Am I Hot Enough?

A Self-Organizing System as Assortative Mating

Luis Alejandro Lee, University of Nottingham

Erika Frydenlund, Old Dominion University

1 Introduction

Most social systems face the problem of how to organize their members for particular goals. In economics, the concept can be related at least as early as Adam Smith. He emphasized the question of how societies manage to align the diverging interests of their members. The market provides a solution by allowing each member one of them to pursue his own interest without affecting the interests of others. But this problem is seems like a ubiquitous issue in most social sciences. As an example, voting mechanisms have the objective of organizing several divergent interests and coming up with one single decision, such as electing a single candidate out of a pool. In these cases, a centralized voting mechanism solves the problem of organization. Some political systems such as communism try to solve the problem by imposing rules from a centralized position to the rest of the system.

The fact that the problem is solved doesn't mean that modern science understands it. For Smith, although the market solved the problem, it wasn't clear the exact mechanism by which this happened, and hence he used the term "invisible hand;" for some invisible reason, the market works. For Hayek, a market price helped solve this question by gathering local information and condensing it into a common signal that can be available to the whole system. However, how organization happens is not addressed. In traditional neoclassical economic thinking and textbooks, supply and demand determines prices and quantities, and the equilibrium they form is a way in which the market is organized. Since the critique to the Walrasian auctioneer until now, how exactly that organization happens is not clear.

To narrow the problem, the focus of this exercise is to understand better how this organization occurs in a decentralized system, particularly one without the use of any common signal such as a price. We are interested in modeling explicitly the process in which agents make their decisions at the individual level, and evaluating what rules of behavior lead to organization at the aggregate level. The key aspect in our modeling approach is that agents don't have any aggregate information about the whole system, and even more, we don't endow them with any a priori preference for a particular kind of organization beyond simple two-agent pairs.

For exploring this problem, the model presented here focuses on the decision process of finding a mate. We focus on two dimensions of the model: one is the individual problem of each agent finding a partner, and the other is the aggregate patterns that are observed from these interactions. For this, we explore different simple behavioral rules at the individual level, and ask whether the interaction of agents, explicitly modeled via computational experiments, can generate an organized system. To measure organization we consider the extent to which the model generates assortative mating: this is an empirical fact in the fields of psychology and economics that humans usually marry partners with similar characteristics such as attractiveness or income.

The interesting part of our model is that there is no exogenous mechanism embedded in the agents that would a priori determine a preference for mating with similar agents, nor common knowledge about the system as a whole. We model simple agents with limited information about the rest of agents with whom they can mate, and by defining explicit rules of local interaction, we ask whether such a system can organize itself by replicating patterns of assortative mating. Our results show that the model can. Since we explicitly model the interaction of agents, the model presents a way in which the organization problem is solved without the need for any centralized mechanism.

2 THE MODEL

Our approach represents agents whose objective is to "match" or "marry" with one other agent. The model is initialized with 100 agents who constitute the dating "pool." Agents can only "meet" and select partners in the pool. The pool changes in each discrete time step, and once two agents marry they are removed from it. All agents are of the same single, undefined sex, meaning that they can marry any other agent in the pool. Agents are heterogeneous in "hotness" levels (fitness measure); they are all randomly endowed with a hotness level that ranges from 0 to 10, drawn from a uniform distribution. An agent's hotness can never be changed.

In each discrete time step, each agent in the dating pool randomly "meets" exactly one other agent also in the pool, which we shall call a "date." At the meeting, both agents' decisions are a binary choice: either "accept" or "reject" their date for marriage. If both agents "accept," they marry and leave the dating pool. Otherwise (if at least one rejects), both go back into the pool and wait until the next time step to have another random encounter.

2.1 Information assumptions and aggregate measures

Agents do not have information about their own hotness level or the distribution of hotness in the larger system. Intuitively this means that agents don't know exactly how they compare to the other agents in terms of hotness (unless some learning occurs, as explained below). So what information do they have?

In each encounter, an agent can accurately assess the hotness level of its date. It also knows whether it was rejected or accepted (i.e. the date's decision). Other than the initial benchmark scenario explained below, we assume that agents choose to accept or reject based on a simple satisficing rule: they accept to marry a date only if the date's hotness level is above a personal threshold. Agents are also heterogeneous in their threshold level; each agent is randomly given a threshold level when the model is initialized, again using a uniform distribution. We examine different behavioral scenarios in the next section. Depending on the rules used in the scenario, agents have the opportunity to update their thresholds. Four rule scenarios are discussed at length below.

In order to evaluate how organized the system is, we estimate a measure of "fairness" (without any normative or political implication). This fairness measure is the difference in hotness level between agents in each married couple, which we will call the "gap." For example, if an agent that is a 10 (in the hotness scale) marries a 1, this is considered to be an "unfair" marriage from an aggregate perspective (perhaps the agent with the lower hotness value would disagree!). Notice that assortative mating means that the gap across married agents should be low (closer to zero) as agents try to marry another with fitness levels close to their own. We consider that the system is more "organized" when the average gap across marriages is lower.

2.2 Behavioral scenarios

We consider different behavioral scenarios. The objective is to analyze which behavioral rules, if any, can lead to assortative mating.

1) Coin-flip: In the base scenario (benchmark), the agents' decisions to accept or reject their date is made by "flipping a coin," accepting with a 50% probability. Notice that here they do not take into consideration their partner's hotness level when making the decision to "marry" or not. In this scenario, we expect that many agents will be inequitably paired since the matching is entirely randomized; hence we expect a higher average gap. Also, since all agents have a positive probability of being accepted, independently of the remaining agents in the pool, we expect all of them to get married with sufficient time steps.

2) Fixed-threshold: The agents are assigned a random threshold value independent of their individual hotness levels, as described above. When meeting a date, they use this threshold to determine if the other agent is "hot enough" to marry. This is the decision satisficing rule. The threshold level remains the same in all periods. In this scenario, some agents that are too low on the hotness scale, that "aim too high" (i.e. have a high threshold value), or that are both too hot and aim too high, might end up never being matched.

3) Updating-threshold: The agents in this version of the model have random threshold values to start (as in last scenario), but they update these values based only on the decision of their date. The updating rule is quite simple, and applies only if at least one of the agents rejects (i.e. if no marriage happens): if the date choses to reject, decrease the threshold by one point (we test other values as well, shown in the results section). If the date choses to accept, increase the threshold by one point. This simple updating rule reflects agents updating their aspiration levels based on experience; agents constantly being rejected will start accepting dates with lower hotness levels. The ones constantly being accepted will keep increasing their threshold values, accepting to marry only hotter dates. Notice that this updating rule reflects agents that learn from experience how high (or low) can they "aim," or how "picky" can they be when selecting a partner with whom to mate. We expect in this case that agents will find much more equitable pairs as they have the potential to dynamically adjust their threshold based on their apparent value expressed through individual interactions.

4) Training: In the previous scenario, the first time step allows agents to pair without much information about the system. In that scenario, the first time step before an updating opportunity may allow agents with very disparate hotness values to pair. In this training scenario, agents have the opportunity to "practice" pairing by meeting others and updating their threshold values without committing to marrying and leaving the system. In this case, agents meet at each time step in the training period, but independent of the decisions, there are no marriages. Intuitively, this can be thought as a period of adolescence, where agents date before actually considering committing to marriage. In this case, the agents should develop a more refined sense of their own value, or at least of how high can they aim. This is expected to lead to more equitable matches.

3 RESULTS

3.1 Measures

Two measures provide insight into the performance of agents in the system to solve the matching problem. First, after pairing, we record the difference in hotness levels between the two married agents (the gap). The gap of all married couples are averaged at the end of each simulation run, and then averaged across runs. Second, we record the number of time steps (ticks) required for the system to converge (to stop creating new marriages). In some cases, the convergence happens when all of the agents find a marriage partner. In other cases, however, not all agents find a pair, but the system stabilizes at a certain percentage of pairs. Perpetually unmatched agents occur when agents' thresholds exceed the available potential partners' hotness levels remaining unpaired in the model.

3.2 Simulation

Each scenario was run 1,000 times using the BehaviorSpace feature of NetLogo. For the threshold updating/learning model, the amount an agent updates at each interaction was tested from one to five at increments of one, each run 1,000 times. For the training period model, agents were allowed training intervals of 10, 50, and 100 ticks, also run 1,000 times each. Tables 1-3 provide the average values for each of these scenarios over 1,000 simulation runs.

Table 1: Average Equity Gap and Time Steps over 1,000 Runs

Rule:	Coin-flip	Fixed	Updating=1	Training Period
Equity Gap	3.333	2.764	2.265	1.709
Ticks to Convergence	16	48	191	174

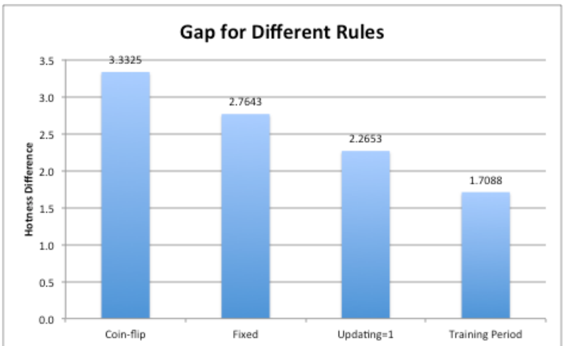

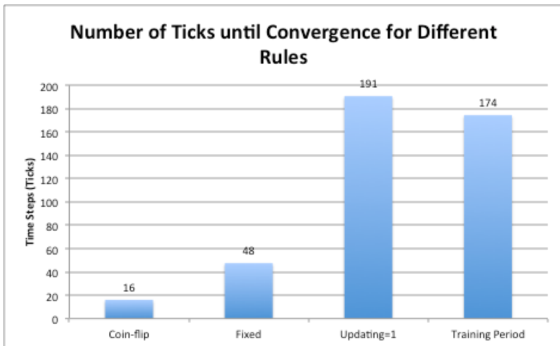

In Table 1, we observe the results of the simulation runs for the four different rule scenarios. The Coin-Flip rule converges fastest, but results in a relatively large average equity gap indicating that while agents found mates quickly, they are not paired very equitably. In the second scenario where agents have randomly assigned thresholds, they still find mates relatively quickly, but the average equity gap is smaller than in the Coin-Flip case. When we introduce learning into the model by allowing agents to update their thresholds based on the decision of the other agent at each meeting, the system takes much longer on average to converge but results in an even more average equitable pairing. Finally, we look at the system where agents are allowed to learn for a predetermined number of iterations before committing to marriage. In this case, agents find suitable mates slightly more quickly on average than the scenario without a training period and significantly improve the average equity gap between paired agents. We see here that with a very simple updating rule and some time for agents to attempt to learn their own 'hotness' value, agents tend to pair with others of a similar hotness value. This result confirms the occurrence of assortative mating as seen in the psychology and economics literature.

Table 2: Average Equity Gap and Time Steps over 1,000 Runs Across Different Threshold Updating Values

Learning Values:	**1**	**2**	**3**	**4**	**5**
Equity Gap	2.265	2.323	2.485	2.640	2.800
Ticks to Convergence	191	105	76	58	47

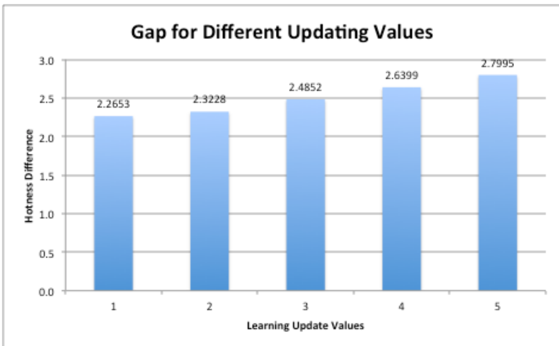

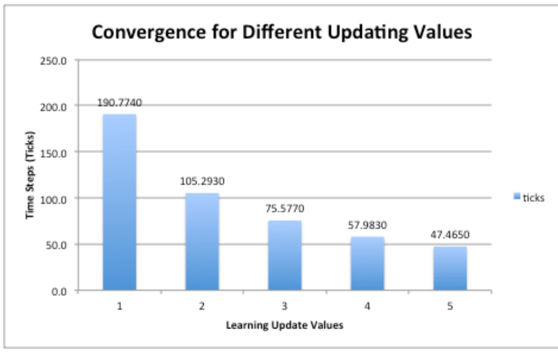

Table 2 shows the results from varying the amount that an agent updates its threshold after an encounter with another agent. Recall that in this scenario, the agent updates its threshold based on the date's decision when they meet. If the agent was 'accepted,' then it increments its threshold by the Learning Value amount; if it was 'rejected,' then it decreases its threshold by that amount. The simulation was run 1,000 times for each of the learning values 1 – 5 and reported in Table 2. Notice that the results in the scenario where learning value is five is very close to the average gap and ticks to convergence in the random assignment scenario in Table 1. A high learning value causes the agent to over-correct, introducing a noise into the system that results in very similar average outcomes to the random scenario. Though it takes on average the longest to converge, the smallest learning value provides the best outcome in terms of equitable pairing.

Table 3: Average Equity Gap and Time Steps over 1,000 Runs for Different Lengths of Training Period

Training Period:	**10**	**50**	**100**
Equity Gap	1.709	1.705	1.705
Ticks to Convergence	164	154	156

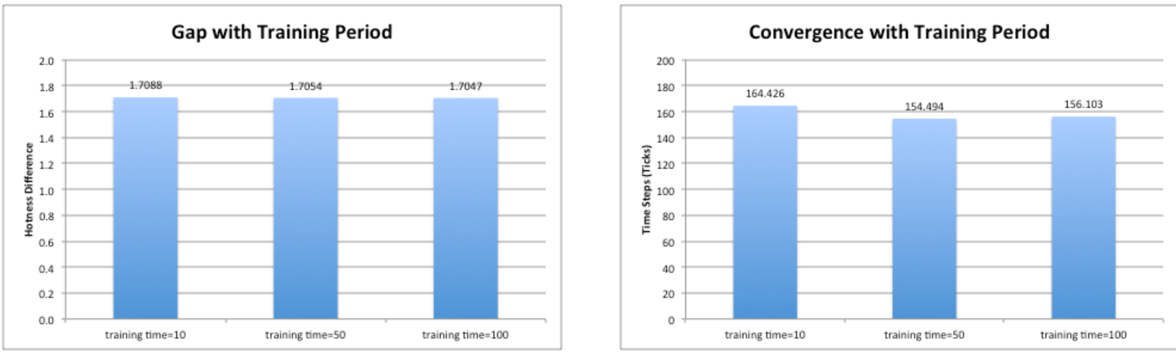

One observation during the initial updating model was that the most pairing happened during the first few iterations. This resulted from agents pairing based on their initial, randomly assigned hotness values before "getting out into the world" to learn more about their own hotness values from interactions with others. To address this issue, we implemented a learning period of either 10, 50, or 100 ticks—or rounds of meeting—before anyone in the model begins to commit to marriage with another agent. Table 1 showed that this strategy improved the outcome over the other rules, but Table 2 shows very little effect past the training period of 10 ticks.

Ultimately, these results indicate that assortative mating comes about through simple agent rules and with very little information about the broader world of agents. Particularly in the case where the agents were allowed even a short a warm-up period, we see that agents find much more equitable partnerships on average that the other scenarios. Similar results have appeared in other psychological studies, most recently in a paper that looked at assortative (Xiea, Chenga, and Zhoua, 2015).

4 Extensions

The model presented in this paper is very simple, with two agents meeting and staying together. Potential extensions for exploration might include allowing agents to divorce and re-enter the dating pool to find an even more suitable mate. This might allow agents to continue learning from interactions and result in even more equitable matches as those who paired very early in the system before learning about their own fitness would have a chance to update partners later. Another extension for future work might allow agents to develop and maintain a reputation that affects their mating preferences or possibilities in future runs. Additionally the potential for asymmetric information among agents could result in some interesting self-organizing dynamics within the model. In the current version of the model, we limit self-organizing to two agents in order to capture the behavior of couples in a mating environment. Future extensions might consider allowing more agents to pair together, which could represent a system where agents form alliances rather than married couple pairs (or polygamist societies!).

5 Applications

We propose here some applications of this generalized model to other contexts within the social sciences.

Academia: In the world of academia, scholars attempt to publish articles in the most prestigious journal possible. In this case, prestige of the journal and quality of the scholarly article could be represented as ‘hotness’ to the other agent. The scholar and journal editor each have a dynamic threshold that they update based on feedback from one another. If the journal exceeds the author’s threshold, he will submit his article for review. If the article exceeds the journal editor’s threshold (setting aside peer referees for the moment), then she will accept the article and publish.

Matching markets: The model can also be related to the matching market literature. Algorithms such as the Shapiro-Galey explore matching under different conditions. Environments such as the job market can be interpreted as agents trying to find the best match they can, without complete information of the whole distribution of candidates (unlike in Shapiro-Galey). As in our model, internships or pos-doc positions can be considered as training periods where both universities and Phd students engage in temporary contracts in order to learn and adapt their threshold level. Afterwards, both parties engage in the search of more stable matching, such as lectureships leading to tenured positions.

Coalition formation: In a theoretical world where agents look to form coalitions, the assortative mating model may have additional applications, particularly in an extension where agents are allowed to pair with more than one partner. In this case, fitness levels determine the interest of each partner in joining the coalition. Those that are accepted regularly may increase their thresholds and become “pickier” about their permanent partner. In a model where agents are allowed to “divorce,” we may see dynamic collation formation where agents pair with those similar to their own fitness level similar to the dynamics seen in the assortative mating model.

6 References

Xiea, Yu, Siwei Chenga, and Xiang Zhoua. 2015. “Assortative mating without assortative preference.” Proceedings of the National Academies of Sciences of the United States of America. vol. 112 no. 19. pp. 5974–5978. DOI: 10.1073/pnas.1504811112