MODEL ROBUSTNESS VERSUS PARAMETER EVOLUTION: ASSORTATIVE INTERACTION WITHIN A BARGAINING GAME

MARK H. GOADRICH, Computer Sciences Department, University of Wisconsin-Madison*

ABSTRACT

Agent-based models exploring aspects of social behavior invariably contain multiple parameters such as population size, heterogeneous makeup and spatial distribution. One common way to validate a model is to ensure robustness, that is, the model must produce consistent results independent of the initial parameter settings. However, when information can be learned about the prior probability of some parameter settings, I believe our robustness requirements on these parameters should be relaxed. We should instead focus on the results produced from using these more likely settings. Brian Skyrms investigates a two-player non-cooperative one-shot bargaining game called "Divide-The-Cake." Placed in an evolutionary setting, where players' claims are genetically hardwired and pairings are made at random, only 67% of initial population distributions result in all players using the "fair" strategy. Skyrms introduces correlation among players and shows this precipitates the evolution of fairness from 100% of initial populations, but critics argue that his exploration of correlation is lacking; other correlation models lead to much worse performance. In this paper, I examine the evolution of these non-random correlations, known as assortative interactions, through two separate agent-based models, a social network and a Schelling segregation model. These experiments show convergence to the "fair" strategy approximately 90% of the time. I believe evolving the assortative interactions between players to find likely correlations, as opposed to guaranteeing model robustness, leads to a much more realistic picture of a model's behavior.

Keywords: model robustness, assortative interaction, social networks, evolutionary game theory, agent-based models.

INTRODUCTION

One of the common dimensions used to classify agent-based models is the degree of complexity. Models can be abstract, such as an iterated prisoner's dilemma (Axelrod, 1984), or realistic, attempting to "investigate where prehistoric people of the American Southwest would have situated their households based on both the natural and social environments in which they lived" (Village, 2003). Abstract models usually have broad applicability and are pursued to explain the general mechanisms underlying a particular process, but these simple models are criticized for not capturing the complex details of the real world.

However, as we move towards realistic models, the size and scope of what is being simulated explodes. Our agents might now have to cope with heterogeneous thresholds and diverse landscapes, among other complications. Each new aspect brings into the simulation new parameters which must be tested, as these models are open to being overly sensitive to any one choice of parameters. Ultimately, we wish our models to be "robust," and produce consistent

* Mark H. Goadrich, University of Wisconsin - Madison 6785 Medical Sciences 1300 University Ave, Madison, WI 53706; e-mail: richm@cs.wisc.edu.

behavior independent of our parameter settings, but how reasonable is this goal? Testing for robustness implies that all parameter settings are equally likely, yet this is not always the case. What if we knew prior information about the parameter likelihood, a situation which brings into question the strict pass or fail test for robustness?

To demonstrate this situation, I will explore a simple bargaining game made popular by Brian Skyrms. In his book "Evolution of the Social Contract," Skyrms explores the use of evolutionary game theory to explain our concepts of fairness (Skyrms 1996). As in other fields, Skyrms hopes this direction will provide explanations for human social behavior when theories relying on rational deliberation are lacking. His initial abstract model quickly becomes complicated when he introduces correlation among his agents, and is no longer robust when taking into account these new parameters. The next two sections provide a brief summary of the current literature on this topic.

Divide-The-Cake

Skyrms' first example involves dividing a chocolate cake between two players, C_1 and C_2. Each player demands a certain amount of cake; when the total cake demand is less than or equal to the whole cake, each player receives her demand. However, should the total demand exceed one, the cake is discarded and the players leave empty-handed. Our natural inclination when presented with this game is to divide the cake evenly, 1/2 for C_1, 1/2 for C_2. But why do we consider this split fair? Skyrms points out that there are an infinite number of polymorphic solutions, or Nash equilibriums (e.g. C_1 demands 30% and C_2 demands 70%). Rational deliberation does not help us distinguish between the "fair" solution and the polymorphic splits. This opens the door to other explanations of fairness, namely, that evolution may have a hand in deciding our social behavior.

An evolutionary model is constructed by creating a finite population of players, each with preset and constant cake demands. I assume the use of D'Arms *et al.* (1998) finite population and discrete simulation rather than Skyrms' continuous equations. Later, we will explore larger numbers of strategies, but to simplify the analysis we start with only three:

S_1: Always demand 1/3 of the cake (*Modest*)

S_2: Always demand 1/2 of the cake (*Fair*)

S_3: Always demand 2/3 of the cake (*Greedy*)

Individual cake games are conducted by independently and uniformly drawing C_1 and C_2 without replacement until all players are exhausted. A player's fitness score is the portion of the cake, if any, received in a game. The next generation of players is determined by the relative success or failure of each strategy for this game in combination with the current population distribution, a selection process known as the *replicator dynamics* (Weibull, 1995). This iterative process is continued until convergence of the population to a steady distribution. Skyrms states that the percent of initial population distributions which evolve to a population where all players demand 1/2 the cake is 74%, not exactly the degree of success that we might expect.[1]

[1] Skyrms' result of 62% as presented in his book is calculated for a population with 10 possible cake divisions, which are explored later. In my simulations of the three above divisions I found this number to be 74%.

TABLE 1 Skyrms' Positive Correlation Matrix

Strategy *i*	pref(*i*, Modest)	pref(*i*, Fair)	pref(*i*, Greedy)
Modest	0.8	0.1	0.1
Fair	0.1	0.8	0.1
Greedy	0.1	0.1	0.8

Skyrms solves the problem by introducing positive correlation among the strategies, or non-random mating of like-minded players. His players are given the ability to determine self versus non-self relationships among opposing players: greedy most likely plays with greedy, fair with fair, and modest with modest. This assumption breaks the polymorphic barrier, and Skyrms reports that only minimal correlation is necessary to cause widespread outbreaks of fairness quickly reaching 100%.

We can incorporate this correlation into our model by allowing the first player to influence the choice of their opponent. The initial finite-size population is still created according to a random population distribution, and C_1 is selected randomly from the current population distribution, $P(S_i)$. A player's preference for other strategies can be as defined by a function *pref*(*i*,*j*), the preference of a player using strategy *i* for a player using strategy *j*. Table 1 shows the correlation matrix when using Skyrms' assumptions of non-random mating. The selection of C_2 is governed by the following formula:

$$P(C_2 = S_j \mid C_1 = S_i) = \alpha < pref\,(S_i, S_j) * P(S_j) >$$

where α is the normalization constant. If the total demand of C_1 plus C_2 is less than 1, a record is made of a successful game for each player's strategy. This process continues to sample without replacement until all players are exhausted. The average fitness for a strategy is calculated based on successful games, and the players are redistributed accordingly for the next round. To assist in evaluating correlations, we will define the *strength* of a correlation matrix in terms of the scale between preferred and non-preferred strategies. For example, Table 1 is of strength 8, since fair is eight times more likely to choose fair over either greedy or modest.

Anti-Correlation Rebuttal

Skyrms' positive correlation is part of a broad class of correlations known as assortative interactions. Assortative interaction is usually discussed in the context of choosing a mate for reproduction as opposed to random mating strategies; in general it describes the tendency for individuals to choose their associates. In "Divide-the-Cake", C_1 is still randomly selected, but the selection of C_2 is now weighted by the preferences of C_1.

TABLE 2 Anti-Correlation Matrix from D'Arms *et al.*

Strategy *i*	**pref(*i*, Modest)**	**pref(*i*, Fair)**	**pref(*i*, Greedy)**
Modest	0.33	0.33	0.33
Fair	0.1	0.8	0.1
Greedy	0.47	0.47	0.06

Justin D'Arms (1996) quickly replies with questions about the assumption of positive correlation. He proposes that a model is robust if the result is virtually independent of the starting parameters. Skyrms' positive correlation makes the model robust with respect to initial population distributions, but correlation is now a parameter and should be examined with the same scrutiny. Finding one particular correlation that works is not a very robust argument.

D'Arms *et al*. (1998) expand this claim into a model allowing for both correlation and anti-correlation as shown in Table 2. A Greedy player using anti-correlation should wish to face anyone but another Greedy in competition for cake. Fair will still use positive correlation and prefer Fair players, and Modest will be happy playing against all three strategies. Unfortunately, this anti-correlation enlarges the basin of attraction for a Greedy/Modest polymorphism to 54%, and their results hold across many strengths. D'Arms *et al*. conclude Skyrms' model is not robust with respect to variations in correlations.

MODELS OF CORRELATION

Both Skyrms and D'Arms *et al.* are using a scatter-shot approach to find reasonable correlation assumptions. While D'Arms *et al.* succeed in their goal of providing a couter-example to Skyrms, this should not be the end of the discussion. What other models of correlation are possible, how do they influence the evolution of fairness, and more importantly, are some more likely than others?

In a separate critique of Skyrms' model, Martin Barrett *et al.* (1999) describe what they believe is the most natural correlation matrix (shown in Table 3) where players choose associates with a mind toward their own utility. Modest players still freely associate with all players equally, but fair players prefer fair and modest opponents, while greedy players exclusively prefer modest opponents.[2] In general, players seek out opponents who will not tip their combined demand over one. Using the preference matrix from Table 3, my experiments show that 90% of initial populations will evolve to all fair. If fair players constitute at least 8% of the initial population, this evolution is guaranteed. As the strength of this correlation increases, fairness approaches 100%.

[2] Barrett et. al. suspect the resulting fairness model will evolve similarly to D'Arms et. al. with a broad basin for polymorphism.

TABLE 3 Utility Preference Correlation Matrix

Strategy *i*	**pref(*i*, Modest)**	**pref(*i*, Fair)**	**pref(*i*, Greedy)**
Modest	0.33	0.33	0.33
Fair	0.47	0.47	0.06
Greedy	0.8	0.1	0.1

TABLE 4 Efficiency Preference Correlation Matrix

Strategy *i*	**pref(*i*, Modest)**	**pref(*i*, Fair)**	**pref(*i*, Greedy)**
Modest	0.1	0.1	0.8
Fair	0.1	0.8	0.1
Greedy	0.8	0.1	0.1

Why is there such a benefit for fair players? Since fair is content with either fair or modest opponents, it will steal away some of the necessary modest players from greedy. The greedy strategy is never able to act with full power, and therefore is at an evolutionary disadvantage. This is opposite of D'Arms' anti-correlation, where greedy players were stealing away fair players and disrupting the average utility of fair. But for realistic anti-correlation, some greedy players must be willing to sacrifice themselves for the good of the strategy. This is an unlikely option considering they are demanding 2/3 of the cake.

Another possible correlation is created by having players search for opponents seeking an complementary proportion of the cake as them, a correlation suggested by Ernst (2001). Table 4 shows this "efficiency" correlation at strength 8. Ernst considers competition between groups of players, as opposed to a single population, and finds that efficient populations will fare better than those which leave cake behind. How this could arise within a population is not exactly clear, since there is currently no benefit to consuming all of the offered cake, nevertheless this is a possible correlation, and it exhibits behavior similar to that of utility preference.

Table 5 lists each correlation matrix discussed and shows the effect on fairness evolving as the strength of correlation increases. With all correlations except anti-correlation, greater strength brings about a greater evolution of fairness. While neither is as successful as positive correlation, utility preference is the closest. We can see that the use of a utility preference correlation matrix would very beneficial to the evolution of fairness.

TABLE 5 Effects of Various Assortative Interaction on Evolution of Fairness

Strength	Positive	Anti-Correlation	Utility	Efficiency
0	74%	74%	74%	74%
2	98%	63%	77%	67%
4	100%	59%	83%	70%
8	100%	56%	90%	79%
16	100%	56%	95%	87%

Instead of only relying on robustness as our criterion for success, it is also important to discriminate between correlations to find those which could arise naturally from player interactions. To properly understand the relationship between assortative interactions and the evolution of fairness, we must consider the evolution of correlations. The rest of this paper explores two models for discovering such correlations. First, players consciously construct a social network to help them learn what types of players will benefit their own claim. Second, players unconsciously employ a Schelling segregation model on a 2D-lattice; players randomly select new locations, without looking at their opponents strategies, when their current utility falls below a threshold.

SOCIAL NETWORK MODEL

Learning preferences amongst players is not as hard as it may seem; in fact, nature provides us evidence that these interactions do indeed exist. Elliott Sober and David Sloan Wilson cite an experiment on the interactions of guppies in the context of altruism.

> A separate experiment allowed three guppies to inspect predators in an aquarium divided into three lanes by transparent panels. The guppies were then placed in an apparatus that allowed the fish that occupied the center lane to indicate a preference for one of the two side fish by swimming over to join it as a companion. The side fish that went closer to the predator was consistently chosen as a future associate (Sober and Wilson 1998, pg. 140).

If simple minded guppies can learn preferences that will increase their utility, Sober and Wilson contend, how much more likely is it that humans with all our faculties in fact do the same. A simple way to learn the preferences for our agents is to randomly pair them with opponents and then record whether a game is successful or not. We assume each strategy is assigned a tag, which can be recognized by other players, and records are kept based on these tags and not individual players. The resulting correlation matrix is consistent with the utility preference matrix from Table 3. But this seems too easy.

TABLE 6 Sample results of Dynamic Social Network Model: resulting fairness 99%

Strategy *i*	**pref(*i,0.1)***	**pref(*i,0.2)***	**pref(*i,0.3)***	**pref(*i,0.4)***	**pref(*i,0.5)***	**pref(*i,0.6)***	**pref(*i,0.7)***	**pref(*i,0.8)***	**pref(*i,0.9)***
0.1	0.08	0.17	0.02	0.12	0.6	0.21	0.14	0.13	0.07
0.2	0.12	0.00	0.07	0.07	0.15	0.24	0.08	0.27	0.00
0.3	0.06	0.03	0.02	0.44	0.41	0.03	0.01	0.01	0.00
0.4	0.55	0.07	0.11	0.07	0.13	0.04	0.00	0.01	0.01
0.5	0.19	0.16	0.18	0.25	0.21	0.00	0.00	0.01	0.01
0.6	0.39	0.03	0.41	0.15	0.00	0.00	0.01	0.00	0.01
0.7	0.01	0.84	0.13	0.00	0.01	0.01	0.00	0.00	0.00
0.8	0.58	0.40	0.00	0.00	0.00	0.00	0.01	0.00	0.00
0.9	0.94	0.00	0.01	0.00	0.01	0.01	0.00	0.01	0.00

TABLE 7 Sample results of Dynamic Social Network Model: resulting fairness 58%

Strategy *i*	**pref(*i,0.1)***	**pref(*i,0.2)***	**pref(*i,0.3)***	**pref(*i,0.4)***	**pref(*i,0.5)***	**pref(*i,0.6)***	**pref(*i,0.7)***	**pref(*i,0.8)***	**pref(*i,0.9)***
0.1	0.08	0.03	0.25	0.03	0.15	0.16	0.01	0.18	0.11
0.2	0.02	0.01	0.06	0.44	0.32	0.06	0.08	0.00	0.00
0.3	0.02	0.21	0.01	0.23	0.04	0.10	0.40	0.00	0.00
0.4	0.17	0.25	0.10	0.02	0.05	0.41	0.01	0.00	0.00
0.5	0.23	0.43	0.22	0.09	0.01	0.01	0.00	0.00	0.00
0.6	0.06	0.15	0.62	0.14	0.00	0.00	0.00	0.00	0.01
0.7	0.25	0.07	0.66	0.01	0.00	0.00	0.00	0.00	0.00
0.8	0.63	0.34	0.00	0.00	0.00	0.00	0.01	0.01	0.00
0.9	0.97	0.00	0.00	0.00	0.00	0.01	0.01	0.00	0.00

Skyrms and Pemantle (2000) suggest a more complex mechanism to dynamically learn a social network between game players. To make things more interesting, the number of strategies is now 9, from 0.1 to 0.9. We redefine “greedy” and “modest” players as those demanding more or less than 1/2 the cake, respectively. Players begin with a uniform preference for all other players. Each player is given 1000 rounds to play “Divide-the-Cake,” choosing opponents according to their preference vector. Whenever a game is successful, the player initiating the visit updates their preference vector to increase the chance of revisiting this cooperative strategy. To add noise, unsuccessful games are recorded favorably 20% of the time.

We might expect this model to evolve similarly to the simple model, yet there are key differences. Two resulting correlation matrices can be seen in Tables 6 and 7. While all of the preference is concentrated on opponents that will provide a positive outcome, this preference is no longer uniform; random choices of initial opponents cause the revisiting and reinforcement of certain players instead of other equally acceptable players.

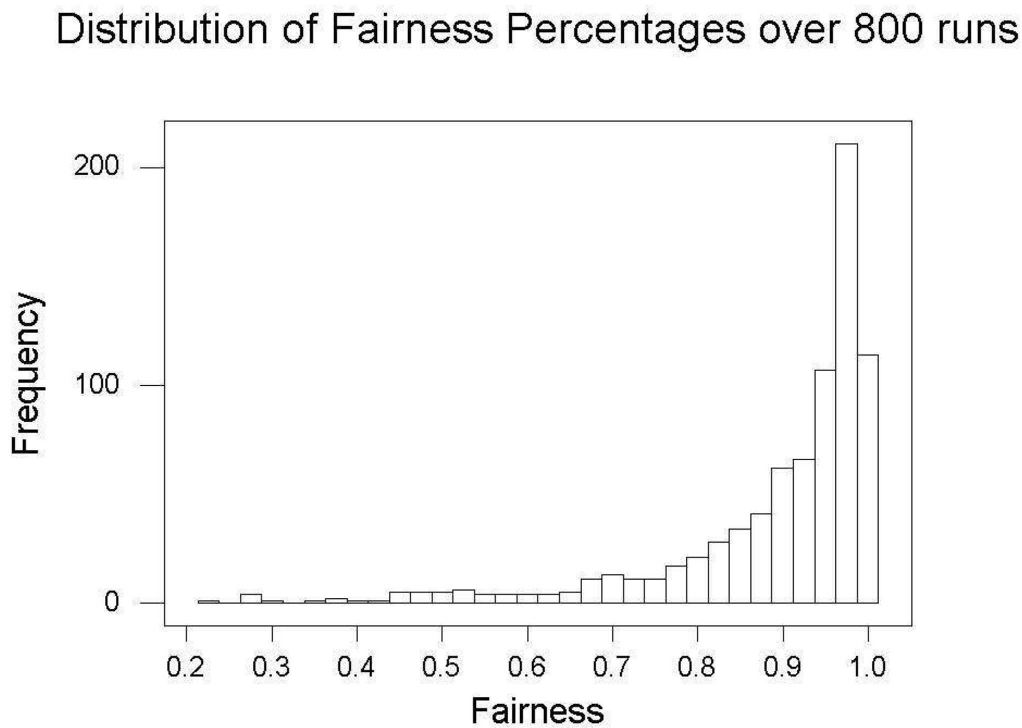

FIGURE 1 Resulting distribution of fairness evolution using 800 learned social networks

For my experiments, I used the dynamic social network model to create 800 networks and then tested each correlation matrix for the resulting fairness with 1000 randomly selected initial populations. Figure 1 shows a histogram for the distribution of fairness percentages. The mean fairness was 89.3%, but the median score was 94.6%, showing evidence of a distribution skewed heavily toward fairness evolving. In those social nets where fairness did not dominate, we can see evidence of a tight network between other demands. Table 7 shows the close preferences of 0.3, 0.4 and 0.6, as well as showing the fair strategy of 0.5 preferring opponents of 0.2. Exactly why these correlation matrices do not evolve fairness is still being investigated.

SCHELLING SEGREGATION MODEL

Evolving assortative interaction matrices can also be approached from the perspective of a spatial model. Again, I will try to make a minimal number of assumptions which I believe are reasonable and relatively benign. Under this model, we will remove the previous assumption that players can distinguish between other players based on strategy. First, players will be spatially distributed. Second, players will be allowed to change their location if they deem it unsuitable. Last, a player's goal is to maximize their utility, a common assumption in game theory. Spatial models of "Divide-the-Cake" have been explored by Skyrms and Alexander (1999), but they allowed players to only change strategies.

These assumptions can be readily modeled in a common framework from economics. Thomas Schelling's famous segregation model demonstrates that minor preferences of satisfaction within your neighborhood can have striking results for the overall distribution of individuals (Schelling 1978). He specifies a simple game to be played with pennies and dimes on a chessboard. First, place about 45 dimes and pennies randomly on the board.[3] The neighborhood of a coin is defined to be the eight surrounding squares, with both the horizontal and vertical edges wrapping around as in a torus. Next, assign certain preferences to both dimes and pennies, for instance dimes prefer neighborhoods with at least 1/3 dimes, and pennies will only be happy when surrounded by at least 1/2 pennies. Then, determine who is unhappy in the initial board, and move them to a new random location.[4] This process is repeated until either all the coins have reached stability or oscillations develop. The overall behavior of the game will gravitate toward patterns of segregation, even though both dimes and pennies would be satisfied under certain layouts of integration.

"Divide-the-Cake" naturally fits into this framework. To continue our simple model of three strategies, there are now three types of players, one for each strategy. A player is defined as unhappy in her neighborhood as follows:

$$unhappy(C_i) = \begin{Bmatrix} true: & \bar{u}(C_i, N_j) < t \times demand(C_i) \\ false: & otherwise \end{Bmatrix}$$

where $\bar{u}(C_i, N_j)$ is the average utility C_i receives against all her neighbors N_j, *demand*(C_i) is the demand of the player, and *t* is a parameter in the range [0,1] indicating the threshold a player has for receiving no cake, with 0 meaning the player is happy no matter how much cake she receives, and 1 meaning the player must fully receive her demand in order to be happy. In other words, players will be looking for neighborhoods to maximize their total possible gain. Repeatedly moving unhappy players and examining the resulting neighborhoods will expose the preferences for each strategy.

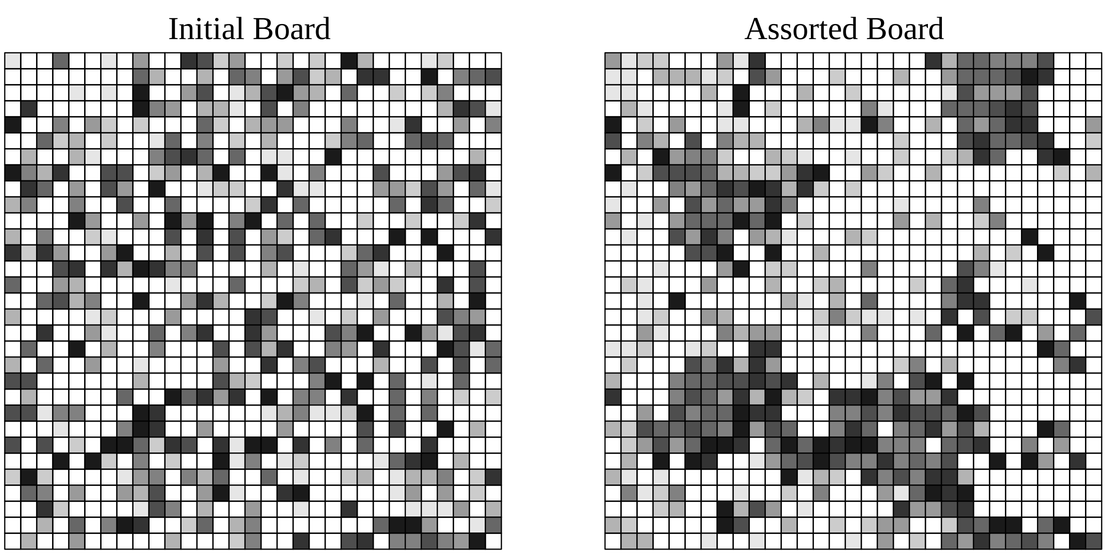

(colors range on the gradient from 0.1 = black to 0.9 = light grey)

FIGURE 2 Rearrangement of players based on Schelling Model

---

[3] Similar segregation behavior should evolve independent of the initial distribution of dimes to pennies.
[4] Schelling recommends starting at the upper-left corner and proceeding row by row. He claims the order of movement is unimportant, however, I have noticed this leads to waves of unhappy players moving down the board.

TABLE 8 Sample results of Schelling Spatial Model: resulting fairness 89%

Strategy *i*	**pref(*i,0.1)***	**pref(*I,0.2)***	**pref(*i,0.3)***	**pref(*i,0.4)***	**pref(*i,0.5)***	**pref(*i,0.6)***	**pref(*I,0.7)***	**pref(*i,0.8)***	**pref(*i,0.9)***
0.1	0.10	0.17	0.14	0.12	0.13	0.12	0.07	0.06	0.07
0.2	0.13	0.18	0.16	0.13	0.13	0.12	0.08	0.03	0.02
0.3	0.10	0.16	0.12	0.18	0.13	0.18	0.06	0.04	0.02
0.4	0.09	0.12	0.17	0.26	0.13	0.17	0.02	0.01	0.03
0.5	0.11	0.14	0.15	0.15	0.23	0.02	0.05	0.07	0.07
0.6	0.11	0.14	0.21	0.21	0.02	0.12	0.05	0.07	0.07
0.7	0.10	0.14	0.11	0.04	0.07	0.07	0.14	0.22	0.11
0.8	0.09	0.06	0.07	0.02	0.11	0.11	0.23	0.13	0.18
0.9	0.09	0.03	0.04	0.05	0.11	0.11	0.11	0.17	0.28

Simulations of the Schelling model were tested for population sizes from 1000 to 5000 using random samples of population distributions. Players were allowed to assort for 20 time steps before evolving into the next generation based on their current fitness levels. Unsettled populations were terminated after 100 generations and recorded as a failure to evolve fairness. Our new parameters for this more complex model are the size of the board and the tolerance *t* at which a person will be unhappy. A sample run can be seen in Figure 2 for initial board size of 31x31, nine player categories, with a distribution of 50 players per strategy, and tolerance value 0.75. Table 8 reports the evolved correlation matrix after 20 time steps. This was calculated by counting the neighboring strategies for each individual and then normalizing to one.

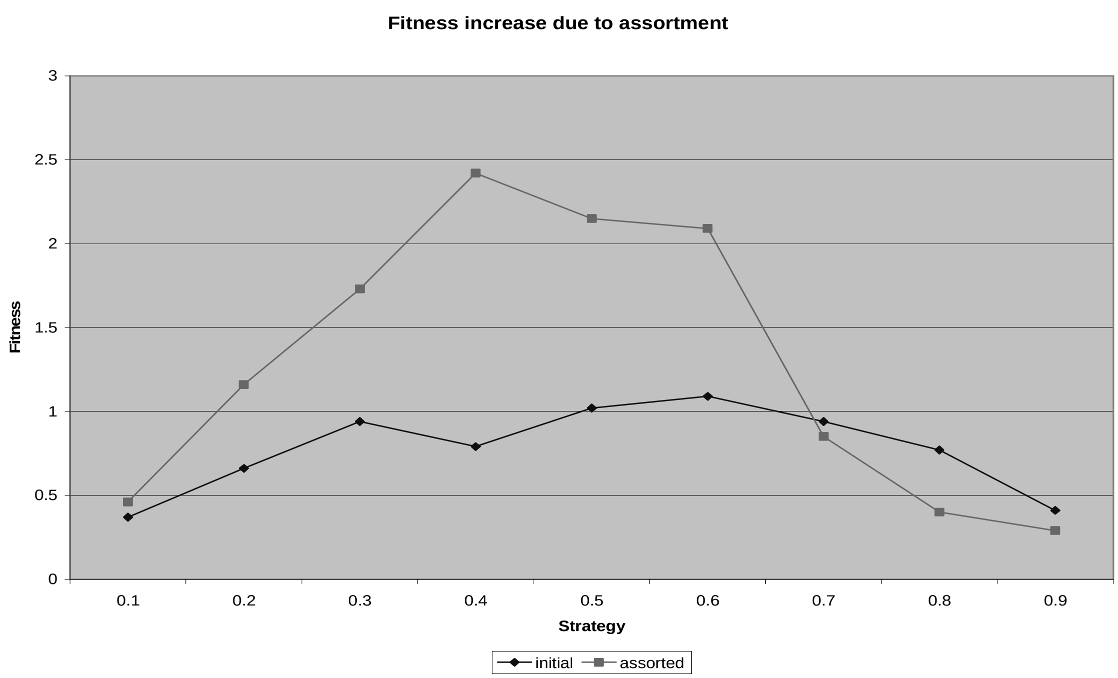

FIGURE 3 Change in Fitness due to Schelling Assortment

When looking at the preference of happy players, this run of the Schelling simulation appears to evolve a utility preference matrix similar to that in Table 3. Since not all players are happy, the utility preference matrix is an asymptote. When all players are considered, greedy individuals display an overall preference to choose themselves, due to the unavailability of suitable modest players. Figure 3 shows the change in fitness scores due to assortment of the players.

Modest players rarely move from their initial random locations. The only reason they would be unhappy is if they are lonely and have no neighbors, otherwise they are content to play against anyone. Also, fair players have a much easier time settling down and finding groups than do greedy players. This is true irrelevant of the initial distribution of players. This could be explained by the fact that fair players can either find neighborhoods of fair or modest players, while greedy players must find near exclusive modest neighborhoods complementary to their own demand to be satisfied. As more and more greedy players surround modest players, the average utility for each greedy player will falter and place them on the move again.

Settings for the board size parameter were tested for up to three times what is necessary to fit all the players. Differences in the resulting evolution of fairness were minimal, however, extra space made it easier for players to find favorable opponents. With more elbow room, greedy players can surround modest players while still avoiding each other. This shifts the correlation matrix closer to the efficiency correlation of Table 4.

Variations of the tolerance threshold produced more interesting results. Figure 4 shows the average fairness evolution when tolerance was varied from 0 to 1. Values from 0.6 to 0.85 result in close to 90% fairness, while higher values, such that a player is only happy with receiving their demand, show a return to polymorphic solutions over fair evolution.

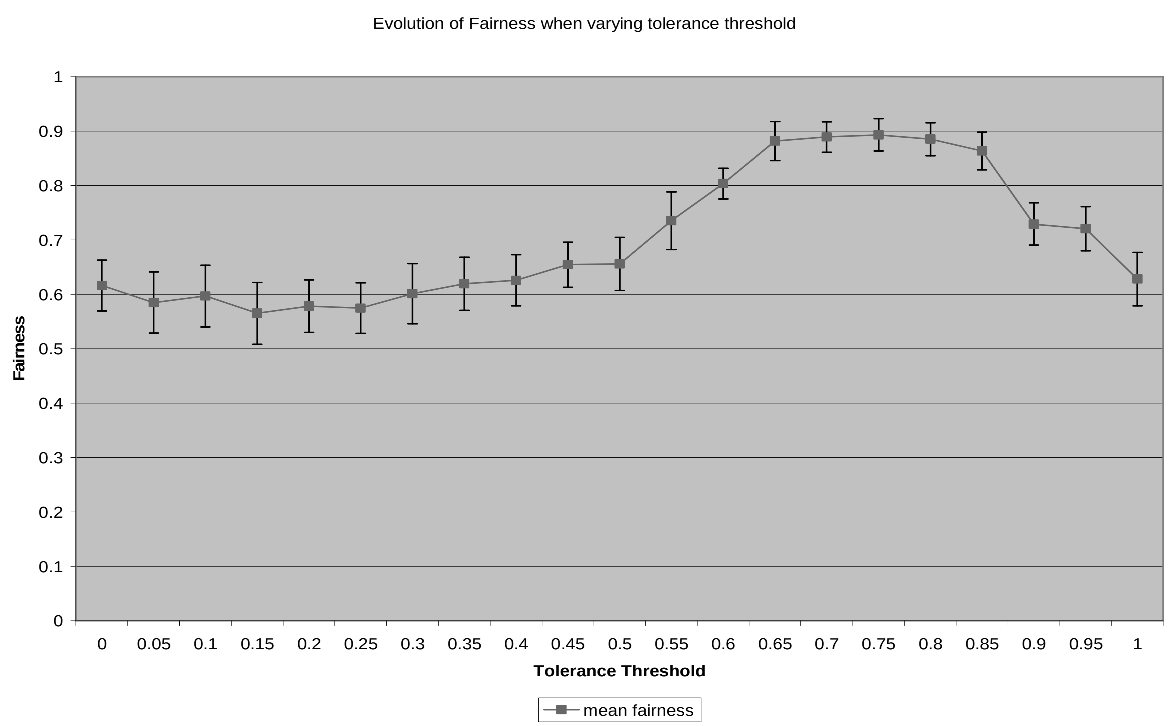

FIGURE 4 Variation of Tolerance threshold parameter from 0 to 1 (Error bars for one st.dev.)

CONCLUSION

Skyrms shows that a certain model of correlation will effectively promote the evolution of fairness across all initial populations. But once he introduces correlation, he is now open to criticisms from D'Arms *et al*. that other correlation schemes produce opposite results. I feel the examination of alternate correlation systems should also proceed in an evolutionary environment to bring out those correlations which could naturally emerge from player interactions. While certainly not robust with respect to alternative correlations, by using the approach of learning our probable parameter values we see a much more accurate picture of the model.

The use of a social network model for player types rather than actual players could be seen as overly simplistic; a more complete model would have each player learn a distinct preference vector for every other player. Also, Skyrms and Pemantle (2000) discuss other variations on their dynamic social network formation, such as reciprocal visiting and decay in memory that should be investigated in the context of the bargaining game.

The results from the spatial Schelling model are very promising. To reinforce the claims made in this paper, a number of extensions to the model should be made. First, the space of possible tolerance values should be examined. With this new parameter, we should examine ways of letting each player learn their own tolerance, as the implications of heterogeneous tolerance by strategy and by player could have very drastic implications and need to be explored. Second, unhappy players are currently randomly relocated to a new location; better relocation packages for displaced players should be explored, such that a player could select the best from *n* randomly chosen new locations. In addition, the cost of obtaining preferences as discussed in D'Arms *et al.* has been totally ignored. A cost could be assessed per player based on how many times they must move to be happy.

Each model was not entirely successful in showing a complete evolution of fairness, however, these results are significantly different than when using a random correlation and bear further investigation. The approach shown here can be readily incorporated into other agent-based models, allowing us to delve deeper into those relevant areas of the model. Although it requires an additional step to tune the model parameters, the benefits as shown can be drastic and focus our attention on essential areas instead of quibbling over irrelevant parameter values.

Acknowledgements
I would like to thank Elliott Sober and Brian Skyrms for their helpful suggestions on an earlier draft of this paper. This work has been partially researched while under the support of NLM training grant 1T15LM007359-01.

References:

Axelrod, R. 1984 *The Evolution of Cooperation* Basic Books, New York.

Barrett, M., E. Eells, B. Fitelson, and E. Sober. 1999, "Models and Reality: a Review of Brian Skyrms' 'Evolution and the Social Contract'," *Philosophy and Phenomenological Research*, 59:237-242.

D'Arms, J., 1996, "Sex, Fairness, and the Theory of Games," *The Journal of Philosophy*, 93:615-627.

D'Arms, J., R. Batterman, and K. Gôrny, 1998, "Game Theoretic Explanations and the Evolution of Justice," *Philosophy of Science*, 65:76-102, 1998.

Ernst, Z. 2001, "Explaining the Social Contract," *British Journal for the Philosophy of Science*, 52.

Schelling, T., 1978, *Micromotives and Macrobehavior*, W. W. Norton.

Skyrms, B., 1996, *Evolution of the Social Contract*, Cambridge University Press.

Skyrms, B., 1999, "Reply to Critics," *Philosophy and Phenomenological Research*, 59:243-254.

Skyrms, B., and J. Alexander, 1999, "Bargaining with neighbors: Is justice contagious?" *The Journal of Philosophy*, 96:588-598.

Skyrms, B., and R. Pemantle, 2000, "A Dynamic Model of Social Network Formation," in *Proceedings of the National Academy of Sciences of the United States of America*, pages 9340-9346.

Sober, E., and D. S. Wilson, 1998, *Unto Others*. Harvard University Press.

Village Project, The, July 28th, 2003 http://www.wsu.edu/~village/

Weibull, J. W. 1995 *Evolutionary Game Theory*, MIT Press