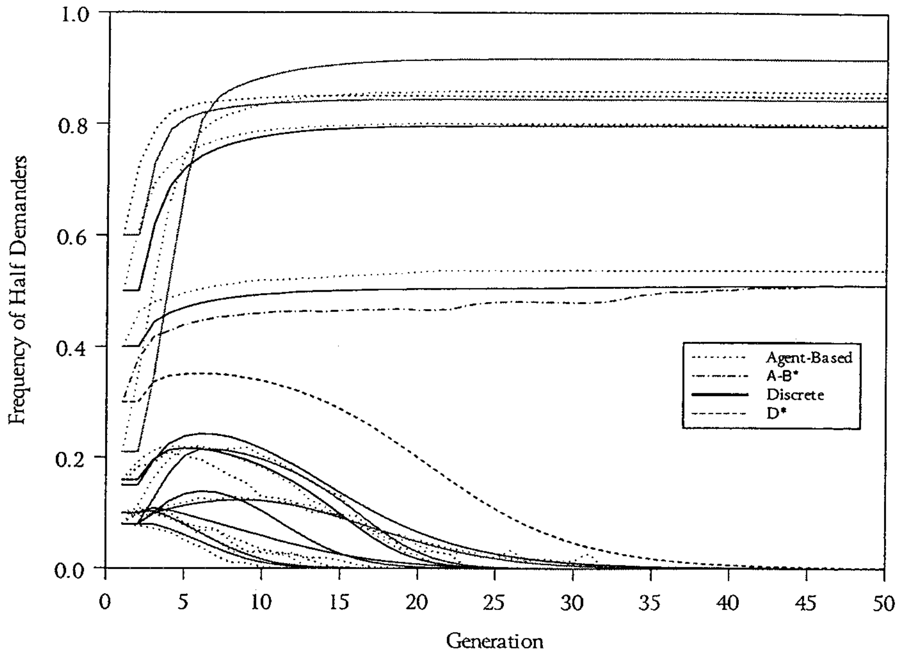

Figure 10. Comparison of agent-based and numerical simulations – responsive model.

As for the performance of the AB method, we see that the AB simulations are actually very good at selecting the same equilibria as the numerical simulations in many cases (e.g. starting with $a_{1/2} = 0.6, 0.5$ or below 0.2). This is interesting because in the responsive model equilibrium selection is much more difficult than in the other two models where the phase space is split between only two equilibria. This demonstrates that the AB method does not summarily do away with weakly dominated strategies. However, there are other cases (e.g. $a_{1/2} = 0.21, 0.3$) where the AB method performs rather poorly. Figure 10 highlights the case of the initial state with 30 percent half-demanders in which the AB method picks the wrong equilibrium entirely (compare line A–B* to D*). As far as our other criteria, time to convergence is concerned, the AB model does quite well when it selects the correct equilibrium[14].

4.3. DISCUSSION

Overall, the agent-based method of simulating complex evolutionary models seems to match rather well the behavior seen in the numerical simulations. However, had we stopped after reporting the results in Figure 5 for the basic discrete model we might have prematurely recommended that agent-based methods accurately approximate numerical simulations thinking they also accurately predict equilibria in more complex evolutionary models. Therefore, exploring our three variants of the basic Nash demand game is an important undertaking because we have been able to stress-test the AB method by digging four pitfalls. Using the systemic model we have shown that the AB approach does quite well in selecting fixed points in a very general sense, but given enough gravity towards the center of the strategy simplex, the accuracy of the AB's final population distribution prediction deviates substantially – it under-predicts the force of this type of gravity. Secondly, by changing the rule that pairs agents, we see that assortation can potentially disrupt the power of the AB method to select equilibrium altogether. Finally, by adding strategies and thereby making some demands weakly dominated, the sophisticated model lays another trap. Despite this, the AB method tracked the numerical simulations very well in that it selected the appropriate equilibrium from many possibilities and accurately predicted the final population distribution.

What might account for the deviations we have seen between the AB simulations and the numerical simulations? For the specific case of the assortative interaction model one reason why the AB simulations are likely to have converged quicker and picked the wrong equilibrium in the low assortation case is because of the structural decision that had to be made concerning how to handle like-type interactions and shuffling given an array of agents (see Section 3.2). On a more general level another strong candidate mentioned in the introduction is the fact that the AB method must work with a finite population while the numerical method can maintain the assumption of infinitely many agents. To examine this hypothesis we can further explore the scenario in which the AB and numerical simulations differ the most – the systemic model where $d = 0.05$. To see whether the size of the population of bargainers matters we can conduct a simple comparative static experiment. We can fix a starting distribution and then evaluate the effect of changing the number of agents in the AB simulation. Figure 11 reports the results of this experiment.

As one can see the most prominent effect of increasing the number of agents is to smooth out the time paths in the AB simulations. This is true in both Frame A, where the starting state is $a_{1/2} = 0.45$, and in Frame B, where the starting state is $a_{1/2} = 0.167$ (the internal unstable Nash equilibrium). Actually, rather than reducing the predicted equilibrium frequency of half-demanders as the population size increases, the AB method does best at predicting the fixed point by using between 500 and 2000 agents. This claim is justified by noticing that the equilibria selected using 500 or 2000 agents is closer to the numerical simulation than are the other cases including the simulations using 1 million bargainers[15].

Axtell *et al.* (1996) discuss what they term *docking*: whether or not two computational models can produce the same result. The current paper is work in this same vein but differs in two ways. First this paper looks as evolutionary bargaining models in particular and second, the work detailed above seeks to answer a slightly different question. Rather than asking whether one computational model can be subsumed by another we are interested in whether the agent-based approach is something to be pursued for more complicated models of bargaining. The answer

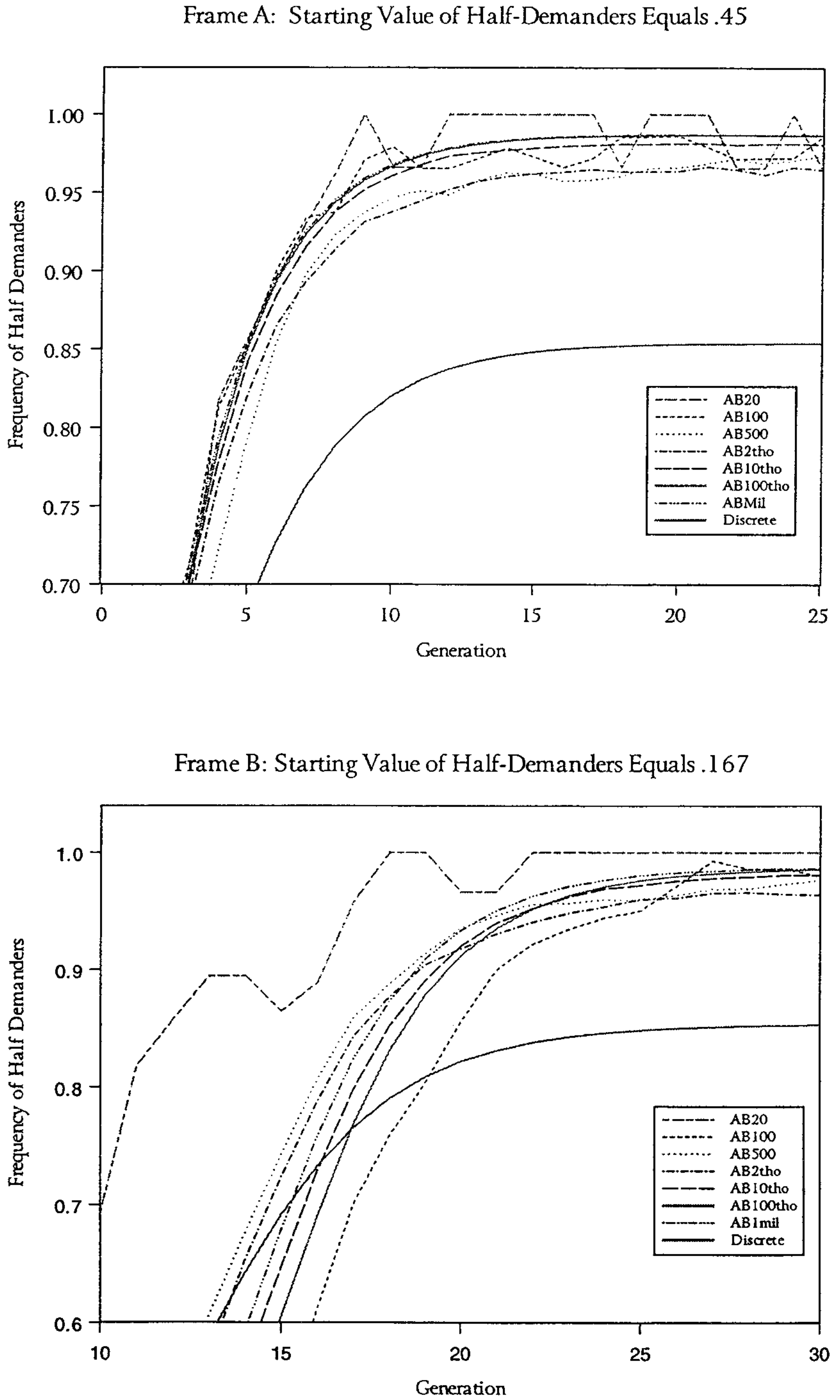

Figure 11. The effect of poluation size in the systemic model ($d = 0.05$).

to this question is a tentative yes and the Axtell *et al*. (1996) framework can help understand why. One of the criteria for determining whether two models can be docked is called *distributional equivalence*. Basically this means that, controlling for starting positions, models are distributionally equivalent if their results can be shown to not differ statistically. Using this criteria the agent-based method performed quite well, but in the end it fails. However, a more loose criteria is *relational equivalence* which means that while the results are not dead-on the same, the two methods largely produce similar results. With reference to the results of this study this can be taken to imply that the two approaches select the same equilibria in a very general sense. In this case, the agent-based method for simulating evolutionary bargaining models does pass the test. As a result we can say that the agent-based method for computing equilibria is substantially equivalent to the analytical/numerical approach.

Acknowledgements

The author wishes to thank Rob Axtell (in particular) for pointing out errors in an earlier version of the paper and for making suggestions that have substantially improved the paper. The author also thanks Rob Boyd for suggestions on the numerical simulations, and an anonymous referee for a critical reading of an earlier draft.

Appendix

Proofs of Propositions 1–5

Proof of Proposition 1. The first condition in the definition of evolutionary stability assures that an ESS is also a Nash equilibrium of the underlying game. This allows us to restrict attention on the Nash equilibria of $\mathcal{D}$ as candidates for evolutionary stability. Suppose there is a Nash equilibrium (e, e) where e is fully supported and e is also an ESS. By construction e must be a strategy that equalizes the payoffs to each pure strategy. Therefore we know that the first condition for an ESS holds as an equality when checked against each pure strategy. This means we can focus on the second or viability condition. The highest payoff of a pure strategy played against itself results when two 1/2 demanders meet. Therefore if e is an ESS, it must be the case that $\pi(e, 1/2) > 1/2$. The payoff of e against the pure strategy 1/2 is $p_e/3 + q_e/2$ where p_e and q_e are the probabilities of demanding 1/3 and 1/2, respectively, when using e. The payoff $\pi(e, 1/2)$ reaches a maximum of 1/2 when q_e equals 1. However this leads to a contradiction because e is fully supported and therefore $p_e + q_e < 1$. Hence, no fully supported strategy can be an ESS of $\mathcal{D}$. □

Proof of Proposition 2. To prove Proposition 2 we must only show that the eigenvalues calculated at the asymmetric and symmetric equilibria have negative real parts. At the asymmetric equilibrium the eigenvalues are –0.22 and –0.03 and at the symmetric equilibrium the eigenvalues are –1/2 and –1/6. □

Proof of Proposition 3. Because closed-form solutions of (6) and (7) are nearly impossible to calculate, analytically we allow the numerical simulation to calculate them for us and then show that such a population distribution is stable. The simulations select an equilibrium distribution of (0.03, 0.96, 0.01) when $d = 0.013$. In this case, the eigenvalues are –0.47 and –0.16 for the selected equilibrium. Mathematica shows the critical value of d where the eigenvalue becomes zero is 0.01296. □

Proof of Proposition 4. There are four closed form solutions to the system (8)–(9) and two non-feasible solutions. We disregard the fixed point (0,0). The three other fixed points are (1/2,0), (1,0) and (0,1). Construct the Jacobian matrix for the system (8)–(9). At the fixed point (1,0) the Jacobian has strictly positive eigenvalues and therefore is a source for any value of z. At the point (1/2,0) the eigenvalues of the Jacobian are $\{-1/6, (3z - 1)/12\}$. The second value is positive for $z > 1/3$. Finally, the eigenvalues of the Jacobian at the point (0,1) are $\{5/6, -(1 + z)/2\}$ which are both negative for any $z > 0$. □

Proof of Proposition 5. (Adapted from Proposition 1 Ellingsen, 1997). First define $\{A|A_{1/2} \geq 1/2 \text{ and } a_r = 1 - a_{1/2}\}$. We first need to show that any population $a \in A$ is a NSS. Any a is robust to invasion by any demand $x < 1/2$ because $\pi(a, a) = 1/2 > a_{1/2}x + (1 - a_{1/2})x = x = \pi(x, a)$. Any distribution a is also robust to invasion by immodest strategies $y > 1/2$ because $\pi(a, a) = 1/2 > (1 - a_{1/2}) > (1 - a_{1/2})y = \pi(y, a)$. Finally, we see that $\pi(a, a) = 1/2 = \pi(v, a)$ for $v \in \{1/2, r\}$. Hence, any $a \in A$ is a Nash equilibrium and therefore a NSS, but the last equality prevents a from being an ESS.

Now the 'only if' part. First, it is obvious that if $a_r > 1/2$, then the greedy strategy $(x - 1)$ can invade. However, if $a_r = 1/2 = a_{1/2}$, then being greedy does just as well as the current population, and therefore a is still a Nash equilibrium, but not a NSS. To see this consider invasion by a small population consisting of both greedy and responsive agents. In this case the fraction of responsive agents will be pushed above one-half and therefore greedys earn an average payoff slightly larger than 1/2. We can eliminate invasion by populations of agents demanding less than half because any mix of these strategies does strictly worse than demanding 1/2. Also Proposition 1 has already shown that fully supported strategies are not evolutionarily stable under D and this result can be generalized to D'. The last possibility to check is that no mixture of responsive and 2/3 can invade. It is clear the greedy agents would dislodge such a population. □

Notes

[1] Axtell, Epstein and Young (2000), for example.

[2] A bargaining convention in this context can be defined as one of multiple stable equilibria of the underlying game (Sugden, 1986).

[3] Adding the two strategies (2/5, 3/5) or (1/5, 4/5) and hence increasing the strategy space only adds equilibria along the diagonal. Numerical simulations show that this does not add stable fully supported equilibria nor does it remove previously stable equilibria.

[4] A strategy is said to be *neutrally stable (NSS)* if the strategy is incumbent and no invading strategy is more fit (Weibull, 1995).

[5] The point of calculating both the ESSs if the finite Nash demand game and the stable fixed points of the corresponding replicator dynamics is that the two sets do not always coincide when the dynamics are nonlinear. Friedman (1991) for example of games in which the two concepts do not coincide.

[6] For rates of experimentation between 0.013 and 0 the asymmetric Nash equilibrium remains stable. This fact was pointed out by Rob Axtell.

[7] This model is adapted from Skyrms (1994).

[8] This model is a simplification of Ellingsen (1997).

[9] All the simulations discussed below were programmed in Visual Basic for an IBM PC. Both the code and executable programs are available from the author upon request.

[10] This method is known as the Euler algorithm for integrating ordinary differential equations.

[11] The AB method fixes at 0.97 rather than 0.95 because, while 5% of the agents are always experimenting, a third of them choose to demand half. Plugging the values for $a_{1/3} - 0.105$ and $a_{1/2} = 0.854$ into Equation (6) results in a value equal to 0.0000422 which, considering rounding, defines and equilibrium of the perturbed replicator dynamic.

[12] For example, even given the difference in end states, the largest difference between the two methods, starting from $a_{1/2} = 0.20$, is insignificant ($ks = 0.1905$, $p = 0.74$).

[13] The Kolmogorov-Smirnov test indicates significant differences (at the 5% level) in convergence times in 8 of the 11 cases.

[14] In 6 of the 9 case where the AB procedure picked the correct ending distribution the time paths were indistinguishable at the 5% level.

[15] The Kolmogorov-Smirnov statistics testing the difference between the 500 agent simulation and the 1 million is Frame A: $ks = 0.87$, $p = 0$; Frame B: $ks = 0.76$, $p = 0$ and the statistic testing the difference between the 2000 and 1 million is Frame A: $ks = 0.90$, $p = 0$; Frame B: $ks = 0.78$, $p = 0$.

References

Axtell, R., Axelrod, R., Epstein, J. and Cohen, M. (1996). Aligning simulation models: A case study and results. *Computational and Mathematical Organization Theory* **1**, 123–141.

Axtell, R., Epstein, J. and Young, P. (2000). The emergence of economic classes in an agent-based bargaining model. In: S. Durlauf and H.P. Young (eds.), *Social Dynamics*. MIT Press.

Ellingsen, T. (1997). The evolution of bargaining behavior. *Quarterly Journal of Economics*, 581–602.

Friedman, D. (1991). Evolutionary games in economics. *Econometrica* **59**, 637–666.

Mailath, G. (1998). Do people play Nash equilibrium? Lessons from evolutionary game theory. *Journal of Economic Literature* **36**, 1347–1374.

Maynard Smith, J. (1982). *Evolution and the Theory of Games*. Cambridge University Press: Cambridge.

Maynard Smith, J. and Price, G. (1973). The logic of animal conflict. *Nature* **246**, 15–18.

Nash, J. (1953). Two-person cooperative games. *Econometrica* **21**, 128–140.

Peyton, Young, H. (1993). An evolutionary model of bargaining. *Journal of Economic Theory* **59**, 145–168.

Rubinstein, Ariel. (1982). Perfect equilibrium in a bargaining model. *Econometrica* **50**, 97–109.

Samuelson, J. (1997). *Evolutionary Games and Equilibrium Selection*. MIT Press: Cambridge.

Samuelson, L. and Zhang, J. (1992). Evolutionary stability in asymmetric games. *Journal of Economic Theory* **57**, 363–391.

Skyrms, B. (1994). Sex and justice. *Journal of Philosophy* **91**, 305–320.

Stahl, I. (1972). *Bargaining Theory*. Economic Research Institute: Stockholm.

Sugden, R. (1986). *Economics of Rights, Cooperation and Welfare*. Basil Blackell: Oxford.

Weibull, J. (1995). *Evolutionary Game Theory*. MIT Press: Cambridge.