The Cultural Red King Effect

Title : The Cultural Red King Effect
Author : Cailin O'Connor
Address : Department of Logic and Philosophy of Science, University of California, Irvine, 3151 Social Science Plaza A, Irvine, California 92697, USA
Electronic address : cailino@uci.edu
Acknowledgements : Many thanks to Simon Huttegger, Edouardo Gallo, Brian Skyrms, James Weatherall, Peyton Young, and Justin Bruner for comments and suggestions. I would also like to thank the UCI Social Dynamics seminar for wonderful feedback.

Abstract

Why do minority groups tend to be discriminated against when it comes to situations of bargaining and resource division? In this paper, I explore an explanation for this disadvantage that appeals solely to the dynamics of social interaction between minority and majority groups—the cultural Red King effect. As I show, in agent-based models of bargaining between groups, the minority group will tend to get less as a direct result of the fact that they frequently interact with majority group members, while majority group members meet them only rarely. I will also show that this effect unifies previous results on the impacts of institutional memory on bargaining between groups.

Key Words

Evolutionary Game Theory, Bargaining, Discrimination, Inequity, Red King Effect, Game Theory, Social Dynamics

The Cultural Red King Effect

Cailin O'Connor

December 8, 2016

Why do minority groups tend to be discriminated against when it comes to situations of bargaining and resource division? In this paper, I explore an explanation for this disadvantage that appeals solely to the dynamics of social interaction between minority and majority groups—the cultural Red King effect. As I show, in agent-based models of bargaining between groups, the minority group will tend to get less as a direct result of the fact that they frequently interact with majority group members, while majority group members meet them only rarely. I will also show that this effect unifies previous results on the impacts of institutional memory on bargaining between groups.

1 Introduction

According to the Red Queen hypothesis in biology, fast evolving species have an advantage over slow evolving ones. Quick adaptation means that they can avoid predation and parasitism, and gain an advantage in mutualistic interactions.[1] In contrast, using tools from evolutionary game theory, Bergstrom and Lachmann [2003] first described what they call the Red King effect in biology. As they argue, under some conditions mutualistic species can actually gain an advantage by evolving more slowly.

Perhaps the easiest way to give an inuitive explanation of this effect is by appeal to an analogous rational choice situation described by Schelling [1980]. Suppose that you and an opponent are playing chicken. You drive towards each other, each hoping the other will swerve first. Neither party wants to be the chicken, but both parties want to avoid a collision even more. One way to win this game is to visibly toss your steering wheel out the window. This means that you are unable to change direction while your opponent maintains their ability to swiftly adapt. We can predict that any reasonable opponent will swerve. Likewise, in some mutualistic settings, a fast adapting species will 'swerve', or adopt strategies that yield ultimately lower fitness because while the other species is not changing they temporarily benefit by behaving in an accommodating way.[2]

[1] The term 'Red Queen' comes from Lewis Carroll's *Through the Looking Glass* where the Red Queen tells Alice, "Now, *here*, you see, it takes all the running you can do, to keep in the same place" [Carroll, 1917].

[2] Thanks to Jean-Paul Carvalho for this analogy.

Given this suggestive example, one might ask: can we see the Red King effect in a cultural setting? Are there interactions where actors who adapt, learn, or culturally evolve more quickly than others ultimately end up disadvantaged as a result? The chicken scenario above illustrates a case where something like this happens in a one-shot interaction between individuals. A more interesting question is: can we observe this effect when conventions of behavior emerge between social groups? Might the cultural Red King effect lead to conventions and norms that disadvantage those of certain races, classes, genders, religious groups, etc.?

In order for this to happen, it must be the case that members of one group learn or culturally evolve more quickly than members of another group. On the face of it, this does not necessarily sound like a reasonable condition for human populations. Bruner [working paper], however, observes that when minority and majority groups interact, minority members meet majority members much more frequently than the reverse as a result of their differential prevalence in the population. This asymmetry in the learning environment of the two types of actors means that minority types will learn to interact with majority members at a faster rate. While Bruner outlines a version of the cultural Red King effect in models of infinite populations evolving according to the replicator dynamics, I will here bring the results into a more explicitly cultural context by replicating them in agent based models using assumptions of bounded rationality. This is an important step since, as I point out, two-population replicator dynamics models, when applied in a cultural situation, rest on the assumption that actors only learn from those in their own social group. As I show here, this is not a necessary assumption to produce the cultural Red King effect. In this new context, as I argue, it is easy to see that the well cited phenomenon of risk aversion strengthens the cultural Red King effect considerably. I will also argue that these results are part of a unified phenomenon where many influences on the learning, or adaptation rate of cultural actors—including differential social network structure, or differing institutional memory—can reproduce the cultural Red King effect.

The results of these models may give insight into why minority groups tend to be disadvantaged by norms and conventions of bargaining in many societies. In particular, they suggest that very bare bones assumptions about interacting populations can give rise to a situation that persistently disadvantages minority groups by dint of their minority status alone. This is not to suggest that thicker explanations of the emergence of inequity and discrimination against minority populations are not important, but that even without psychological phenomena like out-group bias, or stereotype threat, minorities can be at greater risk of disadvantage.

In section 2, I will describe previous relevant results on the Red King and the cultural Red King. In section 3, I will motivate and describe the models analyzed in this paper, and present results from these models. In section 4, I will connect the results here to those outlined by Young [1993b] and Gallo [2014] on the impact of institutional memory and network connectivity on bargaining. And lastly, in section 5, I will conclude.

2 Previous Results

2.1 The Red King

Bergstrom and Lachmann [2003] characterize the strategic scenarios where the Red King effect may arise as ones in which, "multiple Nash equilibria exist, but different players have different preferences over the set of equilibria" (594). The set of games exhibiting these features correspond roughly with what Schelling [1980] called 'mixed motive' games. In these games, actors share some level of common interest in that they do well to choose complementary strategies. Their interests conflict, though, over which equilibria are preferred. Why are these sorts of games the ones of interest? A Red King effect occurs when a speed differential between evolving populations makes it more likely that the evolutionary dynamics carry those populations to an outcome that advantages the slow population. To observe this, it must be the case that there are outcomes that would be preferable for each population.

Figure 1 shows a two player, two strategy version of such a game.[3] The two equilibria are the strategies where one actor chooses A and the other B. Both players prefer the outcome where they are the one to take strategy B, generating the conflict of interest just described, but both players also prefer the equilibria to the off-equilibrium outcome where they receive nothing. As Bergstrom and Lachmann [2003] point out, when $x < 1$, this is a coordination game—one where both actors improve payoff by coordinating their behavior.[4] When $x = 1$, this is a version of the Nash demand game (described in the next paragraph). When $1 < x < 2$, this is a hawk-dove game, where at equilibrium, dove (A) players prefer to meet a dove rather than a hawk, meaning that the interaction is not purely mutualistic.

		Player 2	
		A	B
Player 1	A	x,x	1,2
	B	2,1	0,0

Figure 1: Payoff matrix for a two player, two strategy mixed motive game.

Another set of games which have multiple equilibria and differing preferences over them are versions of the Nash demand game with more demands. In the Nash demand

[3]Bergstrom and Lachmann [2003], Bruner [working paper] use the mini-game approach (see Sigmund et al. [2001]) of investigating small, tractable games that capture the strategic scenario of interest. Throughout this paper, I will do the same.

[4]The version I present is actually often termed an 'anti-coordination' game since actors must take opposite strategies to succeed. However, they still succeed by coordinating their action, so I find this terminology unhelpful. O'Connor [2015] distinguishes between correlative coordination games, where actors must correlate actions to succeed, and complementary coordination games, where actors need to take complementary actions. I use a complementary coordination game here, rather than a correlative one as in Bergstrom and Lachmann [2003], since this makes clearer the connection to the work of Axtell et al. [2000] on bargaining, which I will later elaborate on.

game, two actors divide a resource, and their strategies consist in demands for some portion of it. If the two demands are compatible, each actor receives their demand. If they are incompatible in that they exceed the resource, the actors receive a poor payoff sometimes called the 'disagreement point'. Figure 2 shows a general version of the Nash demand game with three demands labeled Low, Med, and High. For simplicity sake, I assume actors are dividing a resource of size 10. The value of the Med demand is always 5, representing an even split of resources. The Low and High demands can take any values such that $L + H = 10$, and $L < 5 < H$, such as, for example, 3 and 7, or 1 and 9.

		Player 2		
		Low	**Med**	**High**
Player 1	**Low**	L,L	L,5	L,H
	Med	5,L	5,5	0,0
	High	H,L	0,0	0,0

Figure 2: Payoff matrix for a two player, three strategy Nash demand game.

Suppose that two co-evolving species engage in at interaction that is well modeled by some version of the game in figure 1 where the first species always takes the role of player 1 and the second of player 2. For many evolutionary dynamics, there are two possible stable rest points of this model—the outcome where the first species always plays A and the second B, or the outcome where the second species always plays A and the first B.

Let us assume that the populations in the model evolve according to the replicator dynamics.[5] When the two species evolve at the same rate, these two outcomes will always be equally likely. In other words, the two equilibria have equal sized basins of attraction under the dynamics. The phase diagram for this evolutionary population where $k = 1.5$ is shown in figure 3 (a).[6] The x-axis of the diagram shows possible population proportions for species 1, and the y-axis for species 2. The diagram shows, for each joint population state, the direction of change under the replicator dynamics. The two stable rest points are pictured as black dots in the top right, and lower left corners. The basins of attraction for these equilibria are the sets of initial population states that travel towards them. These are shown in white, for the lower left corner, and

[5]These are the most commonly used model of selection in evolutionary game theory. They assume that strategies that beat the population average payoff will become more prevalent, while those that are less successful will contract. The two-population replicator equations specify change in proportional representation, x_i (y_i), for each strategy, i, in a population with proportions $x = \{x_1, x_2...x_i...x_n\}$ ($y = \{y_1, y_2...y_i...y_n\}$). They are $\dot{x_i} = x_i(u_i(y) - \sum_{j=0}^{n} u_j(y)x_j)$. This equation can be read as stating that the rate of change of a particular strategy ($\dot{x_i}$) is equal to the current proportion of that strategy (x_i) multiplied by the difference between the payoff to that strategy given the state of the y population ($u_i(y)$) and the average payoff for the entire x population given the state of the y ($\sum_{j=0}^{n} u_j(y)x_j$)). Strategies for population y update according to analogous dynamics.

[6]These figures, and others like them in the paper, were generated using the program Dynamo [Sandholm et al., 2012].

shaded for the top right corner. The black line represents the separatrix between the basins of attraction.

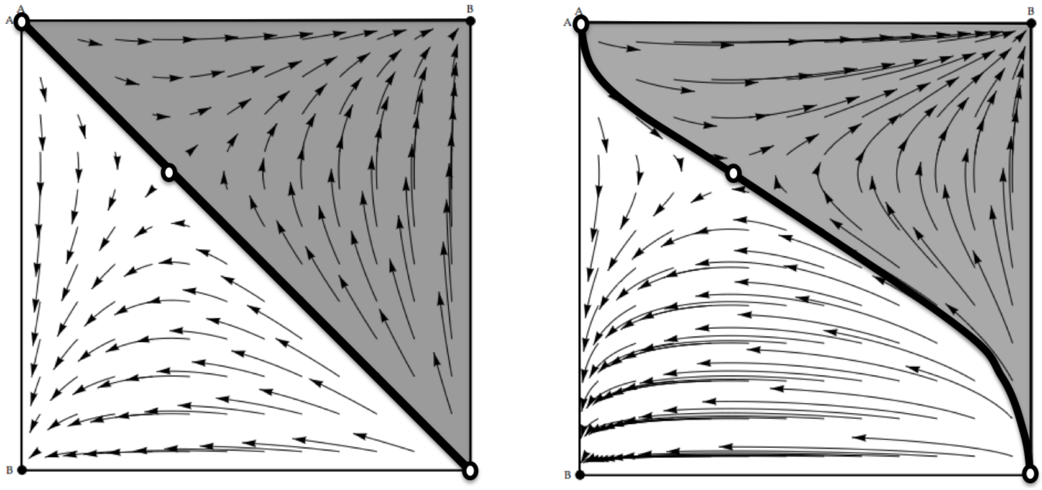

(a) $m = 1$ (b) $m = 3$

Figure 3: Phase diagrams for two populations playing a mixed motive game where one population evolves m times as quickly as the other. The basins of attraction for the two equilibria are shown in white and shaded for each figure. The separatrix is the dark line.

When one species evolves more quickly than the other, the shape of the basins of attraction changes. Figure 3 (b) shows the same phase diagram as (a), but where species 1 evolves at three times the rate of species 2.[7] As is evident, the separatrix in this picture is curved by the increased speed along the x dimension. This curvature means that now the top right equilibrium—where species 1, the fast evolving species, gets a payoff of 2 and species 2 a payoff of 1—has a smaller basin. In contrast, the equilibrium where species 2, the slow evolving species, gets a higher payoff, is now larger.

This sort of speed difference can also create an advantage for the fast evolving population (a Red Queen effect), depending on the underlying strategic scenario. The location of the central, unstable rest point in these diagrams, pictured as the central open dot, will determine whether or not speed increases or decreases the basin of attraction for the preferable outcome for the fast-evolving species. Subsequent authors to Bergstrom and Lachmann [2003] have outlined, in more detail and for further types of interactive scenario, how the speed of evolution can influence mutualistic interactions, though describing these results will be beyond the scope of this paper [Gao et al., 2015, Gokhale and Traulsen, 2012].

For games with more strategies, it will not be possible to show the phase diagram, but a similar effect occurs. For example, consider the Nash demand game where $L = 4$ so that the three possible demands are 4, 5, and 6. If two co-evolving species play this game, there are now three possible equilibria—where the two populations both demand

[7]This is done by adding a multiplier m to the replicator equations for species 1.

5 of the other, and two equilibria where one population demands 6 and the other 4. For this model, the slower evolving species will, again, be more likely to end up at an outcome where they always demand 6 and less likely to end up at an outcome where they demand 4. Figure 4 shows outcomes of simulations for this model where the multiplier m for the faster evolving population ranges from 1 to 10. Data are proportions of simulations arriving at the three equilibria for 10k runs.[8]

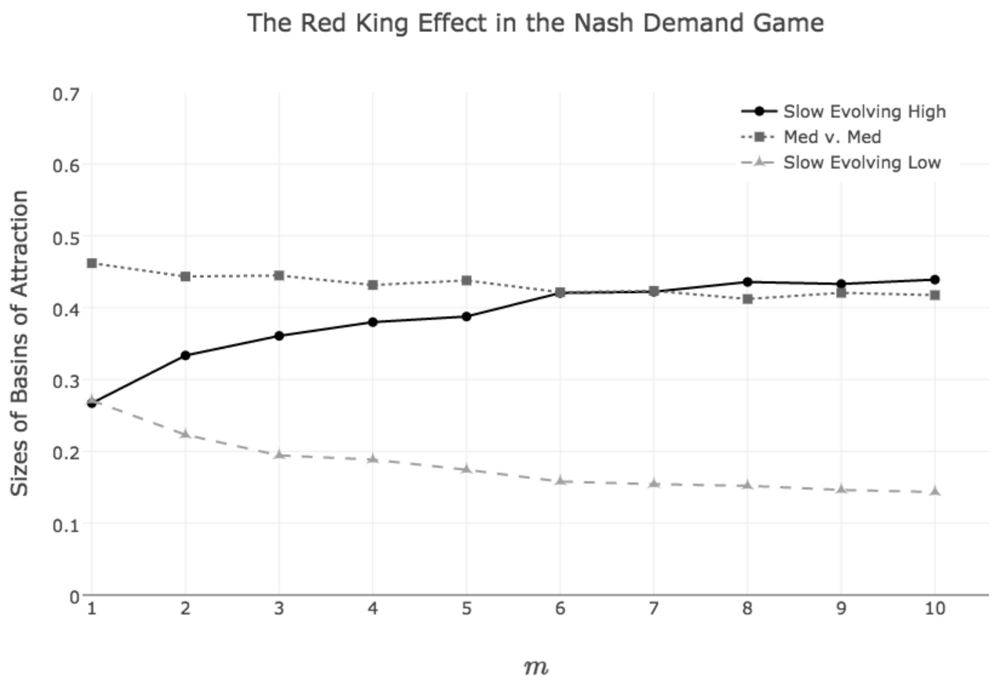

Figure 4: Basins of attraction for two populations playing the Nash demand game where one evolves at a rate m.

As is clear from the figure, as the speed of the faster evolving population increases, the proportion of simulations that go to the favorable outcome for the slower evolving population increases. Notice that this increase is bounded, though. Figure 3 shows why. The curvature of a separatrix is limited by the locations of central rest points, putting an upper limit on how much the Red King effect can impact outcomes in replicator dynamics models.

2.2 Minority Groups and the Cultural Red King

Bruner [working paper] was the first to observe that a version of the Red King effect

[8] These simulation results were generated using the discrete time replicator dynamics, where population proportions update at discrete time steps rather than continuously. The multiplier m then represents the number of replications that the faster population undergoes at each step.

can occur in cultural situations where a population is divided into types and where one type is in the minority. He considers a version of the Nash demand game where actors condition strategies based on the type of their opponent. In this version of the game, strategies consist of an ordered pair such as <Med, High> where the first entry specifies an actor's strategy against an in-group member and the second against an out-group member.[9]

Bruner [working paper] investigates the emergence of bargaining conventions in a model of this game evolved via the replicator dynamics. As with the game, in this model the replicator dynamics must be altered to account for the two-type structure of the population. Now two sub-populations meet both in-group and out-group members, and their evolving strategies specify behaviors in response to each type. The dynamics assume that the two sub-populations replicate separately, but based on payoffs from interactions with the entire population.

In this shifted context, again, a Red King effect occurs. This time, though, there is no multiplier responsible for the speed differential between the populations. Instead, size difference creates a differential in how significant each group is to the other's payoff. This means that the small group will always evolve more quickly. While Bruner [working paper] focuses on the Nash demand game, O'Connor and Bruner [2015] and O'Connor [2015] consider a wider class of games and confirm that results from Bergstrom and Lachmann [2003] and subsequent authors are replicated in this shifted cultural context.

3 The Model and Results

I will first describe and justify the class of explicitly cultural models that I consider in this paper, then specify the details of the particular models I employ and describe results.

3.1 Agent Based Models of Discriminatory Conventions

When employed in a cultural context, the replicator dynamics are commonly interpreted as change via imitation of successful group members. The assumption is that humans are more likely to adopt behaviors that are working for those in their social groups, and less likely to adopt unsuccessful behavior.[10] In order to apply this interpretation to a human group with types (as in Bruner [working paper], O'Connor and Bruner [2015], Bruner and O'Connor [2015]) one must assume that members of each type only imitate their in-group. In some domains, empirical observations support this assumption.[11] As I

[9]This is not necessary in the Bergstrom and Lachmann [2003] models because they consider only interactions between species, so actors only interact with one (out-group) type.

[10]This interpretation comes from the observation that the replicator dynamics are the mean field dynamics of explicit models of cultural imitation [Björnerstedt et al., 1994, Weibull, 1997, Schlag, 1998]. It has been observed in anthropology, for example, that humans do seem to imitate successful behaviors of group members [Richerson and Boyd, 2008].

[11]Wood and Eagly [2012], for example, extensively outline how children are explicitly directed to correctly gendered behavior, including own-gender imitation.

will show, though, it is not a necessary assumption for producing the cultural Red King effect.

To reproduce the effect in an explicitly cultural, agent-based model, without imitative learning, I borrow a framework introduced by Young [1993a], and employed by Young [1993b] to look at the emeregence of bargaining norms and by Axtell et al. [2000] to investigate the emergence of discriminatory norms between agents in different social groups. The version of the model here is closest to that employed by Axtell et al. [2000] who focus on simulation results of short term dynamics. Assume a finite population of N actors with two types of size n_1 and n_2 such that $n_1 + n_2 = N$. In each round of play, two agents in this model are chosen randomly to interact. Each agent has a memory of length m where she stores, for each of the two types, the last m stratagies she has encountered. Whenever an agent is chosen for interaction, she calculates her best response to her limited memory. We can think of this as a type of bounded rationality where the agent assumes that her recent interactions reflect the average proportions of strategies in the population. She then best responds to this assumed average. In addition, assume that actors sometimes err in that they choose a strategy that is not a best response with probability ϵ.

Axtell et al. [2000] investigate models of this sort where agents play three strategy Nash demand games like those described in section 2.1. They find that populations in this model will tend to be near two sorts of states when it comes to in-group interactions—either they all play Med, or else they are in a 'fractious' state where some portion demand High and some Low. When it comes to the out-group interactions, which are of greater interest for our purposes, there are three states populations remain near—the state where all actors play Med against out-group members, and the two states where one population always demands High and the other Low when meeting out-group members. In models without errors, these make up the absorbing states which models will head to, and then remain at. Axtell et al. [2000] identify states where there is an unequal split between the two groups as 'discriminatory' in the sense that individuals treat in- and out-group members differently to the detriment of one out-group. As I will show in the next section, shifting the proportions of the two types of agents in these models will shift the probabilities that each type is discriminated against.

3.2 Results

I begin by examining models like those just described with a subset of the following parameter values: population size, $N = 10, 20, 100$, the proportion of the larger type in the total population $n_1/N = .5, .6, .7, .8, .9$, the value of the Low demand $L = 1, 2, 3, 4, 4.5$, and memory length $m = 2, 5, 7, 10, 20$.[12]

[12] I make one small change from the implementation in Axtell et al. [2000]. They randomly initialize the memories of the agents at the start of simulation. For low values of L, this will mean that on average more agents will start by demanding High than Low, whereas for higher values of L, this will mean that more agents will start by demanding Low than High. These initial demands matter because we are looking at a case where one type learns more quickly on average as a result of its size. If the initial demands are skewed in one direction, the minority type will adapt to this, meaning they will be likely to

If $\epsilon > 0$ so that agents err, populations in these models always have positive probabilities of moving from any absorbing state to any other [Axtell et al., 2000]. This is because strings of errors can occur that shift the best responses of enough members of the population to drive it towards a new absorbing state. To analyze such a model, it makes sense to see, on average, how much time the model spends near each state. This gives a sense of the importance of such states, or of the likelihood that they arise. In practice, once simulations of the model with errors approach an absorbing state, a transition is highly unlikely (assuming a reasonably low error rate). For this reason, I start by looking at models where there are no errors ($\epsilon = 0$) and calculate the relative proportions of simulations that head to each absorbing state under various parameter settings. (The absorbing states, recall, are the ones where between-groups actors either demand Med, or one group always demands High and the other Low.) Once one of these absorbing states is reached, the population remains there.[13]

The key variation from the models explored by Axtell et al. [2000] is the alteration of the population proportions. From this point forward, I will refer to the proportion of the larger type, n_1/N, as p_1 for simplicity sake. Figure 5 shows results for 10k runs of simulation for the model where the low demand is $L = 4.5$, memory length, $m = 10$, population size, $N = 20, 100$, and where p_1 varies across the parameter space. All runs of this simulation ended up at one of the absorbing states described and the figure shows proportions of these. As is evident, as the size of the more prevalent type increases, three things happen. The proportion of outcomes where the two populations make Med demands of each other decreases slightly. The proportion of outcomes where the p_1 type demands High increases, and the proportion of outcomes where p_1 demands low decreases. In the larger population, this effect is more dramatic, though in both cases when $p_1 = .9$ it is essentially certain that the larger population will end up demanding High.

ultimately end up complementing whatever the majority type is doing at the beginning of simulation. To avoid this, I start the agents with no memories and determine their first strategies using random coin flips. Afterwards, they best respond to whatever memories they have.

[13] These are conventions in the sense outlined by Young [1993a].

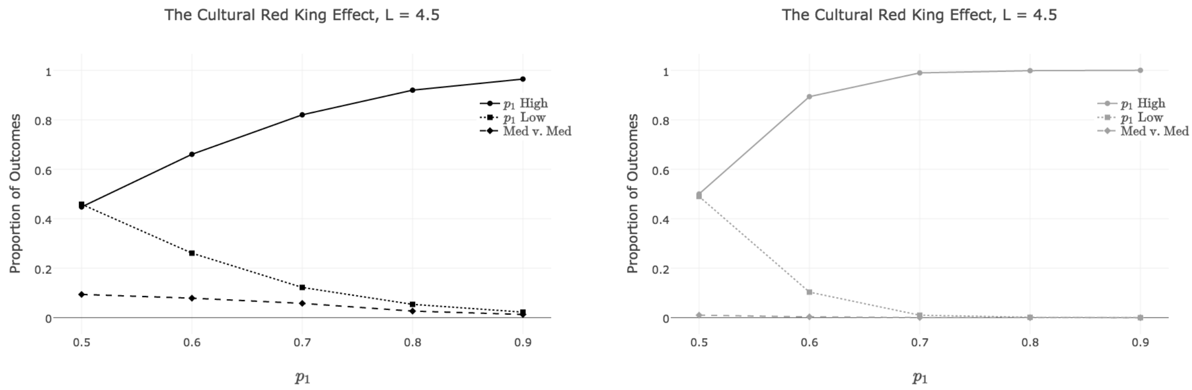

(a) $N = 20$

(b) $N = 100$

Figure 5: Proportions of outcomes for two types playing the Nash demand game where the proportion of the more prevalent type is p_1, $L = 4.5$, $m = 10$.

The cultural Red King effect here consists in the increasing likelihood that the larger population ends up 'discriminating' against the smaller population, and the decreasing likelihood that the reverse is true. For the games where $L = 3, 4, 4.5$, this effect is stable across all parameter values explored, though the strength of the effect varies, as does the likelihood that types end up playing the fair demand. In particular, the effect is stronger when L is higher.

For games where $L = 1, 2$, the effect reverses. As described in section 2.1, this reversal has to do with the location of interior rest points of the model. Intuitively, when L is higher, it is more beneficial for the smaller population, on average, to make Low demands against the larger one at the start of simulation. When L is lower, even though Low is a less risky strategy in that it always pays off, it is still better, on average, for the small population to make High demands at the start of simulation than to demand Low, meaning that they are more likely to end up demanding High as a result of their increased learning speed. This is a type of cultural Red Queen effect. However, there is a difference, in Nash demand game models, between the two effects. When L is lower, making the asymmetric demands more disparate, it is increasingly likely that the two types end up at the Med demand. This means that the proportion of cases where the Red Queen effect matters is much lower than the proportion of cases where the Red King does. Figure 6 shows simulation results from a model where $L = 2$ (results were very similar when $L = 1$). Data is shown for the populations where $N = 10, 20$ because when $N = 100$ the Med v. Med outcome always evolves. As is evident in the figure, the most likely outcome is always the Med v. Med demand. As p_1 increases, it becomes more likely that the large population demands Low, but this effect is much less dramatic than the Red King effect in otherwise similar models. Shortly, I will say a bit more about why the cultural Red King effect is more significant than the Red Queen.

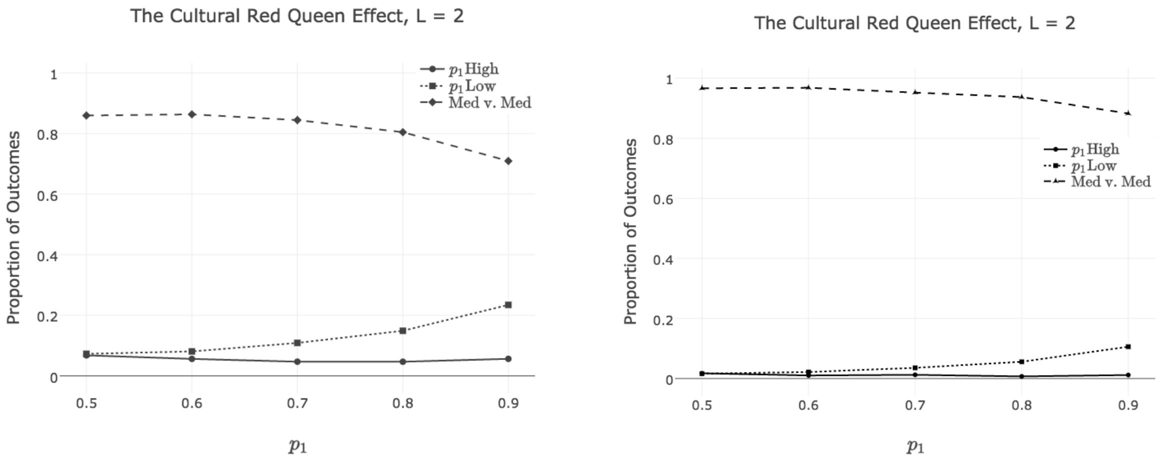

(a) $N = 10$

(b) $N = 20$

Figure 6: Proportions of outcomes for two types playing the Nash demand game where the proportion of the more prevalent type is p_1. $L = 2$, $m = 10$.

Let us also consider models where actors err with some positive probability. In these models, as described above, populations can move from one absorbing state to another (and will do so eventually given enough time).[14] We can analyze these models by considering not what proportion of runs end up at each absorbing state, but by considering the amount of time that the population tends to spend near each state. In particular, I will look here at strategies played by the actors over the entire course of simulation. I consider models with error rate $\epsilon = .1$, memory $m = 10$, population size $N = 10, 20, 100$, low demand $L = 2, 4.5$, and proportion of larger type $p_1 = .5, .6, .7, .8, .9$.

In models with an error rate, again, the cultural Red King and Red Queen effects can be observed. Figure 7 shows the proportions of rounds over which the two types of players take each strategy in simulations of the model. Figure 7 (a) shows these proportions for the game where $N = 20$, $m = 10$, and $L = 4.5$. As the figure shows, as the sizes of the groups become more disparate, the proportion of rounds where the larger types choose High gets larger, the proportion of rounds where they choose Low gets smaller. At the same time, the proportion of rounds where the small type choose High gets smaller, and the proportion where they choose Low gets Larger. Figure 7 (b) shows these proportions for the same model but where $L = 2$. In this case we observe a cultural Red Queen effect though, as in the model without error rates, we can see that it is less significant. For all values of p_1, populations spend most of their time making Med demands. As p_1 gets higher, there is a small increase in the larger type demanding Low and the small type demanding High.

[14]In practice this is very unlikely to occur during 10k runs of simulation for reasonably low error rates.

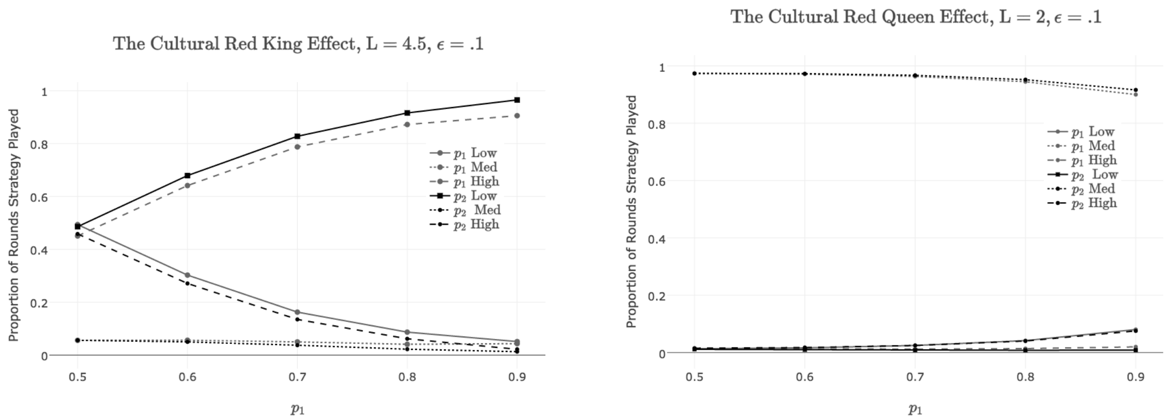

(a) $L = 4.5$

(b) $L = 2$

Figure 7: Proportions of rounds in which strategies are played for each population when $L = 4.5$, the cultural Red King effect, and $L = 2$, the cultural Red Queen effect, $m = 10$, $N = 20$ with $\epsilon = .1$.

We might also be interested in moving away from mini-games to a strategic scenario where actors have the ability to more finely partition a good that they divide. Suppose that actors play a Nash demand game with nine possible demands, from 10% to 90% of the good. Let us again suppose that the error rate, $\epsilon = 0$. In this case, the larger number of demands means that, on average, at the beginning of simulation actors do better to make higher demands. As a result, the smaller type more quickly moves towards these higher demands and we see a Red Queen effect, if a relatively small one. Figure 8 shows results from these simulations. I represent the proportions of outcomes in this figure using a bar graph instead of a line graph to make results more clear because there are now nine possible absorbing states—all those where every member of the two groups makes compatible demands such as 1 vs. 9 or 4 vs. 6. As is evident from the figure as p_1 increases, the outcomes where the p_1 population makes smaller demands expand slightly.

The results of this model, and of smaller games where the Red Queen effect is observed, might lead one to worry that the Red King effect is highly parameter sensitive, and so not of serious concern for real-world populations. As I will show now, though, a small, realistic change to these models strongly strengthens the robustness of the cultural Red King. In the models discussed to this point in the paper, I have assumed that the best response by actors in the Nash demand game involves an expected payoff calculation where the amount of resource they receive directly corresponds to their experienced payoff. In similar models, though, Young [1993b] and Gallo [2014] employ a common assumption from economics about individual preferences over goods. In particular, they assume that actors are risk averse, meaning that their experienced utility decreases over each unit of a good they receive. When this assumption is incorporated in models like those presented thus far, it strengthens the Red King effect because the relative im-

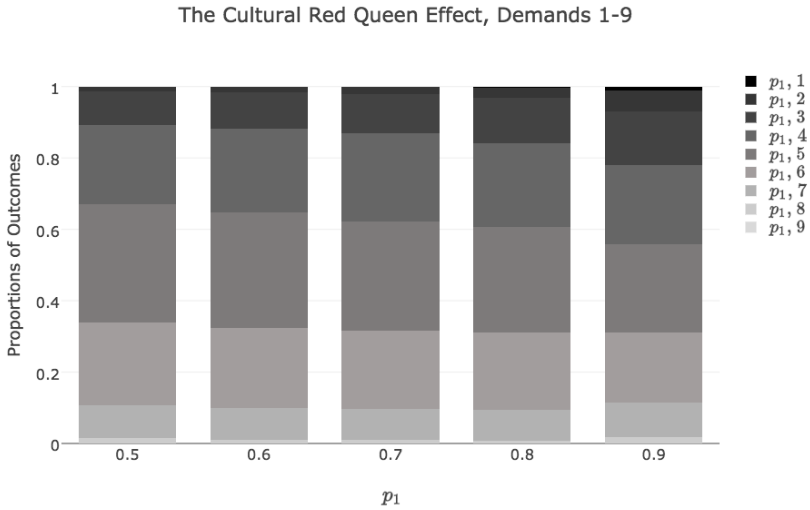

Figure 8: Proportions of outcomes for Nash demand game with demands 1-9, $N = 10$, $m = 10$, $\epsilon = 0$.

portance of earning small, certain amounts of resource outweighs the chance of getting more. This shifts actors' best responses. In such models, initial best responses to the other group tend to involve making low demands, meaning that the minority population tends to move towards these small demands.

Consider the model with nine possible demands, but where actors have risk averse preferences.[15] Figure 9 shows simulation results for this model. As is evident, the addition of risk aversion shifts the effect to a (stronger) cultural Red King.

To drive this point about risk aversion home, let us reconsider the mini-game from the beginning of this section. Suppose that actors play this game, but have risk averse utilities. Figure 10 shows outcomes for two versions of this game. (a) shows outcomes when $L = 4.5$. Without risk aversion, this model displays a cultural Red King (figure 5 (a)). With risk aversion, this is still the case but the effect is arguably somewhat stronger because for this model when the two types are equally proportional about a third of simulations go to the Med v. Med outcome. As p_1 increases it becomes certain that the larger population comes to demand High. (b) shows a game where $L = 2$. Without risk aversion this model has a cultural Red Queen effect (figure 6 (a)). With risk aversion, this becomes a Red King where now as p_1 increases the probability that fair demands emerge decreases and the probability that the larger group demands high increases. Under the particular utility curve employed here, adding risk aversion means that the Red King effect occurs for all values of L except for $L = 1$, which has a nearly

[15] I incorporate this into the model by using a utility function $u(x) = 3ln(x + 1)$. This is a somewhat arbitrary function chosen because it respects the 0 payoff point, is concave, and is monotonically decreasing. Other risk averse utility curves will have a similar effect.

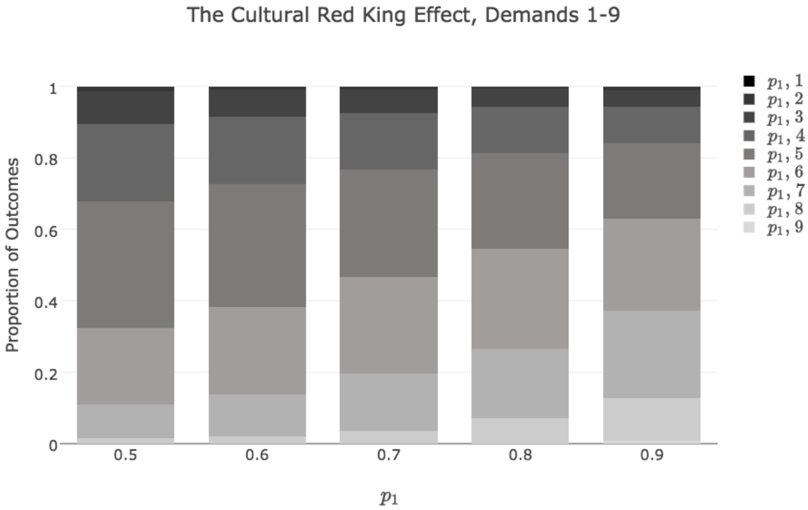

Figure 9: Basins of attraction for Nash demand game with demands 1-9, $N = 10$, $m = 10$, $\epsilon = 0$. Players are risk averse.

indetectable Red Queen.

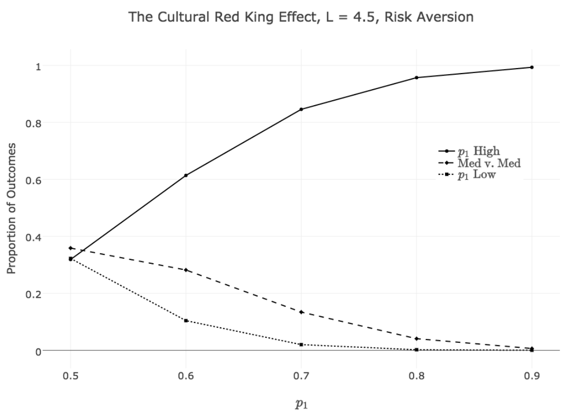

(a) $L = 4.5$

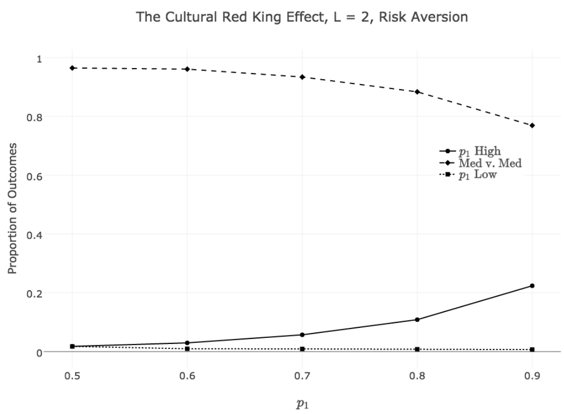

(b) $L = 2$

Figure 10: Basins of attraction for two types playing the Nash demand game where the proportion of the more prevalent type is p_1, $m = 10$, $N = 20$. Players are risk averse.

There are a few things to note about how the results presented in this section differ from those in replicator dynamics models. First, the strength of the cultural Red King effect is not bounded in the same way. In figures 5, 7, and 10 we see parameter values where it becomes nearly certain that the larger group demands High. Second, we can coherently incorporate the assumption of risk aversion. In a model of cultural change via imitation of successful group members, the observable success of others does the work to spread successful variants, leaving no place to include risk averse tendencies.

4 Institutional Memory and the Cultural Red King

In the models discussed thus far, the cultural Red King effect occurs because of a difference in the reactivity of the minority and majority groups. This is a particularly germane sort of case to study because in many societies minority groups tend to be disadvantaged in scenarios of bargaining, and this effect may help explain why. Minority status, though, is not the only way to generate a cultural Red King effect. There are other situations that will tend to make one social group more reactive than another. Suppose, for example, that members of one social group have a longer institutional memory—that they have ways of remembering more past interactions, or ways of sharing these past interactions with each other.

Young [1993b] investigates a set of models embodying such an assumption. The models are much like those explored to this point in the paper with slightly different assumptions about how memory works. Young further assumes, though, that members of the groups may have different memory lengths. This difference changes the reactivity of particular individuals. Those with longer memories will change strategies less readily, while those with short memories will be prone to change based on short strings of interaction. Young analyzes this model for stochastically stable equilibria (SSE). A SSE is defined in an evolving population with some sort of stochasticity—in this case an error rate for the individuals interacting. As mentioned above, such a population will spend time at each absorbing state. As the error rate approaches zero, though, the amount of time the population spends at just one equilibrium will approach 1. This equilibrium is the SSE.

In particular, Young focuses on a version of the game where the set of demands may be much more fine grained than those here. He proves that in these models, under the assumption that each group has homogenous memory lengths, the shorter the memory length, the worse the SSE is for that group. This is a type of cultural Red King because the reactivity of the two populations is what does the work in determining which side gets less in bargaining. As he says, "agents with more information are less likely to respond to mistakes by the other side, so they are steadier" (156).[16]

Gallo [2014] considers a similar set-up where two groups of networked actors bargain with each other. Instead of accessing finite memories to choose their strategies, these actors best respond to a sample of the recent demands that their network neighbors have encountered. Gallo shows that the level of network connectivity influences the SSE between groups in the same way memory length does for Young. A group with less

[16]In particular, Young shows that as the partition of demands for the actors gets finer and finer, the SSE approachs a version of the Nash bargaining solution weighted by memory length. He also considers heterogenous groups and finds that the actor with the shortest memory, i.e., the 'least steady' actor, in each group determines the SSE. Whichever group has the least steady actor of all is expected to be disadvantaged as a result.

connected members will get a lower demand at the SSE than a group with connected members. Again, connection influences the reactivity of individuals. Those with many neighbors see a large sample of demands and so are less likely to shift strategies because of random sampling effects. As he says, "Thanks to this informational advantage, they are less likely to respond to mistakes by the other side, and they are therefore able to maintain an advantageous bargaining position" (14).[17] Gallo backs up his theoretical findings with an experiment where two networked groups of actors play a Nash demand game. He finds that, on average, groups with fewer connections do less well.

We thus observe another type of situation in which two social groups may have an asymmetry in reactivity which leads to disadvantage for the more reactive type. The sort of difference outlined in this section could result when members of one group are admitted to a social club that facilitates networking and information transfer and members of another group are not. Or the difference could arise when members of one group have better access to education, allowing them to create records to transfer information about previous interactions, or access general records of that sort. Or when members of one group are generally tighter knit, meaning that they share information more readily.

5 Discussion

It is no secret that in many societies some social groups get more while others get less. This difference is due, at least in part, to social norms that dictate a bargaining advantage for some sorts of people and not others. Studies of used car sales find that women and people of color receive higher first offers than men and white people [Ayres and Siegelman, 1995]. Studies of hiring find that otherwise identical resumes with male or white sounding names are more likely than ones with female or black sounding names to garner job offers, and to be offered higher pay [Steinpreis et al., 1999, Bertrand and Mullainathan, 2003, Moss-Racusin et al., 2012]. In many societies, women work more hours and have fewer hours of free time than men [Eswaran, 2014]. Evidence suggests that these patterns are normative in the sense that people feel bargaining *should* advantage certain types of people (whether or not this feeling is implicit or explicit). For example, in experimental settings it has been found that women are punished more for making high demands on resources [Bowles et al., 2007].

But where do these inequitable norms come from? Why do they tend to disadvantage some groups and not others? The results presented here, and previous results from Bruner [working paper] and O'Connor and Bruner [2015] suggest that minority status alone may influence the dynamics of discriminatory norms. This effect is worth investigating further, in particular because it presents a very different sort of explanation for discrimination and oppression. This explanation does not appeal to the usual psychological phenomena associated with discrimination—implicit and explicit bias, stereotype

[17] In this set-up, as in the Young models, the least networked agent is the one who matters in determining the expected split between the two sides.

threat, and in-group favoritism. Instead, it depends on three rather bare bones assumptions. First, actors think of members of two social groups as separate, (i.e., they keep memories of these groups separately). Second, they choose actions that they expect to be best for themselves. And third, one group is in the minority. Actors in these models have no particular stereotypes about members of the two groups. Their learning and response rules are identical for everyone they interact with. They do not 'know' that there is a minority group, or even that discriminatory patterns of bargaining have emerged. Instead, simple patterns of choosing best responses lead to norms that can seriously disadvantage one group, and are especially bad for small groups.

References

Robert Axtell, Joshua M Epstein, and H Peyton Young. The emergence of classes in a multi-agent bargaining model, 2000.

Ian Ayres and Peter Siegelman. Race and gender discrimination in bargaining for a new car. *The American Economic Review*, pages 304–321, 1995.

Carl T Bergstrom and Michael Lachmann. The red king effect: when the slowest runner wins the coevolutionary race. *Proceedings of the National Academy of Sciences*, 100 (2):593–598, 2003.

Marianne Bertrand and Sendhil Mullainathan. Are emily and greg more employable than lakisha and jamal? a field experiment on labor market discrimination. Technical report, National Bureau of Economic Research, 2003.

Jonas Björnerstedt, Jörgen W Weibull, et al. Nash equilibrium and evolution by imitation. Technical report, 1994.

Hannah Riley Bowles, Linda Babcock, and Lei Lai. Social incentives for gender differences in the propensity to initiate negotiations: Sometimes it does hurt to ask. *Organizational Behavior and Human Decision Processes*, 103(1):84–103, 2007.

Justin P Bruner. Minority disadvantage in population games. working paper.

Justin P Bruner and Cailin O'Connor. Power, bargaining, and collaboration. In Thomas Boyer, Conor Mayo-Wilson, and Michael Weisberg, editors, *Scientific Collaboration and Collective Knowledge*. Oxford University Press, 2015.

Lewis Carroll. *Through the looking glass: And what Alice found there*. Rand, McNally, 1917.

Mukesh Eswaran. *Why gender matters in economics*. Princeton University Press, 2014.

Edoardo Gallo. Communication networks in markets. Technical report, Faculty of Economics, University of Cambridge, 2014.

Lei Gao, Yao-Tang Li, and Rui-Wu Wang. The shift between the red queen and the red king effects in mutualisms. *Scientific reports*, 5, 2015.

Chaitanya S Gokhale and Arne Traulsen. Mutualism and evolutionary multiplayer games: revisiting the red king. *Proceedings of the Royal Society of London B: Biological Sciences*, 279(1747):4611–4616, 2012.

Corinne A Moss-Racusin, John F Dovidio, Victoria L Brescoll, Mark J Graham, and Jo Handelsman. Science faculty's subtle gender biases favor male students. *Proceedings of the National Academy of Sciences*, 109(41):16474–16479, 2012.

Cailin O'Connor. The dynamics of gender and sex. Book Manuscript, 2015.

Cailin O'Connor and Justin P Bruner. Dynamics and diversity in epistemic communities. Working Paper, 2015.

Peter J Richerson and Robert Boyd. *Not by genes alone: How culture transformed human evolution*. University of Chicago Press, 2008.

W.H. Sandholm, E. Dokumaci, and F. Franchetti. Dynamo: Diagrams for evolutionary game theory. page http://www.ssc.wisc.edu/ whs/dynamo., 2012.

Thomas C Schelling. *The strategy of conflict*. Harvard university press, 1980.

Karl H Schlag. Why imitate, and if so, how?: A boundedly rational approach to multi-armed bandits. *Journal of Economic Theory*, 78(1):130–156, 1998.

Karl Sigmund, Christoph Hauert, and Martin A Nowak. Reward and punishment. *Proceedings of the National Academy of Sciences*, 98(19):10757–10762, 2001.

Rhea E Steinpreis, Katie A Anders, and Dawn Ritzke. The impact of gender on the review of the curricula vitae of job applicants and tenure candidates: A national empirical study. *Sex roles*, 41(7-8):509–528, 1999.

Jörgen W Weibull. *Evolutionary game theory*. MIT press, 1997.

Wendy Wood and Alice H Eagly. 2 biosocial construction of sex differences and similarities in behavior. *Advances in experimental social psychology*, 46(1):55–123, 2012.

H Peyton Young. The evolution of conventions. *Econometrica: Journal of the Econometric Society*, pages 57–84, 1993a.

H Peyton Young. An evolutionary model of bargaining. *Journal of Economic Theory*, 59(1):145–168, 1993b.