

Title : The Cultural Red King Effect

Published in : The Journal of Mathematical Sociology

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, and Peyton Young for comments and suggestions. Thank the UCI Social Dynamics seminar for wonderful feedback. And special thanks to Justin Bruner for his insights and 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 [Bruner, 2017]. 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. This effect is strengthened by certain psychological phenomenon—risk aversion and in-group preference—is robust on network models, and is strengthened in cases where pre-existing norms are discriminatory. I will also discuss how 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 Version of Record of this manuscript has been published and is available in The Journal of Mathematical Sociology, 2017, 10.1080/0022250X.2017.1335723

The Cultural Red King Effect

June 1, 2017

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 [Bruner, 2017]. 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. This effect is strengthened by certain psychological phenomenon—risk aversion and in-group preference—is robust on network models, and is strengthened in cases where pre-existing norms are discriminatory. I will also discuss how 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 intuitive 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

¹The 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].

‘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

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 common condition in human populations. Bruner [2017], however, makes the important observation 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. Using infinite population replicator dynamics models, Bruner, and subsequent authors [O’Connor and Bruner, 2017, O’Connor et al., 2017, O’Connor, 2017], show that although actors in two groups can be completely identical in terms of preferences, skills, and all other factors, the simple fact of minority membership can lead to disadvantage with respect to resource division as a result of these learning speed asymmetries. The goal of this paper will be to outline the generality and robustness of this result and, in particular, to ask: how likely is it that this cultural version of the Red King effect in fact impacts the emergence of real-world bargaining norms and conventions?

I will proceed along several avenues. First, I will replicate these results in a more explicitly cultural context via agent based modeling. In doing so, I instantiate the causal variables responsible for the effect in very different ways, and thus demonstrate a sort of ‘throw-everything-at-it’ robustness to the results. These agent based models are also important because, as I point out, they allow an exploration into the ways certain psychological phenomenon—risk aversion and in-group preference—interact with and strengthen the cultural Red King effect. Furthermore, they admit explicit representations of community structure. As I will outline, something like the cultural Red King occurs, and is in some senses more robust, in network models [Rubin and O’Connor, 2017]. Lastly, I will point out that in cases where existing cultural norms entail discrimination against minority groups, the cultural Red King can increase the chance that these norms are transferred to new arenas. All together these results indicate that the cultural Red King effect should be taken seriously for its potential to impact real-world bargaining. There is more to say, though. As I will discuss, expanding insights from Bruner, these results are part of a unified phenomenon where other influences on the learning, or adaptation rate of cultural actors—including differential social network structure and

²Thanks to Jean-Paul Carvalho for this analogy.

differing institutional memory—can reproduce the cultural Red King.

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 stereotype threat, or confirmation bias, minorities can be at greater risk of disadvantage.

In section 2, I will describe previous relevant results on the Red King, Red Queen, and cultural versions of these effects. 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 discuss the connection between the results here and 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 equilibrium 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

³Bergstrom and Lachmann [2003], Bruner [2017], O’Connor and Bruner [2017] use the mini-game approach (see Sigmund et al. [2001]) of investigating small, tractable games that capture the strategic scenario of interest for similar purposes. Throughout this paper, I will do the same.

⁴The 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 [2017] distinguishes between correlative coordination games, where

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 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 an 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 outcomes—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. In other words, one species or the other will end up at a preferred outcome.

Let us assume that the populations in the model evolve according to the replicator

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.

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




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

⁵These 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 $x = \{x_1, x_2, \dots, x_n\}$ ($y = \{y_1, y_2, \dots, y_n\}$). 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.

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

1 evolves at three times the rate of species 2. This is done by adding a multiplier m to the replicator equations for species 1. 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—one where the two populations both demand 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.⁷

Whether a Red King or a Red Queen occurs in this bargaining model depends on the details of the game. For a three strategy game with $L < 3$, a Red Queen will occur instead. (Later, I will give an intuitive explanation for this switch.) Also, as is visible in figure 4, these effects will be bounded. 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 (or Red Queen) effect can impact outcomes in replicator dynamics models.

2.2 Minority Groups and the Cultural Red King

Bruner [2017] was the first to show that a version of the Red King/Red Queen effect 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

⁷These 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.

The Red King Effect in the Nash Demand Game




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

member.⁸

Bruner [2017] 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 cultural context, as mentioned, a Red King/Red Queen effect occurs. This time, though, there is no multiplier responsible for the speed differential between the populations. Instead, size difference creates an asymmetry in how significant each group is to the other's payoff. Minority members meet majority members commonly, but the reverse is not true. This means that the small group will always evolve more quickly. While Bruner [2017] focuses on the Nash demand game, O'Connor and Bruner [2017] and O'Connor [2017] consider some extensions to it and confirm that results from Bergstrom and Lachmann [2003] and subsequent authors are replicated in this shifted cultural

⁸This 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.

context.

The aim of the current paper is to ask: are these results robust across modeling choices? Are they general? Do we expect them to apply to real populations? As Bruner shows, the cultural Red King will also occur under the logit, Brown-von-Neuman-Nash and Smith dynamics as well as the replicator dynamics because for all three of these the speed with which a variant spreads is impacted by how successful it is compared to other strategies.⁹ This said, all the models previously considered have employed infinite populations and deterministic dynamics. In the next section, I introduce models aimed at expanding the scope of these results.

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

I borrow a well-studied framework introduced by Young [1993a], and employed by Young [1993b] to look at the emergence 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 strategies 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 their recent interactions reflect the average proportions of strategies in the population. They then best responds to this 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

⁹These three dynamics, as well as the replicator and best response dynamics, are described by Sandholm [2010] as some of the most important in evolutionary game theory. As Bruner [2017] outlines, the cultural Red King effect does *not* arise under myopic best response dynamics because the rate of change of a strategy is not directly impacted by its payoff, meaning that both groups adapt at the same rate.

population always demands High and the other Low when meeting out-group members. In models without errors, these make up the absorbing states which populations 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.

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 and O’Connor [2015], Bruner [2017], O’Connor and Bruner [2017]) one must assume that members of each type only imitate their in-group. In some domains, empirical observations support this assumption.¹¹ In others it is less applicable. One nice feature of the agent based models just introduced is that there is no cultural imitation—actors simply develop behaviors in response to their social environments.

Part of the goal here, as discussed, is to check the robustness of the cultural Red King effect. Using these agent based models to do so is a sort of ‘throw-everything-at-it’ approach to robustness analysis. They have finite, rather than infinite, populations, stochastic, rather than determinist, dynamics, and employ boundedly rational best response, rather than cultural imitation, as a change rule. In other words, they maintain the causal variables responsible for the Red King—in- and out- groups, a minority/majority split, a bargaining interaction, and learning—while instantiating quite different modeling assumptions. As we will see, they also allow for the inclusion of several realistic features of psychology—risk aversion and in-group preference—and for the addition of community structure to the model.

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

¹⁰This 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].

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

¹²I 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

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. Note that this proportion will always remain fixed over the course of simulation. 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, $N = 20$, 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. Results shown here focus on between-group interactions, as these are most relevant to understanding discrimination. 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. The cultural Red King effect here, as in previous models described, 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. As in previous models, we see these results because the minority and majority types are learning at different rates. On average minorities will meet majorities more frequently, meaning that they are updating their memories of that group at a higher rate.

For the games where $L = 3, 4, 4.5$, the cultural Red King 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 to a Red Queen (mimicking the Red King/Red Queen reversal in previous models). 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

demands are skewed in one direction, the minority type will adapt to this, meaning they will be likely to 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.

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




Figure 5: Proportions of outcomes for two types playing the Nash demand game where the proportion of the more prevalent type is p_1 , $L = 4.5$, $m = 10$, $N = 20$.

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.

There is a difference, though, 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 impact of the Red Queen effect, when it occurs, is less than the Red King. Figure 6 shows simulation results from a model where $L = 2$ (results were very similar when $L = 1$). 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.

As mentioned above, the versions of these models considered by Axtell et al. [2000] involve errors so that populations can move from one absorbing state to another (and will do so eventually given enough time). We can analyze models with errors 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. Results are very similar to those presented for models without error—for higher values of L the majority population spends more time playing High, and the minority Low, and for lower values of L the reverse is true.¹⁴

We might be interested in moving away from mini-games to a strategic scenario where

¹⁴I ran 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 $p_1 = .5, .6, .7, .8, .9$.




Figure 6: Proportions of outcomes for two types playing the Nash demand game where the proportion of the more prevalent type is p_1 , $L = 2$, $m = 10$, $N = 20$.

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 7 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, and this effect was stable over parameter values.

To summarize, the results described in previous replicator dynamics models are robust across the modeling changes made here. Both the cultural Red King and Red Queen are observed, and the trends occur in the way previous models would predict. One thing to note, though, is that in these agent based models the strength of the cultural Red King effect is not bounded as it is under the replicator dynamics. In figure 5, for example, there are parameter values where it becomes nearly certain that the larger group demands High. The overall effect of this difference is a strengthening of the cul-

The Cultural Red Queen Effect, Demands 1-9




Figure 7: Proportions of outcomes for Nash demand game with demands 1-9, $N = 10$, $m = 10$, $\epsilon = 0$.

tural Red King compared to previous models, but no comparable strengthening of the cultural Red Queen (because of the asymmetry in the effects noted above). In the next three sections, I will expand on this theme—the inclusion of two sorts of psychological effects, plus the inclusion of network structure in the model, plus the inclusion of existing discriminatory norms, all support the potential importance of the cultural Red King in real world populations.

3.2.1 Risk Aversion

As I will show now a small, realistic change to these models strongly strengthens the cultural Red King (and weakens the Red Queen). 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 they have concave utility functions. (Their added 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 importance 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 quickly towards these small demands.

Consider the model with nine possible demands, but where actors have risk averse

preferences.¹⁵ Figure 8 shows simulation results for this model. As is evident, the addition of risk aversion shifts the effect to a (stronger) cultural Red King. Now as p_1 increases, that population is more likely to end up getting more.

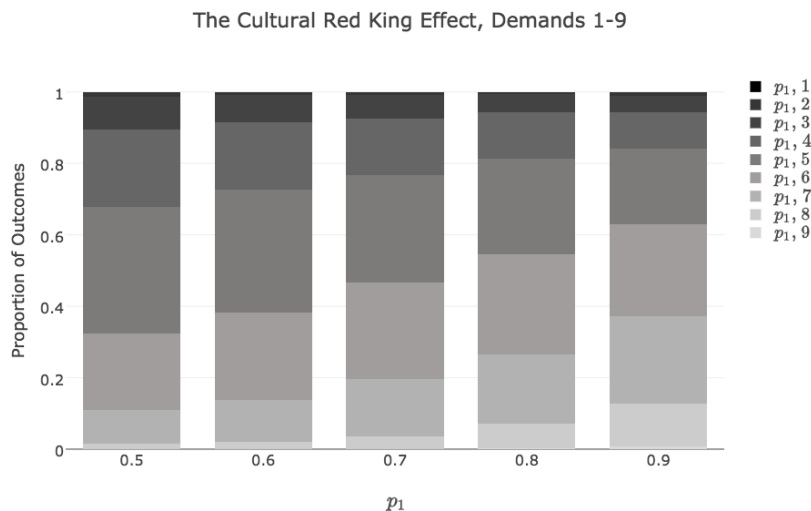


Figure 8: Proportions of outcomes for Nash demand game with demands 1-9, $N = 10$, $m = 10$, $\epsilon = 0$. Players are risk averse.

To drive this point about risk aversion home, let us reconsider the mini-games. Suppose that actors play this game, but have risk averse utilities. Figure 9 shows outcomes for versions of this game where (a) $L = 4.5$ and (b) $L = 2$. Without risk aversion, the model in (a) displays a cultural Red King (figure 5 (a)). With risk aversion, this is still the case but the effect is 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. Without risk aversion the model in (b) 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. Across parameter variations, risk aversion will generally lead to this sort of disadvantage for minority groups.

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

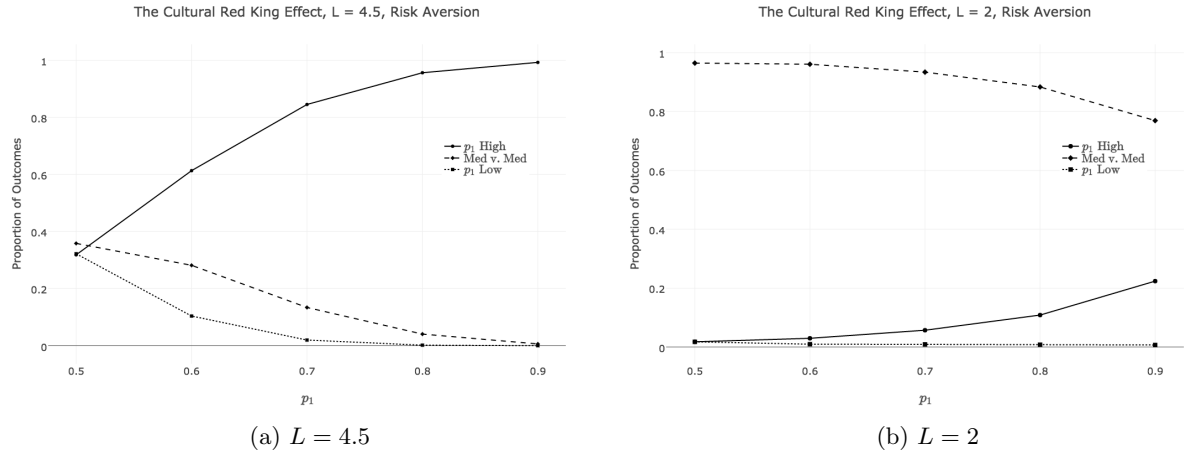


Figure 9: Proportions of outcomes for two types playing the Nash demand game where the proportion of the more prevalent type is p_1 , $m = 10$, $N = 20$. Players are risk averse.

Of course, one might point out that when actors are risk seeking, minority status will be more likely to advantage actors who develop bargaining norms. In such a case, minorities will be more likely to quickly learn high demands. A question we might ask, then, is: are minority groups typically risk averse or risk seeking? While the answer will surely depend on the details of the cultural situation, Guiso and Paiella [2008], in an empirical study, find that respondents with more assets are less risk averse than those with fewer, and that those who face fewer expected future risks also tend to be less risk averse. This implies that minority groups whose members are already disadvantaged in that they hold fewer resources and carry greater financial risk are expected to be further disadvantaged by the cultural Red King. Minority groups who are wealthier (and so less risk averse) should be less disadvantaged by the effect.

3.2.2 In-Group Preference

When it comes to in- and out-group dynamics, one phenomenon that has been widely observed in the social sciences is in-group preference. Experiments using minimal group paradigm have demonstrated that humans will show in-group preference even in cases where group membership is based on something as irrelevant as a coin flip [Tajfel, 1970, 1978, Haslam, 2004].

I incorporate in-group preference into the models explored here in a very minimal way. In particular, I suppose that actors who have no experience with their out-group demand High on their first meeting with probability α , and otherwise choose a random demand. Actors who have experience with out-group members best respond to their memories of these experiences in the standard way. This corresponds to a population where there is some slight suspicion of actors from an completely unknown social group, but where agents otherwise act in their boundedly rational best interest.

This inclusion of minimal in-group preference strengthens the cultural Red King

effect and weakens, or reverses, the cultural Red Queen. When out-group members tend to make high demands, a quick response increases the chances that a group will ultimately demand Low. Majority groups who are unresponsive to the high demands of their out-group are less likely to end up accommodating them. In figure 10, I show simulation results where the proportion of the majority group is $p_1 = .9$, and where α , the level of in-group preference, varies. In general, increasing α decreases the probability that fair bargaining norms emerge—this is because there are more inequitable demands made at the beginning of simulation, leading to more inequitable outcomes. In (a), where $L = 3$, we see the cultural Red King effect strengthen as α increases. The likelihood that the majority group demands High goes from 20% to about 85%. In (b), where $L = 2$, we see the cultural Red Queen shift to a Red King as α increases. When α is low, minority groups are more likely to end up advantaged for the reasons outlined. When it is High, their responsiveness to aggressive out-group members instead yields a disadvantage. Similar effects were observed for all other parameters tested.

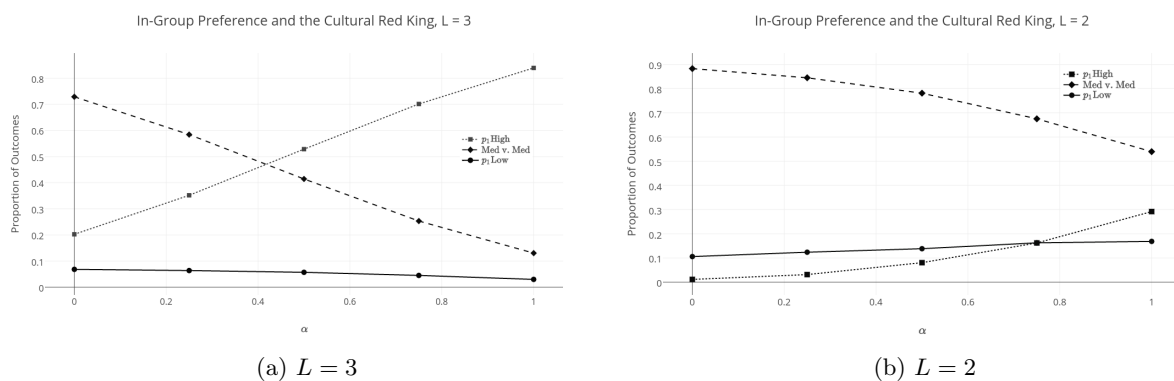


Figure 10: Proportions of outcomes for Nash demand game where $p_1 = .9$, $N = 20$, $m = 10$, $\epsilon = 0$. Players show in-group preference..

3.2.3 Network Models

The next set of results I describe are drawn from recent work with a co-author, Hannah Rubin. We examine the emergence of discriminatory behavior on a network where actors play the Nash demand game and best respond to the current behavior of their partners in each round with some small probability [Rubin and O’Connor, 2017].

We find that minority groups are more likely to end up demanding Low in these models when $L > 3$, just as in the models described above. The reason for this disadvantage is slightly different, however. For a network model with two groups, there will be some number of between group links, b . Given b , it must be the case on average that minority actors have more links to their out-group than majority members do because there are fewer of them. Because they are on average linked to more out-group members than the reverse, this creates a situation where their best response is more likely to be L and less likely to be H , creating an asymmetry that eventually leads to disadvantage.

In these models, we find no analogous Red Queen, because it is highly likely that fair demands emerge in the simulations where we might expect one.

3.2.4 Pre-existing Norms

One arena where the cultural Red King might plausibly arise is the workplace. In many workplaces, it was traditionally the case that almost all workers were from one social group. Consider, for example, the entrance of black people into traditionally white academic fields, or the entrance of women into law, which was traditionally practiced by men. This means that when in-group/out-group dynamics first emerged in these workplaces, the asymmetry in group size was extreme.

The models considered thus far assume that actors play randomly at the beginning of simulation (except when they show in-group preference). In the case of workplaces, though, actors are drawn from populations where norms and conventions of bargaining and status are already established. In particular, traditionally, minority members in many workplaces were part of groups already disadvantaged by discrimination. (Though in some cases, like nursing, the opposite was true.)

As with in-group preference, I incorporate existing norms into the model in a very minimal way. When actors meet an out-group member for the first time, with some probability β majority members demand High and minority members demand Low, and otherwise they make a random demand. This reflects an assumption that some set of the population has normative expectations that majority members should get more. Once actors have experience with out-group members, they behave in their own best interest by best responding to their memories.

What happens to the cultural Red King/Red Queen when actors already follow discriminatory norms? Again, the cultural Red King is strengthened, and the cultural Red Queen reversed. When minority actors tend to meet discriminatory out-group members, they quickly adapt by making lower demands. If they were less reactive, they might instead hold out until majority groups learned to behave fairly, or to accommodate. Figure 11 shows outcomes for simulations where $p_1 = .9$, $N = 20$, $m = 10$, and where (a) $L = 3$ and (b) $L = 2$. In (a), we see that the cultural Red King gets stronger as β gets higher. The greater the probability that actors initially follow discriminatory norms, the worse minority status is for the actors. In (b), we see a cultural Red Queen switch to a Red King as β increases. Normally minority actors would be advantaged by their swift learning in this case, but when discriminatory norms are in play, their swift adaptation to these norms instead causes a disadvantage.¹⁶

¹⁶In O'Connor and Bruner [2017], we discuss a similar observation for a non-agent based model, and in O'Connor [2017] I expand on this observation.




Figure 11: Proportions of outcomes for Nash demand game where $p_1 = .9$, $N = 20$, $m = 10$, $\epsilon = 0$. Players follow existing discriminatory norms.

Taken together, the results presented can significantly strengthen our confidence that the cultural Red King might arise in the real world. First, this phenomenon is robust across quite variable modeling choices. Second, in agent based models, it is not bounded as it is in the infinite population models explored by previous authors. Third, when we incorporate reasonable psychological features—risk aversion and in-group preference—we see the cultural Red King strengthened and the cultural Red Queen weakened. Fourth, in explicit network models of bargaining, the Red King, but not the Red Queen is observed. And last, the presence of existing discriminatory norms against minority groups—prevalent in many societies—leads to greater minority disadvantage when bargaining emerges in new arenas.

In the next section, I will discuss how the cultural Red King is not just robust across variations in the model, but is also deeply related to previous formal results where differences in learning speed or reactivity are generated via asymmetries in information structures. Justin Bruner first made this point in his dissertation work, and I expand on it here.

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. 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. Actors take samples of past play by all individuals, which is why this is best thought of as institutional memory. Young further

assumes that members of the groups may have access to different amounts of history. This difference changes the reactivity of particular individuals. Those with more memories will change strategies less readily, while those with less 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 gets smaller, though, the population will spend more and more time at a subset of equilibria. At the limit, the probability that the population is at one of this subset is 1. These are the SSEs.

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).¹⁷

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 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).¹⁸ 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

¹⁷In particular, Young shows that as the partition of demands for the actors gets finer and finer, the SSE approaches 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.

¹⁸In 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.

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 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 threat, confirmation bias, etc. 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. 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 2017.
- 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.
- Luigi Guiso and Monica Paiella. Risk aversion, wealth, and background risk. *Journal of the European Economic association*, 6(6):1109–1150, 2008.
- S Alexander Haslam. *Psychology in organizations*. Sage, 2004.
- 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. Dynamics of inequity. Book Manuscript, 2017.
- Cailin O’Connor and Justin P Bruner. Dynamics of minority disadvantage in collaboration. Working Paper 2017.
- Cailin O’Connor, Liam K. Bright, and Justin P. Bruner. The evolution of intersectional oppression. Working Paper 2017.

- Peter J Richerson and Robert Boyd. *Not by genes alone: How culture transformed human evolution*. University of Chicago Press, 2008.
- Hannah Rubin and Cailin O'Connor. Discrimination and collaboration in science. Working Paper 2017.
- W.H. Sandholm, E. Dokumaci, and F. Franchetti. Dynamo: Diagrams for evolutionary game theory. page <http://www.ssc.wisc.edu/whs/dynamo.>, 2012.
- William H Sandholm. *Population games and evolutionary dynamics*. MIT press, 2010.
- 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.
- Henri Tajfel. Experiments in intergroup discrimination. *Scientific American*, 223(5):96–102, 1970.
- Henri Tajfel. Interindividual behaviour and intergroup behaviour. *Differentiation between social groups: Studies in the social psychology of intergroup relations*, pages 27–60, 1978.
- 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.