The (Spatial) Evolution of the Equal Split

J. McKenzie Alexander
Department of Philosophy
University of California, Irvine

The replicator dynamics have been used to study the evolution of a population of rational agents playing the Nash bargaining game, where individuals' "fitnesses" are determined by individuals' success in playing the game. In these models, a population whose initial conditions was randomly chosen from the space of population proportions only converges to a state of fair division approximately 62% of the time, unless artificial correlations are introduced into the model. Spatial models of the Nash bargaining game exhibit considerably more robust convergence properties. These properties are considered at length, and a sufficient condition for convergence to fair division is proved.

1. Introduction. In 1974, Nydegger and Owen conducted an experiment to test the validity of the Nash axioms. Thirty pairs of subjects participated in three rounds of bargaining, each successive round of bargaining imposing a different (and increasingly complicated) condition upon the subjects. In all rounds of bargaining, subjects played a variation of "divide the dollar", the simplest version of the general bargaining situation considered in Nash's original treatise on the subject[7]. Although Nydegger and Owen found subjects deviating from one of the axioms under certain circumstances[1], in the special case where the two players were asked to divide a dollar under completely symmetric circumstances, Nydegger and Owen found unanimous agreement on the 50–50 split.

This result is not considered controversial. Nydegger and Owen conducted this experiment primarily to establish a baseline for comparison for

[1]The results of the third round of bargaining suggest that subjects attempt interpersonal comparisons of utility, leading them to violate Nash's fifth axiom of the invariance under linear transformations of utility.

later ones.[2] However, it is precisely the *uncontroversial* nature of the Nydegger and Owen result that I want to focus on in this paper. The question may be put this way: why is the 50–50 split, under completely symmetric circumstances, the only "natural" or "expected" solution to the division problem?

Without further explanation, this question may seem silly. After all, most reasonable moral theories include equal division under completely symmetric circumstances, if not as a fundamental principle, as an easily derived consequence from fundamental principles. It is a consequence of Kant's categorical imperative, Aristotle's rule of distributive justice, and Rawls' theory of justice, to name but a few. Yet philosophers with certain naturalistic inclinations might view with distaste the importation of the categorical imperative (or Aristotle's rule, or the principles underwriting Rawls' theory of justice) to establish something as simple as equal division under completely symmetric circumstances. We should not need to appeal to principles as powerful as the categorical imperative to establish such a simple result. Ideally one should be able to explain the perceived moral phenomenon through nothing more than appeal to informed rational self-interest.

As Skyrms[10] observes, though, this task is more difficult than it first appears. Two player's claims are in an equilibrium in informed rational self-interest when each player's claim is optimal given the other player's claim. I.e., neither player can do any better against the other by changing his or her claim.[3] For the game considered by Nydegger and Owen, many equilbria other than equal division exist. For example, the case where player one receives \$.66 and player two receives \$.34 is another one, as is the case where player one receives \$.02 and player two receives \$.98. How is the equilibrium of equal split selected out of the many possible equilibrium?

Skyrms suggests we view the question of equilibrium selection within an evolutionary game theoretic context. We suppose each individual has a certain strategy (how much of the dollar they wish to keep) and interacts with other players using that strategy. At the end of every generation, strategies

[2] The agreement of two players on the 50–50 distribution under completely symmetric circumstances does not test all of the Nash axioms—only feasibility, individual rationality, Pareto optimality, and symmetry are required. Under these conditions, the Kalai-Smordinsky axioms also lead to a 50–50 split.

[3] Also known as a *Nash* equilibrium, in honor of John Nash. A *strict* Nash equilibrum is one where each player does strictly worse by changing his or her strategy. In the case of "divide the dollar", almost every equilibrium is a strict Nash equilibrium.

replicate with a certain probability determined by their level of success. More successful strategies are used by more people, less successful strategies are used by fewer. If it should turn out that the evolutionary dynamics almost always takes a population from an initially random assignment of strategies to a state where equal division dominates, then we would have a solution to the problem of equilibrium selection. More importantly, though, we would also have the beginnings of a naturalistic explanation for why people, under completely symmetric circumstances, always agree on the equal split.

First, a minor confession: in the above I used the phrase "the evolutionary dynamics", implying that there is one and only one type of evolutionary dynamics. This is false. The type of evolutionary dynamics one uses usually depends upon how one models the population. If we believe that the strategies individuals follow are "genetically" determined;[4] that reproduction occurs asexually; that individuals breed true; that the large size of the population allows us to identify an agent's individual fitness of an agent with the expected fitness of an agent; and that the structure of the population (through random interaction of the agents or various mixing forces) makes any two members of the population equally likely to interact, we arrive at the replicator dynamics of Taylor and Yonker[11]. Other dynamics arise if we select a model using discrete agents, spatial constraints, and so on.

Using the replicator dynamics to model a population of individuals playing "divide the dollar" with a stack of dimes, Skyrms obtained the results for a series of 100,000 trials shown in table 1. Although fair division has

Polymorphism	Count
Fair division	62,209
4–6	27,469
3–7	8,801
2–8	1,483
1–9	38
0–10	0

Table 1: Convergence results for replicator dynamics—100,000 trials

a sizeable basin of attraction (roughly 62%), a significant fraction of the

[4] I use scare quotes here because one does not need assume that the strategies are genetically determined in order to obtain the replicator dynamics. The replicator dynamics apply equally well for the model of cultural evolution I am considering here. However, I find the language of biology the most convenient way to express the traditional assumptions by which one obtains the replicator dynamics.

populations converged to one of the other Nash equilibria. As far as providing an explanation for why people always opt for equal division in symmetric circumstances, this attempt falls a bit short. It requires we postulate that the initial state of the population (was not in the basin of attraction of any other Nash equilibria. Consequently, it removes much of the normative force of the rule. That is, it is not the case that one should always opt for equal division in symmetric circumstances; rather, one should opt for equal division in symmetric circumstances only because, given the particular evolutionary path our society has followed, it is the best thing to do. If our society had followed a different evolutionary path, we would be better off demanding forty or sixty cents of the dollar.

We can recapture some of the normative force of this elementary principle of distributive justice by showing that one should always opt for equal division in symmetric circumstances, regardless of the initial state of the world. In part, this problem can be reduced to eliminating all polymorphic pitfalls other than fair division. If we can construct a model which always converges to fair division, regardless of the initial state of the population, then we have a plausible explanation for why people always opt for equal division in completely symmetric circumstances.

The problem of selecting one equilibrium out of the many possible ones has been given serious thought by game theorists. Peyton Young[13, 12] showed that if we allow members of the population to experiment with new strategies at random, and take the limit as the probability of someone experiementing with a new strategy approaches zero, then the ratio of time the population spends in equal division approaches one. Unfortunately, this result means only that, in the long run, the population will spend most of its time at fair division. If the population happens to get trapped in one of the polymorphic pitfalls, we might have to wait an arbitrarily long time before individual experimentation moves the population outside of that polymorphism. Given the rapid rate at which human populations have converged upon the principle of equal division in completely symmetric circumstances, this explanation does not seem satisfactory. Perhaps we should consider an alternative approach.

2. Spatial bargaining games.

2.1. Definitions. Let $P = \{ 1, \ldots, N \}$ be a population of individuals. Let G be a two-person noncooperative game with a strategy set $\mathbb{S} = \{ s_0, s_1, \ldots, s_M \}$.

If s_i and s_j are strategies held by players p and q, respectively, denote the payoff to player p by $G(s_i, s_j)$ and the payoff to player q by $G(s_j, s_i)$. Notice that in doing so we are assuming the game to be represented by its normal form description. This may not always be a correct assumption. Fictitious play dynamics, allowing players to take previous moves into consideration when deciding what to do for their current move, would be more naturally modelled using the extensive form representation.

Π is a *population structure* if $\Pi = (P, E)$ is a connected graph. The requirement of connectivity simply prevents us from having two separate populations models "inside" of a single model. (This would correspond to a spatial variant of the patch models of Durrent and Levine[1].) The general shape of the graph also determines the topology of the world, of which we only distringuish between whether the world is *bounded* or *unbounded*. A spatial model is bounded if the corresponding population structure has a well-defined perimeter (e.g., E is a planar graph), and a spatial model is unbounded if it is not bounded. Unbounded spatial models appear quite frequently in the literature; simple examples use population structures based a square lattice with the top and bottom connected, forming a cylinder, or with both pairs of opposite sides connected, forming a torus.

Although unbounded models do not model spatial populations accurately, there are advantages to using them. The absence of a boundary often means all individuals have the same number of neighbors, which simplifies the analytic treatment. Moreover, this also means one does not have to worry about edge effects affecting the behavior of the model. On a more pragmatic level, the absence of a boundary often simplifies coding since one does not need to consider separate cases for each stage of the model (individuals on the boundary and individuals in the interior). Each of these advantages may also be viewed as a disadvantage. In real life, everyone does not have the same number of neighbors, and edge effects do have a proper place in analyzing some phenomena (people do live on the edge of town). Therefore, all models considered in this paper use bounded population structures in an attempt to increase the level of accuracy of the model.

The *neighborhood* of a player p, denoted $N(p)$, is the set of all players q such that $(p, q) \in E$. In a single round of play, a given player p interacts with every player in his *interaction* neighborhood $N_i(p)$ and updates his strategy at the end of each generation by comparing his success level with that of every player in his *update* neighborhood $N_u(p)$. I distinguish between the interaction and update neighborhoods as there is no *a priori* reason

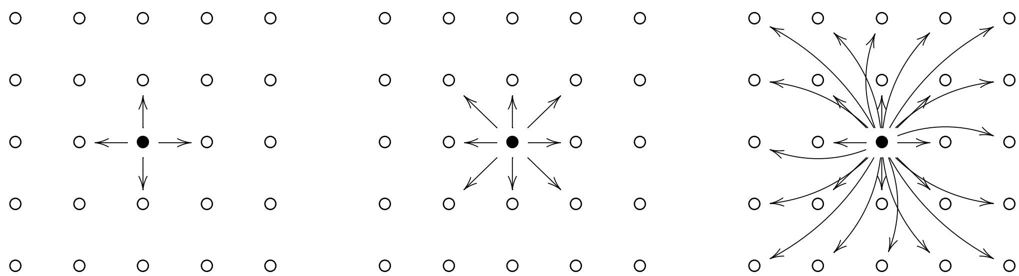

(a) von Neumann (b) Moore 8 (c) Moore (24)

Figure 1: Three common neighborhoods defined on a square lattice.

to assume that these neighborhoods are equal. However, I also follow the majority of papers in this literature by assuming $N_i(p) = N_u(p)$, which I denote by $N(p)$.

Commonly studied neighborhood types when Π has the form of a rectangular lattice are listed in figure 1. These diagrams specify directional offsets identifying the neighbors of a player (identified in the diagram by '$\bullet$'). Since the models of this paper are bounded, players on the boundary have fewer neighbors than those in the interior. In general, the model of this paper allows the neighborhoods to be any arbitrary subset of the Moore (24) neighborhood.

2.2. Dynamics. Let $\sigma_t(p)$ be the score of player $p \in P$ at the end of generation t, and denote the strategy held by a player $p \in P$ at the end of generation t by $\xi_t(p)$. The score of player p is defined as:

$$\sigma_t(p) = \sum_{q \in N(p)} G(\xi_t(p), \xi_t(q)).$$

This assumes that players do not make mistakes during interactions.

In this paper we allow for three different update rules, each rule having a certain degree of plausibility, and attempt to trace their affect on the limit form of the model. The general question of how one's choice of the update rule affects the limit form of the model remains an open and difficult problem.

Imitate the best neighbor.. This is the most common update rule in the spatial modeling literate (see Nowak and May[8, 9], Lindgren and Nordahl[6],

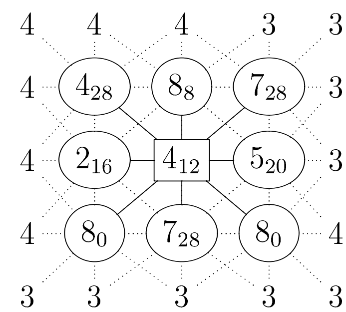

Figure 2: The possibility of ties among neighbors when updating (subscripts indicate the score).

Huberman and Glance[4], and Epstein[2].) Each player p looks at her neighbors and mimics the strategy of the neighbor who did the best, where "best" means "earned the highest score." As figure 2 shows, though, ties can occur between several players in the neighborhood of p; an additional rule needs to be given which, in such circumstances, selects a unique strategy to adopt. (In all cases it is assumed that p does not change her strategy unless one neighbor did strictly better than her.) Formally, let

$$M_t(p) = \{\, \xi_t(q) : \text{for all } r \in N(p),\ \sigma_{t+1}(q) \geq \sigma_{t+1}(r) \,\}$$

That is, $M_t(p)$ is the set of all maximally scoring strategies held by players in the neighborhood of p after the most recent round of interactions. Then,

$$\xi_{t+1}(p) = \begin{cases} \xi_t(p) & \text{if } \sigma_{t+1}(p) \geq \sigma_{t+1}(q) \text{ for each player } q \in N(p). \\ \xi_t(q) & \text{if } q \in N(p) \text{ and } \sigma_{t+1}(q) > \sigma_{t+1}(p) \text{ and } \sigma_{t+1}(q) > \sigma_{t+1}(r) \\ & \text{for all } r \in N(p) \setminus \{\, q \,\}. \\ \text{rand}(M_t(p)) & \text{otherwise} \end{cases}$$

where "rand(Q)" means to randomly choose a strategy from the set Q.

The model of this paper assumes that the number of players in $N(p)$ who follow a given maximal strategy s affects the likelihood that p will choose to adopt s. This seems reasonable since, if several neighbors of p follow s and earn the maximal score of $N(p)$, it would be foolish of p to ignore this information. The simplest way to take this information into account would be for p to let the probability of choosing a maximal strategy s be a linear function of the number of people in her neighborhood who follow that strategy. (More complicated functions could be used to model risk-averse players who require a certain number of neighbors to follow a maximal strategy before they consider adopting it.) For simplicity, I assume that if the number of players in $N(p)$ using maximal strategy s is n_s, then the probability of p choosing to adopt s is $\frac{n_s}{|M_t(p)|}$.

Imitate with probability proportional to success.. Here, as before, each player p compares her score with those of her neighbors, modifying her strategy only if at least one neighbor did strictly better than her. However, instead of ignoring those players who did better than p but not well enough to include their strategy in the set of maximal strategies, this update rule assigns to every neighbor who did better than p a nonzero probability that p will adopt their strategy. Formally,

$$\xi_{t+1}(p) = \begin{cases} \xi_t(p) & \text{If } \sigma_{t+1}(p) \geq \sigma_{t+1}(q) \text{ for each player } q \in N(p), \\ \xi_t(q) & \text{With probability } \Pr(p \leftarrow q), \text{ if } \sigma_{t+1}(q) > \sigma_{t+1}(p). \end{cases}$$

where "$\Pr(p \leftarrow q)$" denotes the probability of p adopting q's strategy. Exactly how probable it is that p will adopt q's strategy (where q is a neighbor of p) depends on how players treat the relative success of q.

The relative success of q is simply the difference $\sigma_t(q) - \sigma_t(p)$ between the two scores. The greater this value, the more likely it is that p will adopt q's strategy. As before, there are a number of ways p might use this information when making her decision, the simplest being to let the probability of adopting q's strategy vary proportionally with the magnitude of $\sigma_t(q) - \sigma_t(p)$. We define this probability as follows:

$$\Pr(p \leftarrow q) = \frac{\sigma_{t+1}(q) - \sigma_{t+1}(p)}{\displaystyle\sum_{\substack{r \in N(p) \\ \sigma_{t+1}(r) > \sigma_{t+1}(p)}} \sigma_{t+1}(r) - \sigma_{t+1}(p)}.$$

Imitate best expected payoff.. Under these dynamics, players calculate the expected payoff of each strategy in their neighborhood and select the one with the highest value. As with *imitate the best neighbor*, though, the possibility of ties exists and so some kind of tie-breaking rule needs to be given. Let $S_p(i) = \{\, r \in N(p) : \xi_t(r) = i \,\}$ denote the set of all players in

$N(p)$ who follow the strategy i. Formally, we have

$$\xi_{t+1}(p) = \begin{cases} \xi_t(p) & \text{if } \sigma_{t+1}(p) > \sigma_{t+1}(q) \text{ for all } q \in N(p), \\ \xi_t(q) & \text{if for all } u \in N(p) \text{ with } \xi_t(u) \neq \xi_t(q) \\ & \displaystyle\sum_{\substack{r \in N(p) \\ \xi_t(q) = \xi_t(r)}} \frac{\sigma_{t+1}(r)}{|S_p(\xi_t(r))|} > \sum_{\substack{r \in N(p) \\ \xi_t(u) = \xi_t(r)}} \frac{\sigma_{t+1}(r)}{|S_p(\xi_t(r))|} \\ \text{rand}(M_t(p)) & \text{otherwise.} \end{cases}$$

where $\overline{M_t}(p)$ is the set of all strategies in the neighborhood of p which had the maximal average score of $N(p)$. We assume the same tie-breaking method as used for *imitate the best neighbor*, which increases the probability of selecting a strategy from the set $\overline{M_t}(p)$ based on the number of players around p who follow it.

2.3. Synchronicity. We assume all updating occurs synchronously. There has been considerable debate over the appropriateness of using synchronous dynamics. The original papers of Nowak and May[8, 9] on the spatialized prisoner's dilemma used synchronous dynamics. However, this assumption was later criticized by Huberman and Glance[4] on the grounds that synchronous dynamics lead to stable equilibrium states which did not appear when asynchronous dynamics were used. Since then, asynchronous dynamics have typically been preferred, as the more recent papers of Hegselman[3] and Epstein[2] show.

However, I am not convinced that asynchronous dynamics necessarily provide a more accurate model. Most implementations of asynchronous dynamics construct a list of all agents in the model and walk through the list, one agent at a time, performing the required calculation or computation. At the end of each generation, the list is permuted, modifying the update order for the next generation. This type of dynamics does not seem any more "realistic" than synchronous dynamics; although agents do not update their strategies in the rigid lock-step manner suggested by synchronous dynamics, neither do they update their strategies in this carefully orchestrated manner, where only one person updates at any given point in time.

If we really want a more realistic way of handling the update dynamics, we need to strike a middle ground between these two extremes. One possibility would be to partition the set of agents and place each partition into a list—updating all the agents belonging to the same partition synchronously, but update only one partition at a time (using the list ordering). This method captures the fact that *many* agents in a population do update simultaneously (or nearly simultaneously), without requiring that everyone in the population does. However, I leave the examination of that dynamic for another paper.

3. Convergence Results.

3.1. Dependence on initial spatial configuration. Spatial models, unlike models using replicator-like dynamics, allow for wide variation in the initial conditions even when the initial population frequencies are held constant. Consider a world initialized according to the vector:

$$\langle .1, .1, .1, .1, .1, .1, .1, .1, .1, .05, .05 \rangle$$

(this vector has eleven elements since players are permitted to demand nothing). In this world, exactly 1,000 of the 10,000 members, begin by following the strategy of demand nothing.[5] Since players are assigned unique and distinguishable spatial positions, there are $\binom{10000}{1000}$ possible ways of assigning the strategy demand 0 to population members. For each of these assignments, there are $\binom{9000}{1000}$ ways of assigning the strategy demand 1 to players. All in all, there are

$$\binom{10000}{1000}\binom{9000}{1000}\cdots\binom{1000}{500}\binom{500}{500} \approx 10^{2817}$$

distinguishable ways of assigning strategies to players which conform to the initial vector above.

Given this, one might wonder to what extent the initial spatial distribution of strategies affects the final convergent state (if one exists). One can easily construct cases which exhibit extreme sensitivity to the initial spatial distribution. As a trivial example, consider a world containing 9991 players who demand 9, eight players who demand 1, and exactly one player who demands 5. Assume this model uses the Moore (8) neighborhood with imitate the best neighbor dynamics. A model initialized in accordance with

[5] In practice, the model used in this paper does not attempt to match the population proportion *exactly* since many vectors chosen at random from the space of initial population proportions do not correspond to actually attainable states in a discrete world. For example, the vector whose first coordinate is .101010101 is an admissible element in the space of initial population proportions (provided the sum of the coordinates equals 1), yet it fails to correspond to an attainable state in a discrete world with 10,000 members. (Notice, though, that it is attainable in a larger discrete world.)

this initial population vector exhibits sensitive dependence on initial conditions: most of the approximately 10^{2817} strategy assignments will lead to a unusual polymorphic state containing the strategies demand 1, demand 5, and demand 9.[6] In the vast majority of cases, the demand 5 strategy will not disappear as it will be surrounded by players who demand 9; since the strategy of demand 9 is incompatible with itself, and with the strategy of demand 5, neither the player following demand 5 nor his neighbors will receive nonzero scores. Consequently, the player following demand 5 has no incentive to switch strategies and will not; a similar argument applies for the players following demand 9. Other cases, though fewer in number, lead to the elimination of fair division, leaving the world in a 1–9 polymorphism, which *is* a stable polymorphism of the replicator dynamics. But in both of these cases, elimination of fair division and the three-strategy polymorphism, the world will settle into a stable cycle of length two. Now consider the case where all eight demand 1 strategies surround the sole demand 5. In this case, the demand 5 strategy will spread to all of its surrounding neighbors at the end of the first generation, leaving a 3×3 block of demand 5 players in a world consisting solely of demand 9. From this point, the strategy of fair division will spread to the entire world as none of the demand 9 strategies will ever earn a nonzero score, whereas all of the demand 5 strategies will.

The above example strikingly demonstrates the sensitive dependence on initial conditions often associated with spatial models. Given the extraordinarily large number of possible spatial configurations for a single vector, one might wonder whether this would make it impossible for us to say anything about the general convergence properties of such models; if, given a particular vector from the space of initial population proportions, the final state that world converges to depends entirely upon the spatial distribution of strategies, we cannot say much at all about the general convergence properties of the model. There are simply too many initial spatial distributions to consider. At best, one could say that *when* the model is started in state X_1 it converges to final state Y and that *when* the model is started in state X_2 it also converges to final state $Y \ldots$ Does the sensitive dependence on initial conditions noted above prevent us from legitimately drawing inferences about general convergence properties of the model from the outcomes of sample trials?

[6] This polymorphism is unusual in that it is not a stable polymorphism of the replicator dynamics.

I do not think so. We must recognize that the question of whether models exhibit sensitive dependence on initial conditions (as described above) can be separated from the question of how widespread such behavior is. If the model exhibits sensitive dependence on initial conditions very infrequently, we may draw inferences about the general convergence properties of the model by considering the outcome of sample trials. We may determine the extent to which the model exhibits sensitive dependence on initial conditions by considering the outcomes of repeated trials where we hold the initial vector of population proportions constant for all trials.

To do this, one hundred vectors from the space of population proportions were randomly selected. Each vector specified the initial frequencies of strategies in the population for one hundred trials; however, each trial used a different spatial distribution of strategies among members of the population. Of the 100 vectors, 97 showed no disagreement on their final convergent state. The three initial vectors which did have disagreement had remarkably little. Two vectors converged to a state of fair division 91 and 96 times out of a hundred (the remaining trials converged to the 4–6 polymorphism); one vector converged to the 4–6 polymorphism 99 times out of a hundred. The remaining trial for the last vector did not converge to any of the "standard" polymorphisms of the replicator dyanamics: this last trial converged to a stable cycle of length two containing the strategies demand 3, demand 4, and demand 6! These polymorphisms, though rare, can exist in spatial models.

The three vectors showing sensitivity to the spatial distribution of strategies are listed in table 2. Obviously, the first vector was the one which led to 4–6 polymorphisms (almost) exclusively. The second and third vectors allowed the spatial distribution of strategies to influence the final convergent state because, with only .17% or .09% of the players following demand 5, unlucky positioning of this strategy could lead to its elimination in the first few generations. Given the extremely few number of players assigned the strategy of demand 5, I find it surprising that worlds initialized according to these initial frequencies agreed so often on their final convergent state. This shows, in part, that the initial spatial distribution on the final convergent state becomes important only when the frequency of certain strategies becomes extremely low (i.e., when the initial vector lies close to the boundary of the simplex space of initial population proportions).

As a further check, twelve extended series of 10,000 trials each were performed using the initial state vectors $\bar{v}_i = (s^i_0, \ldots, s^i_{10})$, for $i = 0, \ldots, 11$. (These runs were performed exactly like the previous ones, the only exception

dem. 0	dem. 1	dem. 2	dem. 3	dem. 4	dem. 5	dem. 6	dem. 7	dem. 8	dem. 9	dem. 10
0.0294	0.0201	0.3676	0.0354	0.0713	0.0000	0.1687	0.1364	0.0454	0.0297	0.0960
0.1123	0.0159	0.0424	0.0836	0.0214	0.0017	0.3365	0.0273	0.0964	0.0786	0.1839
0.1396	0.3670	0.1396	0.0871	0.0662	0.0009	0.0842	0.0709	0.0270	0.0136	0.0040

Table 2: Three vectors showing sensitivity to the spatial distribution of strategies.

being that instead of checking only 100 different distributions for each initial frequency, we examined 10,000.) The initial conditions were as follows, for $i = 0, \ldots, 10$:

$$s^i_j = \begin{cases} 0.099 & \text{if } i \neq j \\ 0.01 & \text{otherwise} \end{cases}$$

when $i = 11$, $s^0_0 = s^0_1 = \cdots = s^0_{10}$. Unlike the previous experiments, here we found that each series *always* converged to the same polymorphism. We conclude, then, that although the initial spatial distribution of strategies can have an effect on the final state of convergence, the relatively small effect allows the use of Monte Carlo methods to determine the size of the basins of attraction for various polymorphisms.

3.2. Dependence upon the neighborhood and underlying dynamics. Table 3 summarizes the final convergent state of the world for several different combinations of neighborhoods and dynamics. The neighborhoods examined include the three most common in the literature (von Neumann, Moore (8), and Moore (24)), as well as the three nonstandard types displayed in figure 3. The row identified as "R(8)" used a different method: at the start of every generation, each player p randomly selects eight players from the world to serve as p's neighborhood for interaction and updating. Thus, the model of row R(8) serves as an intermediary between the fixed neighborhood structure of the other models and the replicator dynamics.

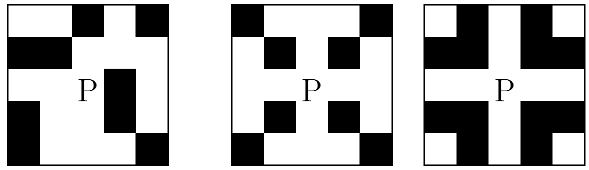

(a) Type 1 (b) Type 2 (c) Type 3

Figure 3: Three nonstandard neighborhoods used in table 3

		Polymorphism						
Nbhd	Dynamics	0–10	1–9	2–8	3–7	4–6	5	Other
VN	Mimic with proportion relative to success	0	0	0	0	29	9970	1
	Mimic best neighbor	0	0	0	0	26	9966	8
	Imitate best average strategy	0	0	0	0	13	9984	3
M(8)	Mimic with proportion relative to success	0	0	0	0	26	9973	1
	Mimic best neighbor	0	0	0	0	26	9908	66
	Imitate best average strategy	0	0	0	0	24	9970	6
M(24)	Mimic with proportion relative to success	0	0	0	8	110	9879	3
	Mimic best neighbor	0	0	0	21	220	9721	38
	Imitate best average strategy	0	0	0	0	62	9934	4
R(8)	Mimic with proportion relative to success	0	0	57	556	2418	6964	5
	Mimic best neighbor	0	0	54	550	2560	6833	3
	Imitate best average strategy	0	0	0	1	1523	8439	37
Type 1	Mimic with proportion relative to success Mimic best neighbor Imitate best average strategy	0	0	0	0	43	9868	89
Type 2	Mimic with proportion relative to success							
	Mimic best neighbor	0	0	0	3	62	9933	2
	Imitate best average strategy	0	0	0	0	28	9924	48
Type 3	Mimic with proportion relative to success							
	Mimic best neighbor	0	0	0	3	62	9933	2
	Imitate best average strategy	0	0	0	0	32	9965	3

Table 3: Convergence results based on neighborhood and dynamic.

In general, mean times to convergence are quite rapid. Models using the Moore (8) neighborhood usually converged within sixteen generations to fair division. This is a considerable improvement over the results of Skyrms[10], and a significant improvement of that of Kandori, Mailaith, and Rob[5], whose stochastically stable equilibrium only selects the equilibrium of fair division in the limit. As one might expect, the larger Moore (24) neighborhood leads to faster convergence times because the radius of influence of any given single player has increased.

		Population composition					
Nbhd	Dynamics	0–10	1–9	2–8	3–7	4–6	fair
	Mimic with proportion relative to success	–	–	–	–	25.4	23.9
VN	Mimic best neighbor	–	–	–	–	22.7	26.3
	Imitate best average strategy	–	–	–	–	32.2	23.9
	Mimic with proportion relative to success	–	–	–	–	17.9	16.4
M(8)	Mimic best neighbor	–	–	–	–	28.0	15.4
	Imitate best average strategy	–	–	–	–	17.2	14.6
	Mimic with proportion relative to success						
M(24)	Mimic best neighbor	–	–	–	32.3	22.8	12.8
	Imitate best average strategy	–	–	–	–	18.7	10.6
	Mimic with proportion relative to success	–	–	38.4	24.8	13.5	6.2
R(8)	Mimic best neighbor	–	–	15.5	13.9	8.5	4.5
	Imitate best average strategy	–	–	–	28.0	16.1	5.37
	Mimic with proportion relative to success						
Type 1	Mimic best neighbor	–	–	–	24.0	10.7	12.69
	Imitate best average strategy						
	Mimic with proportion relative to success						
Type 2	Mimic best neighbor	–	–	–	–	16.3	15.3
	Imitate best average strategy						
	Mimic with proportion relative to success						
Type 3	Mimic best neighbor						
	Imitate best average strategy						

Table 4: Mean convergence times

Figures 4, 5, and 6 illustrate the evolutionary path followed by worlds using three different neighborhoods. In these three figures, the initial conditions set all strategies equally likely and had players update their strategies using imitate the best neighbor dynamics. In the first two worlds, the strategy of fair division emerges from the initial random conditions even without any means to globally coordinate such an outcome. The third figure illustrates the effect of a degenerate (one-person) neighborhood in which all players use only their northern neighbor for interaction and updating.[7]

[7]Empirical tests suggest that the minimal neighborhood allowing fair division to successfully emerge in almost all initial conditions is the von Neumann neighborhood. The next

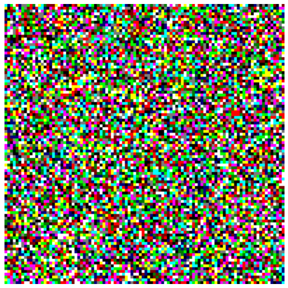

(a) Gen. 0

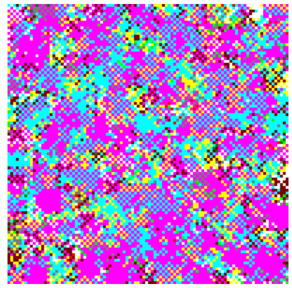

(b) Gen. 2

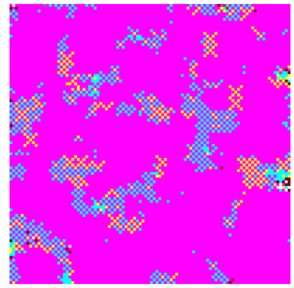

(c) Gen. 4

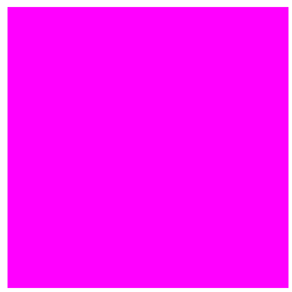

(d) Gen. 10

Figure 4: Evolution under neighborhoods of type 1

likely candidate, a neighborhood containing only the N, SE, and SW neighbors, often leaves small sections following the 4–6 polymorphism along the perimeter. This effect would not appear in models with periodic boundary conditions.

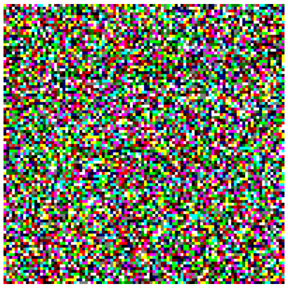

(a) Gen. 0

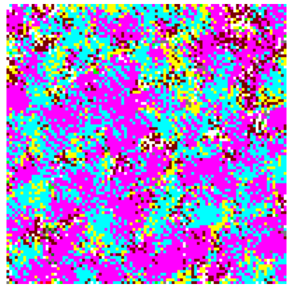

(b) Gen. 2

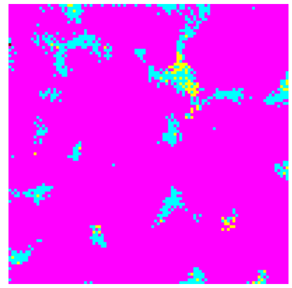

(c) Gen. 4

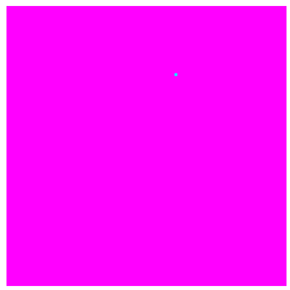

(d) Gen. 6

Figure 5: Evolution under type 2 neighborhoods

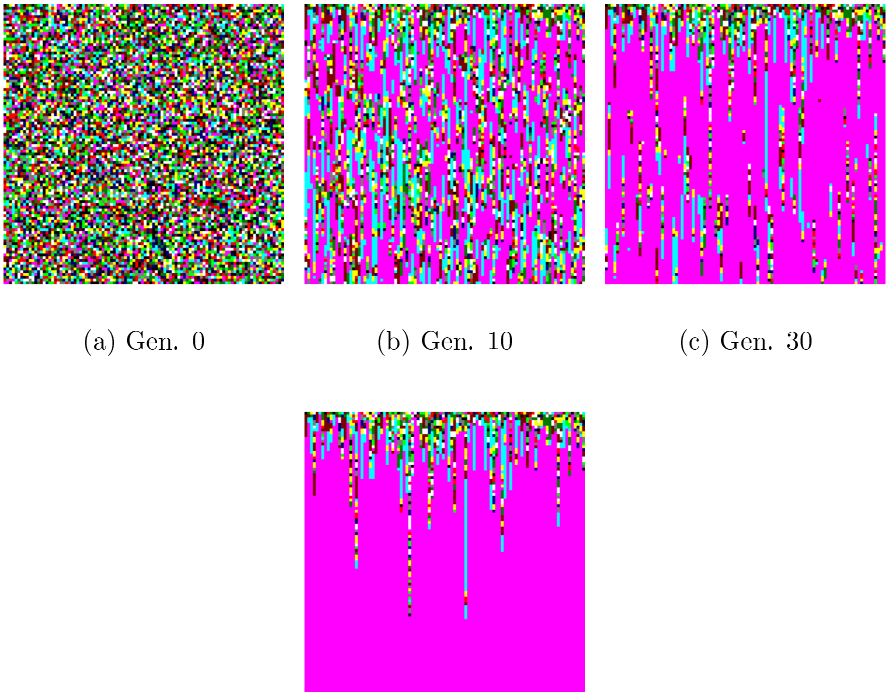

(a) Gen. 0 (b) Gen. 10 (c) Gen. 30

(d) Gen. 100

Figure 6: Evolution under degenerate (one person) neighborhoods

3.3. Dependence upon cake size. In *Evolution of the Social Contract*, Skyrms reported an interesting relationship between granularity of the good and the distribution of the resulting polymorphism. If we assume that players divide a cake consisting of ten slices, we find that fair division takes over the population roughly 62% of the time with some percentage of the population falling into one of the 1–9 or 2–8 polymorphic traps. However, increasing the total number of pieces the cake is sliced into leads to an increase in the total number of populations that will evolve into something "near" fair division. In particular, Skyrms found that a cake divided into 200 pieces went to fair division $\pm$ 3 pieces approximately 94.1% of the time; all trials went to fair division $\pm$ 11 pieces.

Since most populations evolving under spatial constraints lead to a pure state of fair division already, one natural question inverts the one considered by Skyrms: how *coarse* can we slice the cake while stilll getting fair division? Table 5 lists the results as the cake size varies from ten to two pieces, for each of the three dynamics considered, under the Moore (8) neighborhood.

4. Discussion. This paper began with a summary of Nydegger and Owen's findings on bargaining behavior in human subjects, noting that the behavior of the subjects was—though not inconsistent with—not entirely explicable by traditional game theory. The overwhelming preference for people to prefer equal division under completely symmetric circumstances suggests that people acted in accordance with a norm which strongly regulates individual behavior. If so, the etiological question of *where* this norm came from arises. Although traditional moral theories purport to offer explanations of such norms these explanations do not usually mesh well with naturalistic methods in philosophy.

Cake size	Dynamics	Polymorphism 0–10	1–9	2–8	3–7	4–6	5	Other
10	Mimic best neighbor	0	0	0	0	2	998	0
	Imitate best average strategy	0	0	0	0	2	998	0
	Imitate using relative success	0	0	0	0	3	997	0
9	Mimic best neighbor	0	0	0	0	5[10]	0	995[11]
	Imitate best average strategy	0	0	0	0	1*	0	999[12]
	Imitate using relative success	0	0	0	0	17*	0	983[13]
8	Mimic best neighbor	0	0	0	0	999[1]	0	1[1]
	Imitate best average strategy	0	0	0	0	998*	0	2[3]
	Imitate using relative success	0	0	0	0	1000*	0	0
7	Mimic best neighbor	0	0	0	3	0	0	997[14]
	Imitate best average strategy	0	0	0	3	0	0	997[15]
	Imitate using relative success	0	0	0	4	0	0	996[16]
6	Mimic best neighbor	0	0	0	998*	0	0	2[4]
	Imitate best average strategy	0	0	0	1000*	0	0	0
	Imitate using relative success	0	0	0	995*	0	0	5[5]
5	Mimic best neighbor	0	0	1	0	0	0	999[17]
	Imitate best average strategy	0	0	1	0	0	0	999[18]
	Imitate using relative success	0	0	2	0	0	0	998[19]
4	Mimic best neighbor	0	0	1000*	0	0	0	0
	Imitate best average strategy	0	0	999*	0	0	0	1[6]
	Imitate using relative success	0	0	997*	0	0	0	3[7]
3	Mimic best neighbor	0	1	0	0	0	0	999[20]
	Imitate best average strategy	0	0	0	0	0	0	1000[21]
	Imitate using relative success	0	0	0	0	0	0	1000[22]
2	Mimic best neighbor	0	997*	0	0	0	0	3[8]
	Imitate best average strategy	0	1000*	0	0	0	0	0
	Imitate using relative success	0	998*	0	0	0	0	2[9]

[1]Of these, 973 were pure states of demand 4. [2]A 3–5 polymorphism, $(3609, 6391)$. [3]Two 3–5 polymorphisms: $(5112, 4888)$, $(5196, 4804)$. [4]Two 2–4 polymorphisms: $(3736, 6264)$, $(3813, 6187)$. [5]Five 2–4 polymorphisms: $(3273, 6727)$, $(3484, 6516)$, $(3380, 6620)$, $(3476, 6524)$, $(3589, 6411)$. [6]A 1–3 polymorphism, $(2563, 7437)$. [7]Three 1–3 polymorphisms: $(2147, 7853)$, $(2233, 7767)$, $(2135, 7865)$. [8]In all three worlds, the strategy of demand 1 went extinct early on, leaving the population in an unstable equilibrium of $(1, 0, 9728, 34, 34, 36, 1, 90, 28, 0, 48)$, $(22, 0, 5682, 129, 571, 884, 556, 430, 960, 128, 638)$, and $(69, 0, 5771, 1646, 321, 86, 187, 626, 844, 115, 335)$. [9]Both worlds contain unstable equilibrium in which all strategies are present: $(4, 7, 5241, 212, 90, 195, 280, 495, 2387, 572, 517)$ and $(2, 11, 2440, 523, 646, 988, 702, 1831, 105, 1040, 1712)$. [10] Four of these states contained only demand 4. [11] One 3–6 polymorphism, one 3–5 polymorphism, with the rest being 4–5 polymorphisms. [12] Three 3–6 polymorphisms, the rest 4–5 polymorphisms. [13] Three 3–6 polymorphisms, the rest 4–5 polymorphisms. [14] One 2–5 polymorphism, one 2–4–5 polymorphism, the rest 3–4 polymorphisms. [15] Five 2–5 polymorphisms, the rest 3–4 polymorphisms. [16] Two 2–5 polymorphisms, the rest 3–4 polymorphisms. [17] Two 1–3–4 polymorphisms, the rest 2–3 polymorphisms. [18] One 1–3–4 polymorphism, the rest 2–3 polymorphisms. [19] All 2–3 polymorphisms. [20] One world containing the unstable equilibrium $(15, 3, 108, 4220, 254, 146, 525, 855, 2335, 1520, 19)$, the rest 1–2 polymorphisms. [21] One world containing the unstable equilibrium $(0, 0, 1074, 1070, 1011, 1111, 1125, 1083, 1110, 1175, 1241)$, the rest 1–2 polymorphisms. [22] Three unstable equilbriums of the following form: $(6, 0, 8764, 9, 10, 806, 60, 95, 116, 53, 81)$, $(86, 0, 2958, 1357, 648, 2611, 263, 1159, 618, 81, 219)$, and $(10, 10, 129, 1650, 4478, 4, 289, 253, 964, 1080, 1133)$, the rest 1–2 polymorphisms. *All states contain only the strategy making the lowest demand of the pair.

Table 5: Convergence results for a shrinking cake

Skyrms[10] explores an alternative approach to answering the etiological question using methods from evolutionary game theory. However, he uses the replicator dynamics as his evolutionary model, which makes assumptions of dubious validity for human populations (especially for human populations of the size when our moral norms were first forming). Moreover, Skyrms' models do not predict that a population of players engaged in the Nash bargaining game will always (or almost always) converge to fair division: the basin of attraction for fair division comprises only about 62% of the population. Skyrms found that the sizes of the basin of attraction for fair division could be significantly improved by introducing a small degree of correlation into the population, but justifying this correlation requires complicating the underlying story.

Many of these concerns with Skyrms' original model disappear if we consider a spatialized version of the Nash bargaining game. This model drops the assumption of the replicator dynamics that the population is essentially infinite and that any two players are equally likely to interact. In this setting, randomly initialized populations converge almost always to fair division. These convergence results persist even if we shrink the size of the cake. Although randomly initialized populations do not always converge to fair division, introducing a small amount of mutation virtually guarantees that a population will converge to fair division within a reasonable amount of time.[8] Figure 7 illustrates how a pure 4–6 polymorphism may be taken over by fair division in the presence of a little mutation.

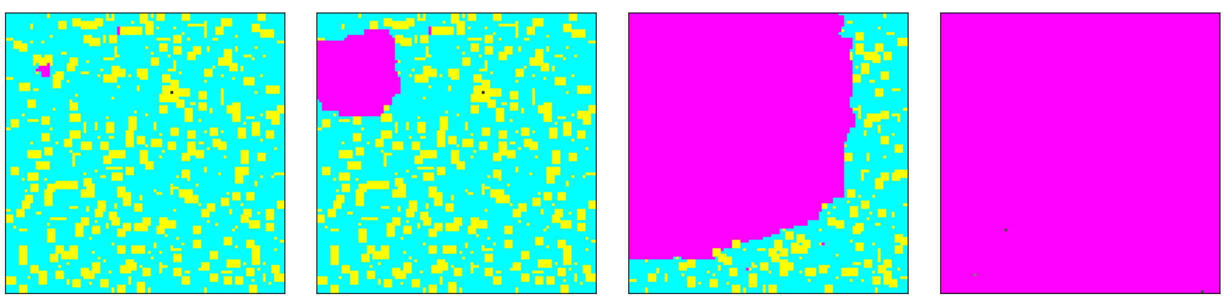

Figure 7: Emergence of fair division out of a 4–6 polymorphism due to mutation

[8]One, of course, needs to adapt the definition of "converge" accordingly to allow a population to converge to a state of fair division even when ϵ of the players follow another strategy.

The amount of time required to move a population out of a polymorphic pitfall to a state where everyone (or almost everyone) follows fair division depends on the frequency of mutations μ. Inspection of figure 7 will reveal that the key step in the emergence of fair division is the introduction of the demand 5 strategy into a site surrounded by sufficiently many compatible strategies. If μ is large, we will not have to wait very long for such a mutation to occur. If μ is small (or if there are not many sites following compatible strategies), then one will obviously have to wait longer. In comparison with the time required for the model of Kandori, Mailaith, and Rob to return to fair division if it gets trapped in a polymorphic pitfall, the wait in this case seems hardly significant: the run portrayed in figure 7, which had $\mu = .0001$, took just over one hundred generations for the critical mutation to occur.

Appendix A contains a proof of the claim that, in any competition between fair division and any other polymorphism, where the region occupied by fair division is of sufficient size and players update using mimic best neighbor dynamics, the population will converge to a state of fair division. Does this proof show that human populations will always converge to a state where people make equal demands in completely symmetric circumstances? Clearly not; spatial models, like the replicator dynamics, abstract away many important features of real communities. However, spatial models may be a closer approximation to the actual situation than the replicator dynamics. If so, the dynamics of the spatial Nash bargaining game moves us closer to a naturalistic explanation of the norm of fair division.

A. Analysis. Let $\mathbb{S} = \{0, 1, \ldots, 10\}$ be the set of strategies. Then $\Pi \subset \mathbb{S}$ is a *polymorphic pair* iff $\Pi = \{\, 5 \,\}$ or $\Pi = \{\, s, c \,\}$ where $s + c = g$ and $c < s$ (where g is the total amount of good available). (We often refer to s as the "greedy" strategy and c as the "modest" strategy.)

Let P be an infinite population of players with the population structure of an infinite square lattice. (We will often denote a given player $p \in P$, located at (x, y), by $p_{x,y}$.) A *strategy distribution for the population P at time t* is a function $\xi_t : P \to \mathbb{S}$. A *frontier competition* is strategy distribution $f : P \to \mathbb{S}$ such that there exist two polymorphic pairs Π_1 and Π_2 and some integer i such that for every $p_{j,k} \in P$, if $j \leq i$ we have $f(p_{j,k}) \in \Pi_1$ and $f(p_{j,k}) \in \Pi_2$ otherwise.

Let ξ be a strategy distribution for the population P, where the population structure is that of a square lattice, and let f be a frontier competition.

We say that the *frontier competition f is embedded in the strategy distribution* ξ (or that ξ has an embedded frontier competition) if there exist integers i and j, $i \leq j$ such that for all $p_{k,l} \in P$:

1. If $k \leq i$ or $k \geq j$, $\xi(p_{k,l}) = f(p_{k,l})$;
2. If $i < k < j$ and $\xi(p_{k,l}) \in \Pi_2$, then $\xi(p_{k,l}) = f(p_{k,l})$.
3. Otherwise, $\xi(p_{k,l}) \in \Pi_1$.

(Notice that every frontier competition has an embedded frontier competition.) According to this definition, saying that the strategy distribution ξ has an embedded frontier competition f between fair division and the c.s.s. Π is equivalent to saying that ξ may be constructed from f by changing the strategy assignments of a set of players whose x-coordinates lie between i and j to demand 5.

Lemma 1. Let ξ_t be a strategy distribution for the population P at time t, and let f_t be a frontier competition between $\Pi = \{ 5 \}$ and the polymorphic pair $\Pi' = \{ s, c \}$ embedded in ξ_t. Let $i+1$ be the x-coordinate of the leftmost member of Π' in f_t. If the x-coordinate of the leftmost member of Π' in f_{t+1} is ℓ, then the x-coordinate of the leftmost member of Π' in ξ_{t+1} is at least ℓ.

(This lemma says that we can predict certain qualitative features of the evolution of the strategy distribution ξ_t by considering the evolution of the embedded frontier competition, if there is one.)

Proof. Since f_t is a frontier competition, player p_{ik} follows strategy 5. Let S be the score of player p_{ij} computed according to the strategy distribution ξ_t, and let S' be the score of player p_{ij} computed according to f_t. We claim that $S' \leq S$.

According to the equivalent formulation of the definition of a frontier competition, ξ_t differs from f_t only in the assignment of strategies to a set of players whose x-coordinates lie between i and j, for some integer j. If none of these players lie within the Moore neighborhood of p_{ik}, then $S' = S$. But suppose that some of these players lie within the Moore neighborhood of p_{ik}. Denote the set of these players by $\{ q_1, \ldots, q_n \}$. If $q \in \{ q_1, \ldots, q_n \}$ had a strategy compatible with p_{ik} in f_t, changing q's strategy to demand 5 still leaves q with a strategy compatible with p_{ik} and, consequently, will not affect the score of p_{ik}. However, if q's strategy was incompatible with p_{ik}, changing it to demand 5 will increase p_{ik}'s score by 5. Consequently, $S' \leq S$.

Now suppose that p_{ik} did not change her strategy from demand 5 to a strategy in the c.s.s. Π' at the end of generation t when we consider the evolution of the frontier competition f_t. We claim that p_{ik} will not change her strategy in the evolution of the strategy distribution ξ_t. Let $q \in \{ r : r \text{ is a neighbor of } p_{ik} \}$. We know that q's score was insufficient to make p_{ik} replace her strategy at the end of generation t when we consider the evolution of the frontier competition f_t. We claim that the score of q, when computed using the strategy distribution ξ_t, cannot be any higher than the score computed using the strategy distribution f_t, unless q's strategy is demand 5 (in which case, even if q earns a higher score under ξ_t than in f_t it does not matter as it will not cause p_{ik} to change strategies). One minor complication in this argument is that there is a possibility that q's strategy is different in ξ_t than in f_t. Thus, we need to consider cases on q's strategy.

Suppose that $\xi_t(q) \neq f_t(q)$. Given the definition of an embedded frontier competition, it follows that q's strategy must be demand 5. The claim follows trivially.

Suppose that $\xi_t(q) = f_t(q)$. Let $r \in \{ t : t \text{ is a neighbor of } q \}$. If $f_t(r) = \xi_t(r)$, then the interaction with r contributes the same amount to q's score regardless of whether we consider the evolution using strategy distribution f_t or ξ_t. So suppose that $f_t(r) \neq \xi_t(r)$. According to the definition of an embedded frontier competition, if $f_t(r) \neq \xi_t(r)$, it must be the case that $\xi_t(r) = 5$. If q has the strategy of demand c, the fact that r's strategy is different in the distribution ξ_t from what it was in the distribution f_t cannot affect q's score since r's new strategy (demand 5) is still compatible with r's strategy. If q has the strategy demand s, the interaction with r under the distribution ξ_t will result in a lower score than under the distribution f_t if $f_t(r) = c$; in the other two cases q's score under ξ_t will be the same as under f_t. If q's strategy is demand 5, then her score will *increase* if r is assigned an incompatible strategy in f_t and will remain the same if r is assigned a compatible strategy in f_t. Thus, q can only earn a higher score under ξ_t than under f_t if q's strategy is demand 5.

Now, we know that p_{ik}'s score under ξ_t is greater than or equal to p_{ik}'s score under f_t. Additionally, none of the neighbors of p_{ik} can earn a higher score under ξ_t than f_t unless they follow the strategy of demand 5. Given these conditions, then, it follows that p_{ik} will not switch strategies under ξ_t if p_{ik} did not switch strategies under f_t.

However, we can say more than this: let ℓ be the x-coordinate of the leftmost member of Π' under f_{t+1}. It follows that the leftmost member of Π'

under ξ_t must have an x-coordinate of at least ℓ. To see why, suppose not. I.e., suppose that the leftmost member of Π' under ξ_t has an x-coordinate less than or equal to ℓ. Let p denote this player. Notice that $f_{t+1}(p) = 5$ and $\xi_{t+1}(p) \neq 5$.

Although p followed the strategy of demand 5 at the end of generation t, it is not guaranteed that p followed that strategy at the beginning of generation t. So we need to consider cases. Suppose that $f_t(p) = 5$. The reader may check that, according to the previous argument, in this case p's score cannot be lower under ξ_t than under f_t. Since the only players who earn higher scores under ξ_t than under f_t are players who follow the strategy of demand 5, it is impossible for $\xi_{t+1}(p) \neq 5$ if $f_{t+1}(p) = 5$.

Now suppose that $f_t(p) \neq 5$. Since $f_{t+1}(p) = 5$, it must be the case that p adopts the strategy of demand 5 at the end of generation t from some other player q. However, q's score cannot be lower under ξ_t than under f_t. Since no player with a strategy belonging to the polymorphic pair Π' earns a higher score under ξ_t than under f_t, it is impossible for p to adopt the strategy of demand 5 at the end of generation t under f_t but not under ξ_t.

As the existence of such a player p is impossible, we conclude that the leftmost member of Π' under ξ_t must have an x-coordinate of at least ℓ. $\square$

Lemma 2. Let f_t be a frontier competition between Π_5 and $\Pi_{s,c}$. Then $\lim_{t\to\infty} f_t = \mathbf{5}$.

Proof. We establish this claim by first showing it to hold for several special cases, and then argue that the case of the full frontier competition cannot evolve differently from these special cases.

Case I: To begin with, consider the following simple frontier competition:

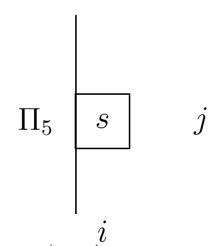

with a single s-strategist located at (i, j) with the rest of the region occupied by the $\Pi_{s,c}$ polymorphism consisting of c-strategists. We now consider cases on the type of polymorphism Π.

• $\Pi = \Pi_{4,6}$: Since p_{ij} demands 6 and has $4r(r+1) - r(2r+1) = 2r^2 + 3r$ compatible neighbors, p_{ij} will earn a score of $6(2r^2+3r)$. But p_{ij}, who follows the strategy of fair division and has $4r(r+1)-1 = 4r^2+4r-1$ compatible neighbors, will earn a score of $5(4r^2+4r-1)$. $\Pi_{4,6}$ will spread into the region occupied by Π_5 if and only if

$$\begin{aligned} 5(4r^2+4r-1) &< 6(2r^2+3r) \\ 20r^2+20r-5 &< 12r^2+18r \\ 8r^2 &< 5-2r \end{aligned}$$

which cannot be satisfied, given the range of the variable r. Thus the polymorphism $\Pi_{4,6}$ cannot spread into the region occupied by fair division.

• $\Pi = \Pi_{3,7}$: Proceeding as above, we obtain the inequality $6r^2 < 5+r$ which also cannot be satisfied.

• $\Pi = \Pi_{2,8}$: As above, we obtain the inequality $4r^2 < 5+4r$. Here, we see that the inequality holds when $r = 1$. In this case, we calculated the evolution of the frontier as follows (suppose we begin at generation t):

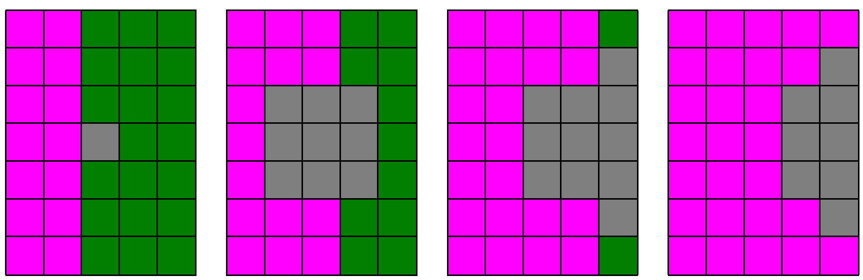

t $\quad$ $t+1$ $\quad$ $t+2$ $\quad$ $t+3$

By the end of generation $t+3$, the boundary between Π_5 and $\Pi_{2,8}$ has shifted two squares to the right, and the frontier competition between the two polymorphisms has disappeared. Notice, though, that we know how the evolution will proceed from this generation on: the 8-strategists proceed to the right, eliminating the less successful 2-strategists. At the same time, the 8-strategists at the boundary between Π_5 and $\Pi_{2,8}$ will adopt the strategy of fair division. (Surrounded only by 5-strategists and 8-strategists, these players earn scores of zero. Following the Imitate Best Neighbor dynamics, they adopt the strategy of the highest scoring player in their neighborhood. The general course of the evolution, as shown above, guarantees that the highest scoring player in the neighborhood of a boundary 8-strategist will be a player who divides fairly.) Thus, $\lim_{t\to\infty} f_t = \mathbf{5}$.

• $\Pi = \Pi_{1,9}$:   As above, we obtain the inequality $2r^2 < 5 + 7r$, holding for $r = 1, 2, 3$, and 4. As before, in generation $t+1$ the singular demand 9 neighborhood will expand to occupy every square in the Moore neighborhood of radius r around him (for the values of r specified.) However, none of the players following demand 9 in generation $t+1$ earn a score greater than $5(4r^2 + 4r)$ (the maximum possible score for a 5-strategist). Since all 5-strategists have at least one 5-strategist earning the maximum possible score, none of the 5-strategists will be replaced at the end of generation $t+1$, regardless of the value of the radius r. Furthermore, by the end of generation $t+3$, the boundary will have shifted at least $2r$ squares to the right. An argument similar to that for the case where $\Pi = \Pi_{2,8}$ applies, and so $\lim_{t\to\infty} f_t = \mathbf{5}$.

• $\Pi = \Pi_{0,10}$:   As above, we obtain the inequality $0 < 5+10r$, which holds for all r. Applying the argument for the case where $\Pi = \Pi_{1,9}$, *mutatis mutandis*, we conclude that $\lim_{t\to\infty} f_t = \mathbf{5}$.

Case II: Consider the frontier competition shown below, generalizing the frontier competition of case I:

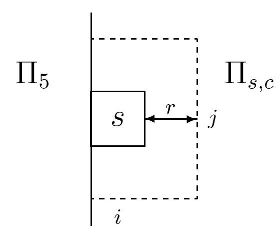

where within the Moore neighborhood of radius r about the player p_{ij} we position n s-strategists. We assume, at this point, that there are no other s-strategists outside of the neighborhood of p_{ij}.

In all cases, introducing additional s-strategists into the neighborhood of p_{ij} will *reduce* the score of player p_{ij}, since s-strategists are incompatible. For the polymorphisms $\Pi_{4,6}$ and $\Pi_{3,7}$, introducing s-strategists cannot affect the growth of the region occupied by fair division as these strategies could not, even before the introduction of additional s-strategists, successfully invade the region held by Π_5; for these polymorphisms, then, it is still true in this case that $\lim_{t\to\infty} f_t = \mathbf{5}$.

For the polymorphisms $\Pi_{2,8}$, $\Pi_{1,9}$, and $\Pi_{0,10}$, introducing additional s-strategists may affect the growth of the region occupied by fair division in several ways. The diagrams below show the possible effects for the Moore (8) neighborhood and the polymorphism $\Pi_{1,9}$ (the case for the polymorphism $\Pi_{2,8}$ will be similar). The diagram contains four two-column pairs in which the left column contains a schematic indicating the distribution of strategies in the generation prior to the one displayed in the right column. The strategies of fair division and demand 1 are omitted from the diagram as they may easily be inferred (all players left of the frontier demand 5, players to the right of the frontier demand 1 unless indicated otherwise):

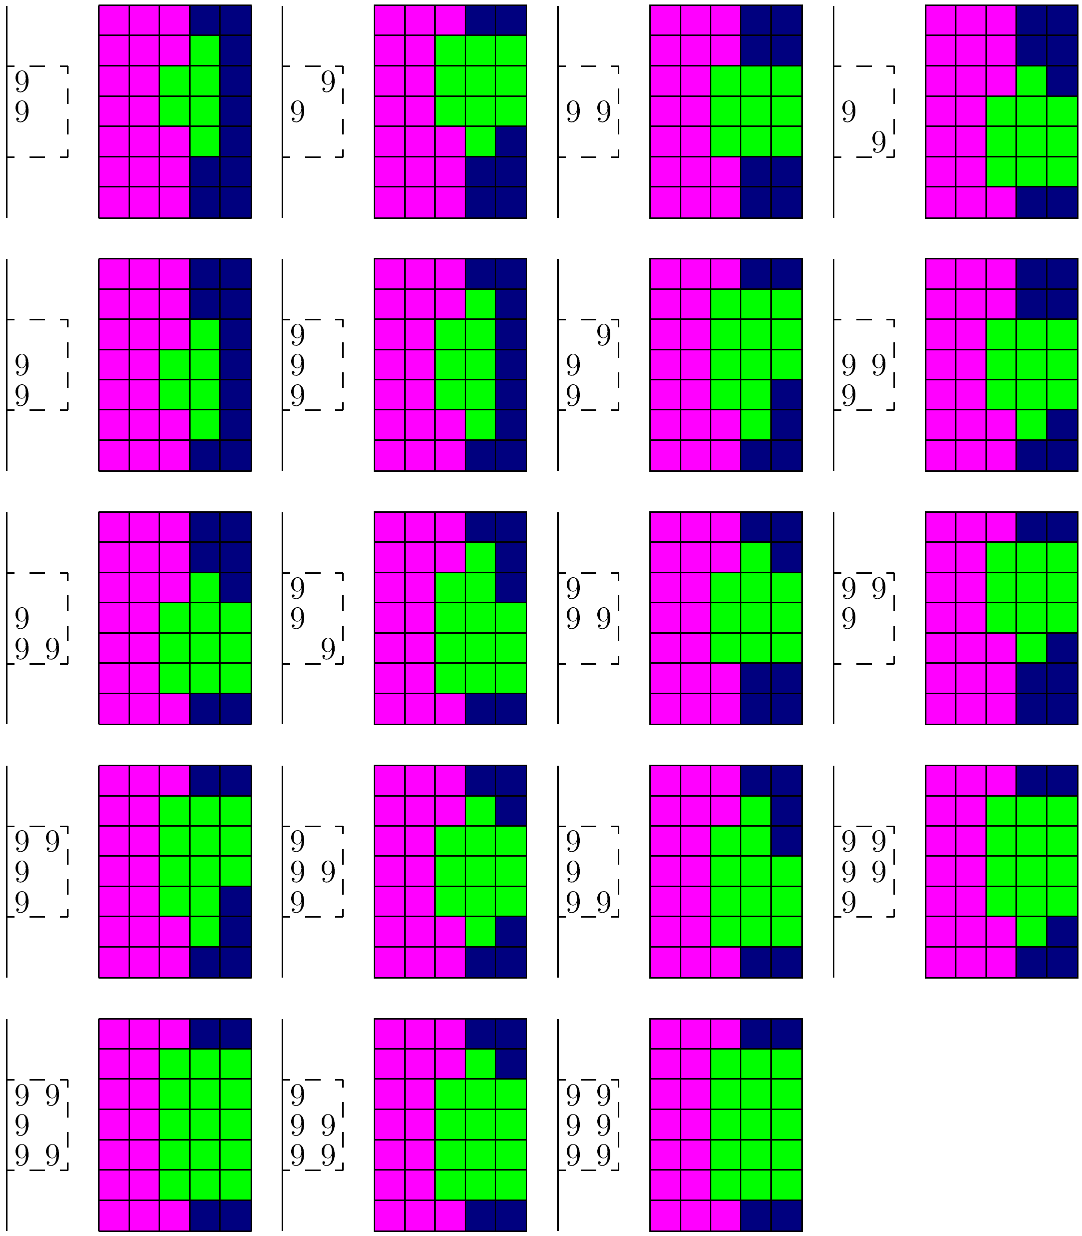

Notice that out of the nineteen possible configurations, not a one successfully invaded the region occupied by fair division. In addition, for each of the configurations above, the following properties hold:

1. Players following fair division located on the boundary between Π_5 and $\Pi_{1,9}$

will not be replaced in later generations since they all have in their neighborhood a 5-strategist who obtains the maximal possible score of $5(4r^2 + 4r)$.

2. All 9-strategists on the frontier, except the northernmost and southernmost will earn scores of 0 as they are isolated from any compatible strategy.

3. 1-strategists lying within the Moore neighborhood of a 9-strategist or fair divider will switch strategies to adopt either demand 5 or demand 9.

Thus we see that the demand s strategy will spread into the $\Pi_{s,c}$ region, turning c-strategists into s-strategists; s-strategists on the boundary between Π_5 and $\Pi_{s,c}$, cut off from compatible strategies, will become demand 5 strategists. Occasionally a s-strategist p with a score of 0 will remain an s-strategist since another s-strategist q in p's neighborhood will earn a score high enough to prevent p from switching strategies. However, in the next generation, q's success will have led the c-strategists around q to become s-strategists, preventing q from obtaining a high score in the next generation and, consequently, leading to p adopting the demand 5 strategy in the next generation. Thus, $\lim_{t\to\infty} = \mathbf{5}$.

Case III: Now consider the special case:

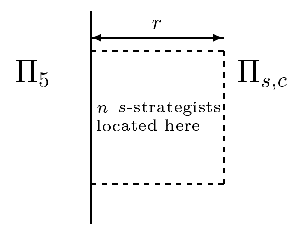

where s and c-strategists are distributed within r cells of the boundary between Π_5 and $\Pi_{s,c}$, with all other cells in the right-hand side of the world being c-strategists. This case lifts the require that we have an s-strategist located exactly on the boundary. We assume that no such s-strategist is so located (otherwise we have case II again). Notice that the spread of s-strategists into the region occupied by fair division, at the end of generation t, can be, at best, a subset of that occuring in case II.

The spread of fair division into the region occupied by $\Pi_{s,c}$ will occur as in case II, except that, since s-strategists are not located on the boundary, more c-strategists in the region occupied by the $\Pi_{s,c}$ will become s-strategists at the end of generation t. We note that properties (1)–(3), as listed in case II, still hold at the end of generation $t+1$. By an argument similar to that of case II, we conclude that $\lim_{t\to\infty} f_t = \mathbf{5}$.

Case IV: Now consider the special case:

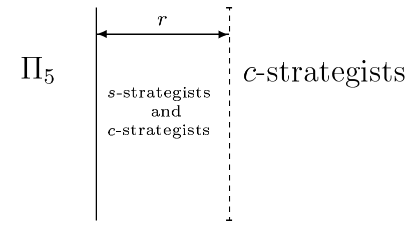

In this case, we allow for the possibility of having multiple regions of the sort identified in cases I, II, and III occuring simulataneously in the corridor of width r indicated in the diagram. Notice the addition of multiple cases does not significantly affect the argument given previously since the inclusion of more s-strategists can only reduce the maximum possible score of present s-strategists.

In addition, note that, whereas before, fair division would spread along the top and bottom of the region occupied by s-strategists, this may not happen in this case (at the end of generation t, those cells may be occupied by s-strategists). However, one may check to see that the only property of those cells that we used in all preceeding arguments was that the strategy employed by those cells wsa incompatible with demand s. Since the strategy of demand s is incompatible with itself, the same effect occurs here. Consequently, $\lim_{t\to\infty} f_t = \mathbf{5}$.

Case V: Now consider the general frontier competition:

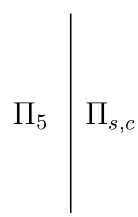

One can check that the argument given in case IV did not rely on the presence of only c-strategists in the half-infinite plane in the region occupied by $\Pi_{s,c}$. Furthermore, inclusion of s-strategists there will only serve to *aid* the spread of fair division, since adding additional s-strategists can only decrease the maximum possible score held by any particular s-strategists.

Therefore, $\lim_{t\to\infty} f_t = \mathbf{5}$. $\square$

Lemma 3. Let ξ_t be a strategy distribution for a population P. If ξ_t has an embedded frontier competition between $\mathbf{5}$ and $\Pi = \{\, s, c \,\}$, then $\lim_{t\to\infty} \xi_t = \mathbf{5}$.

Proof. Follows immediately from lemmas 1 and 2. $\square$

Now consider the case where we have a competition between the polymorphisms Π_5 and $\Pi_{s,c}$ spatially positioned on an infinite plane as follows:

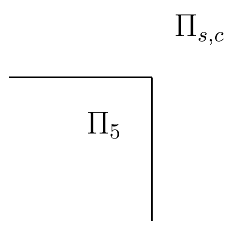

Although I do not included the details here, I claim that one can carry out an argument similar to the one given previously for the case of the frontier competition; namely, that regardless of the spatial positioning of s and c strategies in the region occupied by $\Pi_{s,c}$, the strategy of fair division will have expanded its territory by one square to the north and east within three generations. I also claim that a lemma similar to lemma 1 can be established for the case of the "corner competition" illustrated above.

Finally, notice that in the argument of Lemma 2 (and in the analogous argument(s) which can be given for the case of the corner competition), we did not need the assumption that the strategy of fair division occupied the entire half-infinite plane (or a quadrant of the plane). In both cases, we only needed the block of players following fair division to be large enough to survive the initial advance of the competing polymorphism. Consequently, if there exists a block of players of this size, the population will converge to a state of fair division.

References

[1] Richard Durrett and Simon Levin. The importance of being discrete (and spatial). *Theoretical Population Biology*, 46:363–394, 1994.

[2] Joshua A. Epstein. Zones of cooperation in demographic prisoner's dilemma. *Complexity*, 4(2):36–48, 1998.

[3] Rainer Hegselmann. Social dilemmas in lineland and flatland. In Liebrand and Messick, editors, *Frontiers in Social Dilemmas Research*, pages 337–361. Springer, 1996.

[4] Bernardo A. Huberman and Natalie S. Glance. Evolutionary games and computer simulations. *Proc. Natl. Acad. Sci.*, 90:7716–7718, August 1993.

[5] Michihiro Kandori, George J. Mailath, and Rafael Rob. Learning, mutation, and long run equilibria in games. *Econometrica*, 61(1):29–56, January 1993.

[6] Kristian Lindgren and Mats G. Nordahl. Evolutionary dynamics of spatial games. *Physica D*, 75:292–309, 1994.

[7] John F. Nash. The bargaining problem. *Econometrica*, 18:155–162, 1950.

[8] Martin A. Nowak and Robert M. May. Evolutionary games and spatial chaos. *Nature*, 359:826–829, October 1992.

[9] Martin A. Nowak and Robert M. May. The spatial dilemmas of evolution. *International Journal of Bifurcation and Chaos*, 3(1):35–78, 1993.

[10] Brian Skyrms. *Evolution of the Social Contract*. Cambridge University Press, 1996.

[11] Peter D. Taylor and Leo B. Yonker. Evolutionary stable strategies and game dynamics. *Mathematical Biosciences*, 40:145–156, 1978.

[12] H. Peyton Young. The evolution of conventions. *Econometrica*, 61(1):57–84, January 1993.

[13] P. Young. An evolutionary model of bargaining. *Journal of Economic Theory*, 59:145–168, 1993.