

Minority (dis)advantage in population games

Justin P. Bruner¹

Received: 20 December 2016 / Accepted: 30 June 2017 / Published online: 6 July 2017
© Springer Science+Business Media B.V. 2017

Abstract We identify a novel ‘cultural red king effect’ that, in many cases, results in stable arrangements which are to the detriment of minority groups. In particular, we show inequalities disadvantaging minority groups can naturally arise under an adaptive process when minority and majority members must routinely determine how to divide resources amongst themselves. We contend that these results show how inequalities disadvantaging minorities can likely arise by dint of their relative size and need not be a result of either explicit nor implicit prejudices, nor due to intrinsic differences between minority and majority members.

Keywords Social philosophy · Social norms · Game theory · Evolutionary game theory · Bargaining · Distributive justice

1 Introduction

We develop a simple game-theoretic model and identify a novel “cultural red king effect” which, in many cases, results in stable arrangements that are to the detriment of minority groups. In particular, we show that when members of two groups must routinely interact so as to divide resources amongst themselves, it is exceptionally likely that an adaptive process such as the replicator dynamic or social learning will lead to a minority disadvantage equilibrium in which members of the larger group demand the bulk of a contested resource when interacting with members of the minority, who in turn acquiesce to the demands of the majority group. We contend that this minority disadvantage is due to a cultural analogue of the red king effect from evolutionary

✉ Justin P. Bruner
justin.bruner@anu.edu.au

¹ School of Politics and International Relations, Australian National University, Canberra, Australia

biology, whereby when co-evolving groups interact, the group which evolves more slowly in response to their counterpart gains an advantage. While this effect was first discussed by [Bergstrom and Lachmann \(2003\)](#), we identify a cultural red king effect whereby differences in response rate between groups is due not to different exogenous rates of biological evolution, but is instead driven by the relative size of the competing groups. In short, the members of larger groups are slower to respond to the behavior of out-group members, thereby placing them at a distinct advantage in a variety of important strategic scenarios.

We believe our results demonstrate how minorities could possibly come to be systematically disadvantaged in virtue of their minority status alone. In the models we develop, members of minority and majority groups are similarly situated and, for instance, share the same utility function and cognitive capacities. Furthermore, members of the majority harbor neither prejudice nor malice toward minority members, nor are they implicitly biased against minority members. In this sense, our results echo previous work by [Schelling \(1978\)](#), who demonstrated how it is possible for undesirable macro-phenomena (such as racial segregation) to come about despite the fact that individual agents are not motivated to individually bring about (and may in fact actively oppose) said macro-phenomena. To clarify, we do not argue that the mechanism uncovered in this paper is in fact responsible for most real-world instances of minority disadvantage. Instead, we establish that it is *possible* for minorities to be systematically disadvantaged by dint of their minority status alone.

This paper proceeds as follows. In Sect. 2 we discuss previous work on bargaining and introduce the basic tools from evolutionary game theory we employ throughout the paper. Section 3 houses our central results. We demonstrate that minorities are at a distinct disadvantage in a variety of strategic settings. In Sect. 4 we explore how robust these results are, and uncover that in certain circumstances minorities may in fact gain the upper hand. Section 5 concludes.

2 The dynamics of division

The question of how agents come to divide resources amongst themselves is well-worn territory in the social sciences, biology and game theory. We consider the simplest non-trivial strategic interaction, the so-called mini-Nash demand game ([Skyrms 1996](#)). In this scenario, two individuals must simultaneously make one of three possible demands regarding how a resource should be split. If the sum of their joint demands is equal to or less than the total amount of the resource on offer, then both agents receive the amount they demand. If, however, the sum of the joint demands exceeds the total amount of the resource, then neither receive their demands. This latter scenario corresponds to a situation in which negotiation breaks down and the agents are both left empty handed.

An example of the mini-Nash demand game is presented in normal form in Fig. 1. In this case, a total of 10 units of the resource are on offer and individuals can select to demand four, five or six units for themselves. Any combination of demands which sums to 10 is a Nash equilibrium, meaning neither agent has incentive to unilaterally change their behavior by demanding more or less of the resource. Thus the equal

	Demand 4	Demand 5	Demand 6
Demand 4	4, 4	4, 5	4, 6
Demand 5	5, 4	5, 5	0, 0
Demand 6	6, 4	0, 0	0, 0

Fig. 1 Nash demand game with three strategies

split in which both agents receive five units of the resource is stable. But so is the asymmetric split in which one individual demands the high amount of six while the other concedes and demands the low amount of four. There is also a mixed Nash equilibrium where agents mix between all three demands.

Note that since there are multiple equilibria in this game, rational choice theory does not provide a determinate answer as to which equilibria will be reached. The equal split as well as the asymmetric split are consistent with both participants behaving rationally. To shed light on how play will unfold in the mini-Nash demand game, Sugden (1986) and Skyrms (1996) have independently studied how strategic behavior in the mini-Nash demand game will change over time through either individual or social learning. Skyrms, for instance, considers the so-called replicator dynamic, a commonly used model of biological and cultural evolution (Sandholm 2010). The replicator dynamic specifies how the behavior of a population of agents playing a particular game will change over time as a function of how well the various strategies perform. In particular, under the replicator dynamic strategies that do better than the population average will increase in number, while those that fare poorly will decline in number.¹

Returning to the game of Fig. 1, Skyrms (1996) has shown that under the replicator dynamic a population of agents tasked to play the mini-Nash demand game will end up at one of two stable endpoints.² The first of these corresponds to the equal split. In this case, all members of the population will demand five of their counterpart. This arrangement is stable, for invaders who demand less than this will fare worse than natives (receiving a payoff of four instead of five) and greedy invaders will walk away empty handed (since their demands are incompatible with those of the natives). The second arrangement is a so-called polymorphism, or mixed equilibrium, in which a particular mixture of agents who demand four and six is possible. Once again, this arrangement is stable and cannot be perturbed by the introduction of a small handful of mutants since an individual demanding five will butt-heads with natives demanding six but will leave ‘money on the table’ when interacting with those demanding four. Note, however, that while this second arrangement is indeed stable, it is also inefficient: when both agents demand four, a portion of the resource is left unclaimed, and when both agents demand six, bargaining breaks down.

¹ While the replicator dynamic was initially constructed to model evolution by natural selection, it is formally equivalent to models of social learning via imitation, as well as models of individual learning such as the Herrnstein reinforcement model used in psychology.

² See Ernst (2001) for alterations to the replicator dynamic that lead to other evolutionary possibilities.

How to avoid these ‘polymorphic traps’ is a question that has received considerable attention in recent years (see, [Zollman 2005](#); [Alexander 2007](#)). One means of sidestepping these inefficient arrangements involves segmenting the population into groups of individuals who then interact with one another. [Skyrms and Zollman \(2010\)](#) consider the case in which agents are split into two groups, labeled ‘Blue’ and ‘Green’. If agents are tasked to just interact with out-group members (i.e., Blues never interact with fellow Blues), then inefficient arrangements can be completely avoided. In this case, play evolves to either the equal split, in which both those from the Blue and Green group demand five, or an asymmetric split, where Blues (Greens) always demand the high amount of six and Greens (Blues) acquiesce and demand four.³ [Axtell et al. \(1999\)](#) reinforce this point via computer simulation and demonstrate the asymmetric split is possible in a slightly different setting.

It is easy to see why interactions will head toward efficient outcomes when agents only interact with out-group members in the way described above. Consider the case in which both Blue and Green populations consist of a mixture of those who demand four and demand six. If 60% of Blues demand the high amount of six while only 30% of Greens demand six, the vast majority of Blues who demand the high amount will interact with a Green agent demanding four. Thus the payoff associated with being a Blue who demands six will be particularly high. Likewise, Greens who demand six will interact for the most part with agents who similarly demand six, resulting in a rather paltry payoff. This gives rise to a positive feedback loop, whereby the proportion of those who demand six in one group rises, in turn increasing the average payoff these agents receive, which leads to them further proliferating.

Things are more complicated, however, if we allow for interactions between both in-group and out-group members. In this case we assume individuals can accurately determine the group membership of their counterpart (i.e., whether they are Green or Blue) and group membership is fixed (that is, one cannot switch from Green to Blue). Additionally, individuals can condition their behavior on the group membership of their strategic partner. In other words, agents can select to offer an equal split to Blue agents, and demand the bulk of the resource when interacting with Green agents. In our simple mini-Nash demand game, there are a total of nine such type-conditional strategies available to the agent.

It is worth noting the significance of allowing for both in-group and out-group interactions as it enables us to investigate, among other things, the emergence of in-group bias and ethnocentric behavior (see [Hammond and Axelrod 2006](#)). For instance, play can unfold such that Green agents treat their fellow Greens in an equitable fashion but demand the bulk of the contested resource when interacting with out-group members. Furthermore, by allowing for in-group interactions we also reintroduce the possibility of polymorphic traps. Interactions between members of the Blue group, for instance, can lead to the inefficient mixture of agents demanding 4 and demanding 6 discussed above. Thus while interactions between Blues and Greens will always result in the resource being exhausted, play within groups can result in inefficient outcomes. In what follows, we consider the situation in which Green and Blue groups are not of the

³ [Skyrms and Zollman \(2010\)](#), p. 271.

same size. We find that this has rather dramatic effects that can, in many circumstances, result in the minority systematically demanding the low amount when interacting with members of the majority.

3 Minority disadvantage and the red king effect

To see how minority members are at a distinct disadvantage in the mini-Nash demand game we modify the two-population replicator dynamic. As mentioned earlier, the replicator dynamic is a standard model of cultural evolution that has recently been appropriated by social philosophers to better understand the dynamics of social norms pertaining to, among other things, fair division, collective action and altruism.⁴ Generally, under the replicator dynamic individuals are randomly paired with a partner to then play a simple game (such as the mini-Nash demand game). Then the average payoff of the group as a whole is calculated, as well as the average payoff for each strategy. Strategies which do better than the group average increase in frequency, while those failing to do better than average decrease in frequency. We below consider a version of the so-called two-population replicator dynamic, where individuals are tasked to play with in-group and out-group members and can, if they so choose, condition their behavior on the group-membership of their counterpart. The standard equation for this is as follows:

$$x_i(t + 1) = x_i(t) \frac{F_i}{F}$$

where $x_i(t)$ refers to the proportion of agents from the Blue population utilizing strategy i at round t , F_i denotes the average payoff of Blue members utilizing strategy i and F is the average payoff of members of the Blue population (a similar sort of equation governs the dynamics for the Green population). Note that when calculating F_i , the average payoff to Blue agents utilizing strategy i , we consider the composition of both Blue and Green populations since, as we mentioned earlier, individuals interact with both groups. In other words, F_i is the sum of two terms, the average payoff the agent receives against in-group and out-group members. Formally, this is $F_i = \sum_{l=1} \pi(i, l)x_l(t) + \sum_{k=1} \pi(i, k)y_k(t)$, where $x_l(t)$ is the proportion of Blue agents utilizing strategy l , $y_k(t)$ is the proportion of Green agents utilizing strategy k and $\pi(i, l)$ is the payoff a Blue agent utilizing strategy i receives when interacting with an agent utilizing strategy l . Note that since the above summations are given equal weight and (as is the case under the standard replicator dynamic) individuals are randomly paired, the above equation captures the fact that the Green and Blue groups are of equal size. Half of an individual's payoff is determined by their interactions with out-group members, while the remaining half of their payoff is determined by their interaction with in-group members. This, however, can be easily changed to reflect the fact that agents of one population are more numerous. For instance, we can weigh the two

⁴ See, for instance, Bruner (2015), Skyrms (2004), Alexander (2007) and Smead (2008).

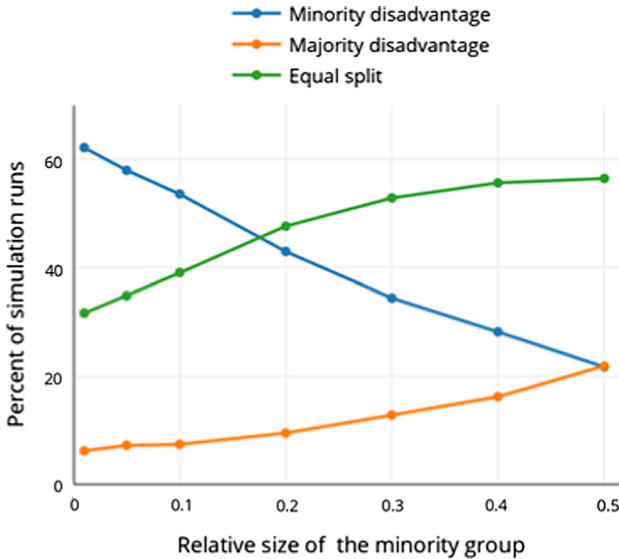


Fig. 2 Simulation results of the modified two-population replicator dynamic

summations in the above expression to reflect the fact that, say, Blues are twice as numerous as Greens.⁵

With this modification in hand, we can now proceed to explore via computer simulation how minority and majority members fare when tasked to bargain over contested resources. Simulations are carried out as follows. We first randomly sample the space of possible population compositions and feed these initial population statistics into the replicator dynamic. This generates the population composition for the next time period which are, in turn, input back into the replicator equations once again. This process continues until the population arrives at one of the stable endpoints or ‘equilibria’ in the underlying divisional game. We refer to the basin of attraction of a particular equilibrium as the proportion of simulations which result in play evolving toward said equilibrium.

Our central simulation results are displayed in Fig. 2. When the two populations are of the same size, nearly 60% of the time an equal split emerges between the two populations. The remaining cases are equally split between Blues demanding the high amount from Greens and vice versa. In other words, when Blue and Green groups are of the same size, the even split is the most likely outcome, and both groups have the same chance of gaining the upper hand.

This changes as one group shrinks in size. When one group is in the minority, the even split is less likely to hold and it becomes increasingly likely that the equilibrium in which the minority demands the low amount of four from the majority comes to fruition (we call this the minority disadvantage equilibrium). This effect becomes all

⁵ In this case, the average fitness of Blue individuals utilizing strategy i is $2 \sum_{l=1} \pi(i, l)x_l(t) + \sum_{k=1} \pi(i, k)y_k(t)$.

the more pronounced as the minority shrinks in size. For instance, when the minority makes up only one-percent of the population, approximately 62% of simulations result in the minority disadvantage equilibrium, while less than 7% of simulations lead to the majority disadvantage equilibrium in which minorities demand the bulk of the resource from members of the majority.

Thus it appears that in bargaining contexts smaller groups are sometimes at a distinct disadvantage. In particular, it is extraordinarily likely that they receive less than half of the contested resource when paired with out-group members. Moreover, since under the replicator dynamic individuals are randomly paired to interact, minority members are exceptionally likely to encounter and bargain with members of the majority (in virtue of their sheer numbers), meaning a large proportion of their interactions will involve them conceding the bulk of the contested resource to their counterpart. Finally, as mentioned in the previous section, this asymmetric split is stable. Once the minority disadvantage equilibrium is reached, neither members of the majority nor minority have incentive to deviate and demand a different fraction of the resource.

We contend that what drives this ‘minority disadvantage effect’ is an *asymmetry of response* which holds between the two groups. Consider the fact that since agents are randomly paired, members of both the minority and majority interact for the most part with majority members. As a result, how well an individual fares in the Nash demand game in large part hinges on her interaction with members of the majority. How individuals behave when paired with members of the minority, on the other hand, is not particularly important from a strategic point of view since only a relatively small fraction of interactions are with members of the minority. As a result, minority members quickly adapt to the behavior of the majority, while members of the majority only very slowly adjust their behavior in response to the minority. This results in an asymmetry of response: the minority population quickly adjusts to the initial behavior of the majority while members of the majority only slowly alter their behavior in light of the actions of the minority group.⁶

In many circumstances, such as the mini-Nash Demand game of Fig. 1, the minority’s best recourse given the initial random composition of strategies in the majority population is often to utilize a ‘safe’ strategy which guarantees a modest payoff regardless of what their counterpart does. In the context of the Nash Demand game, this results in the minority quickly adopting the ‘demand four strategy when interacting with members of the majority. Once this occurs, the majority slowly learns to demand the bulk of the resource from the minority, resulting in the minority disadvantage equilibrium.

This dynamic is an instance of a more general phenomena in evolutionary biology called the Red King Effect. Bergstrom and Lachmann (2003) consider situations in which two mutualistic species interact, co-evolving at different rates. They demonstrate with an evolutionary model that in many cases the species which evolves at a slower rate gains an advantage in the long run. As Bergstrom and Lachmann put it, in these mutualistic settings, the slower runner ‘wins the evolutionary race’. This occurs for the reasons we discussed above: the group that rapidly evolves to match the behavior of

⁶ Once again, the initial composition of both populations was determined by randomly sampling the space of all possible population compositions.

their counterpart often ends up ‘playing it safe’ and forgoing benefits in many strategic settings.

Note that while Bergstrom and Lachmann are agnostic as to why the two groups evolve at different rates, our scenario provides a straightforward and intuitive story as to why there exists an asymmetry of responsiveness between the two groups.⁷ In particular, one group is less important from a strategic point of view than the other due to its small size. Hence both groups will quickly respond to the behavior of the more strategically important group, meaning minorities respond to the behavior of the majority at a faster rate than the majority responds to the minority. As we’ve seen, this ‘cultural red king effect’ tilts the scales in favor of the majority by pushing play toward the minority disadvantage equilibrium. In what follows we consider how robust this effect is.

4 How robust is the red king?

In this section we consider alterations to our baseline model and explore whether minority groups are still systematically disadvantaged upon relaxing various assumptions. We explore four alterations in particular. The first involves the introduction of positive assortment (correlated interactions). Perhaps, the thought goes, minority members will do better if they are more likely than chance to interact with in-group members. Next, we examine in what way our results depend on our choice of dynamic and explore one popular alternative to the replicator dynamic, the so-called *best response dynamic*. This is followed by the introduction of ‘noisy signals’ to our model. Finally, we explore whether minority disadvantage still exists under alternate payoff structures. We find that in some conflictual coordination games, such as the battle of the sexes, there exists a *majority disadvantage effect*.

4.1 Positive assortment

The effects of correlated interactions have been long known in evolutionary biology. [Hamilton \(1964\)](#), for instance, was already aware that correlation is key to promoting altruistic behavior: self-sacrificing and cooperative behavior can thrive if altruists are more likely than chance to interact with fellow altruists. [Skyrms \(1996\)](#) has likewise shown that when interactions are correlated in the mini-Nash demand game it is less likely the population ends up at an inefficient polymorphic trap. Perhaps, then, the addition of positive correlation to our setting will steer the population away from the minority disadvantage equilibrium by ensuring minority members are more likely than chance to interact with other in-group members. Furthermore, the inclusion of correlation makes for a more realistic model of social behavior. Recent work in sociol-

⁷ As I first noted in my dissertation ([Bruner 2014](#)), an asymmetry of response can be the result of differential population size, differential network connectivity ([Gallo 2015](#)) as well as the result of differences in institutional memory ([Young 1993](#)). Both Gallo and Young investigate the evolution of bargaining in these different contexts and arrive at results which are similar to those uncovered in this paper. Namely, in many bargaining contexts, asymmetries between parties results in one party responding at a slower rate than their counterpart, which in turn, yield benefits.

ogy suggests many personal networks are highly homogeneous with respect to certain demographic factors (such as race or ethnicity).⁸ That said, it should be noted that there are some important strategic contexts where interactions appear to not necessarily be positively correlated. For instance, interactions and collaboration in the workplace may not involve positive assortment, as individuals may be limited in their ability to choose who they interact and collaborate with. If this is correct, then our baseline model applies to such cases.⁹

We incorporate positively correlated interactions into our model as follows (see, Skyrms 1996). The chance an agent interacts with an in-group member is no longer just the proportion of in-group members in the population. Instead, the likelihood a Blue agent interacts with a fellow Blue agent is $p(B) + ep(G)$, where $p(B)$ and $p(G)$ refer to the proportion of Blue and Green agents in the population at large, respectively, and e refers to the level of correlation.¹⁰ When $e = 1$ interactions are perfectly correlated, i.e., the chance of interacting with an in-group member is simply 1. For $e = 0$, interactions are completely uncorrelated and the probability a Blue agent meets another Blue agent is simply the proportion of Blue agents in the population ($p(B)$).

There are two points worth noting here. First, positive assortment (typically) helps members of the minority. In particular, at the minority disadvantage equilibrium, positive correlation has the effect of *insulating members of the minority* from the unequal divisions they receive when interacting with out-group members. Of course, positive assortment does not benefit the minority at the *majority* disadvantage equilibrium. In this setting, minority members do best to interact with as many members of the majority as possible.

Second, while positive assortment reduces the proportion of interactions which take place with out-group members, the minority disadvantage equilibrium is still extraordinarily likely.¹¹ In fact, the level of correlation does not seem to affect the likelihood of the minority disadvantage equilibrium. Figure 3 clearly illustrates this: divisions disadvantaging members of the minority are just as likely to take root in the population regardless of the level of correlation.¹² This is due to the fact that positive assortment does not change the minority or majority group's *relative responsiveness* to out-group members. As we saw in the previous section, the reason why the minority group is at a distinct disadvantage is because the minority is more responsive to the majority than the majority is to the minority. The inclusion of correlation preserves this asymmetry of responsiveness.¹³ Although like-types are more likely to meet, it is still

⁸ See, for instance, McPherson et al. (2001).

⁹ For more on how norms disadvantaging minorities could develop in the workplace or in academia, see Bruner and O'Connor (forthcoming) and O'Connor and Bruner (forthcoming).

¹⁰ Likewise, the chance of interacting with a Green out-group member is $p(G)(1 - e)$ instead of $p(G)$.

¹¹ This was determined by computer simulation.

¹² That said, the average payoff of the minority is higher when interactions are correlated since, as mentioned earlier, minorities are more likely to interact with in-group members.

¹³ Of course, since interactions are correlated, both minority and majority are interacting with out-group members less frequently, which means both are less responsive to the behavior of out-group members. Yet the minority is still significantly more responsive to the majority than the majority is to the minority. To

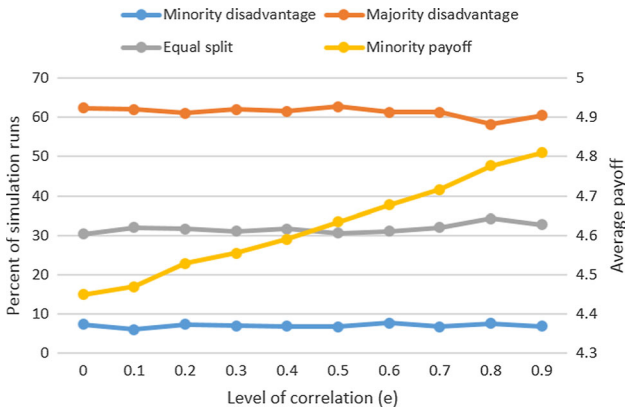


Fig. 3 Simulation results of the modified two-population replicator dynamics with correlation. Size of basin of attraction for minority disadvantage equilibrium (*blue*), majority disadvantage equilibrium (*orange*), equal split equilibrium (*grey*), payoff to minority members (*yellow*). Minority constitutes 1% of the total population. (Color figure online)

the case that a higher proportion of interactions for the minority are with out-group members, meaning there is still more pressure for the minority to quickly adapt to the behavior of the majority.

4.2 Best response dynamics

We next explore what effect choice of dynamic has on our results. While the replicator dynamic is a common mathematical model of evolution and social learning, there are a variety of alternative learning rules that are routinely employed by social scientists and biologists. Moreover, it is widely recognized that different dynamics may result in drastically different patterns of behavior. For instance, [Wagner \(2013\)](#) has shown how certain imitative learning dynamics need not lead to an equilibrium in even the simplest of settings. For these reasons, we explore whether the cultural red

Footnote 13 continued

see this, consider the situation in which Greens and Blues constitute 20 and 80% of the total population, respectively. In the absence of correlation, 80% of the minority's interactions are with out-group members while only 20% of the majority's interactions are with out-group members. Thus the minority is four times as likely as the majority to interact with an out-group member (80/20). Now consider the situation involving positive correlation. In particular, set e to 0.2. Recall that the likelihood of interacting with an out-group member is $(1 - e)pr(Out)$, where $pr(Out)$ is the proportion of out-group members in the population. Thus members of the Green minority will interact with out-group members with probability $(1 - .2)0.8$ or 64% of the time. Blue majority members will interact with members of the minority with probability $(1 - .2).2$ or 16% of the time. Note that while both groups are now less likely to interact with out-group members it still remains the case that members of the minority are four times more likely than members of the majority to interact with out-group members (64/16). Thus the relative responsiveness of the two populations remains the same despite the increased level of correlation. Since this asymmetry in responsiveness is what determines whether the minority or majority disadvantage equilibrium is reached, the cultural red king effect uncovered in Sect. 3 still holds.

king effect holds under a particularly popular dynamic, the so-called best response dynamic.

Under the best response dynamic some fraction of the population at each time period best responds to the composition of the population from the previous time period. In other words, the best response dynamic assumes a proportion of individuals myopically best respond to the aggregate behavior of the recent past. Models of myopic best response can be traced back to [Brown \(1951\)](#), who first proposed the learning rule known as ‘fictitious play’ in which individuals adjust their behavior to an interactive partner by myopically best responding to a history of past play.

We find there is no cultural red king effect under the two-population best response dynamic.¹⁴ This is because minority and majority populations are equally responsive to each other: since the same proportion of individuals from both populations are chosen to best respond to the behavior of both in- and out-group members, the minority and majority adjust to each other at the same rate.¹⁵

While the cultural red king effect is not possible under the best response dynamic, it does appear under a variety of other adaptive dynamics tailored to model cultural evolution or social learning. This is due to the fact that for many dynamics, the rate at which strategies increase or decrease in frequency in a population is determined by how well a particular strategy does when compared to its competitors. For instance, under the replicator dynamic, all strategies that do better than the population average increase in frequency and the higher the average payoff of a strategy, the greater the rate of increase. It is this property of the dynamic which allows for an asymmetry of response between minority and majority populations. It is worth noting that the replicator dynamic is not unique in this regard. For instance, the logit dynamic, Brown-von Neumann–Nash dynamic and Smith dynamic all allow for a cultural red king effect since the rate of change for a particular strategy is a function of how that strategy fares when compared to the rest of the population.¹⁶

4.3 Noisy signals

Up to this point we have assumed individuals can clearly and without error determine the group-membership of their counterpart. Accurately determining group-membership is particularly important given our agents utilize conditional strategies, allowing them to treat members of minority and majority groups differently. We now consider a modification of our baseline model and explore a situation in which signals of group-membership are susceptible to noise. In other words, individuals cannot with surety ascertain the group-membership of their counterpart.

¹⁴ This was established on the basis of computer simulations.

¹⁵ However, this would not be the case if the best response dynamic was modified so the same *number* of individuals from each group was selected to best respond. In this case, a larger proportion of the minority group would update their strategic behavior than the majority group, and as a result the minority would quickly learn to accommodate the majority.

¹⁶ See [Sandholm \(2010\)](#) for a discussion of these various adaptive dynamics and their properties. It is worth noting that these three dynamics along with the best response dynamic and replicator dynamic are among the most popular deterministic dynamics used in evolutionary game theory.

This is done simply enough. Let the reliability of a signal of group-membership, r , be the probability an individual will correctly identify the group of origin of her counterpart. If an individual fails to correctly recognize another, then they mistakenly treat their counterpart as a minority when they are in fact a majority member, or a majority member when they in fact belong to the minority. Hence, the likelihood of mutual (accurate) recognition is r^2 . The probability both bargainers fail to correctly recognize each other is $(1 - r)^2$ and with probability $2r(1 - r)$ exactly one of the two individuals fails to correctly recognize their counterpart.

We modify the payoff table of the Nash demand game to account for the possibility of noisy group-membership signals. This is a rather straightforward process. For instance, when a Green individual using the strategy ‘demand five if counterpart is Green and demand six if counterpart is Blue’ interacts with a Blue agent using the strategy ‘demand four if counterpart is Green and demand four if counterpart is Blue’, the resulting allocation will provide the Green agent with six and the Blue bargainer with four. If signals are not entirely reliable, however, then the expected payoff to the Green agent will instead be $6r^2 + (6 + 5)r(1 - r) + 5(1 - r)^2$ while the expected payoff for the Blue agent is 4.

We find that the red king effect still holds for moderate levels of noise. For instance, when the minority only constitutes 25% of the total population the minority is still at a disadvantage as long as r is greater than 0.95.¹⁷ In other words, as long as signal error is kept below 5%, our central result holds.¹⁸ This suggests the red king effect is especially likely to hold when groups are determined on the basis of race or gender. In such cases, we would expect signals of group membership to be on the whole quite reliable.

4.4 Alterations to the payoff table

We have shown in the previous section that minorities face a distinct disadvantage when tasked to bargain over resources. Yet bargaining scenarios constitute just one small—albeit very important—slice of the total number of possible strategic situations individuals are confronted with daily. Does this minority disadvantage effect spill over to other settings? We see that it need not, and moreover, there are circumstances in which it pays to be in the minority.

Consider the conflictual coordination game, the battle of the sexes (BoS). In this game two individuals must arrange a meeting place. We will assume that agents can either meet at the coffeehouse or the teahouse (Fig. 4). While both agents prefer to meet at the same place than to not meet at all, individuals from the Brown group prefer a meeting at the coffeehouse to a meeting at the teahouse. On the other hand, those from the Black group prefer the teahouse to the coffeehouse.

¹⁷ To be clear here, what we mean is that for $r > 0.95$ the minority disadvantage equilibrium is more likely to emerge than compared to the ‘baseline’ model in which $r = 1$ and both groups are of equal size. When r dips below this 0.95 threshold the equal split becomes overwhelmingly likely.

¹⁸ This threshold, of course, depends on parameters in the model (such as the size of the minority group). In general, however, a relatively high level of signal noise still allows for the red king effect.

		Brown group	
		Teahouse	Coffeehouse
Black group	Teahouse	4, 2	0, 0
	Coffeehouse	0, 0	2, 4

Fig. 4 Battle of the sexes

Just as in the previous section, individuals can distinguish between in- and out-group members and can further condition their strategic behavior on this knowledge. Not surprisingly, a number of equilibria are possible. For example, members from both groups can ignore the group membership of their counterpart and simply meet at the coffeehouse. Likewise, Brown individuals can concede and meet Black counterparts at the teahouse, while patronizing the coffeehouse when interacting with fellow Browns. Each group desires to meet at their favorite establishment. Which group wins out in the end?

The answer depends on the relative size of the Brown and Black groups. If the groups are of the same size, then the coffeehouse and teahouse both have the same chance of becoming the de facto meeting place when Brown and Black individuals are paired to interact. Surprisingly, if we vary the relative sizes of the two groups we uncover that the minority group has an advantage. When Browns only constitute a small five percent of the overall population, over 64% of simulations result in an equilibrium in which Browns and Blacks meet at the coffeehouse, the preferred location of the minority Browns. Once again, an asymmetry of response explains this ‘majority disadvantage effect’. The majority is very unresponsive to the behavior of the minority, and this once again means the minority rapidly adapts to the initial distribution of strategies in the majority population. In the Nash demand game, this meant the minority more often than not did best to demand the low amount of four. In the BoS, it is often the minority Browns best response to simply patronize their preferred establishment, the coffeehouse.

In light of these results it is apparent that we cannot say that the minority is always at some sort of disadvantage in virtue of its minority status. In certain classes of games, such as the Nash demand game from Fig. 1, the minority is more likely to be disadvantaged, stuck demanding a small portion of the contested resource. In other circumstances, such as the BoS, the minority does well, in the sense that the population frequently goes to the equilibrium they favor. It is not the case that in all strategic situations the minority is at a disadvantage.

Situations where the minority gains the upper hand, such as the BoS, are instances of the so-called Red Queen effect from biology (Van Valen 1977). The red queen effect occurs when two species interact in a strategic setting, the group that evolves more quickly gains an advantage. As in Sect. 3, the reason one population responds more quickly than the other is due to its size. Larger groups are less sensitive to the behavior of the minority since the minority is strategically not very important. Likewise, members of the minority quickly best respond to the majority since how well they perform in these strategic settings in large part hinges on how they interact with members of the majority. Yet while the minority behaves in an accommodating

fashion in the Nash demand game, the asymmetry due to group size actually results in the minority getting what they most desire in the BoS.

5 Discussion

In the course of this paper we have uncovered a novel cultural red king effect, whereby stable arrangements disadvantaging members of the minority are likely to arise in virtue of the group's minority status. We take this work to have shown how it is possible for minorities to become systematically disadvantaged without presupposing widespread prejudice or implicit bias, or appealing to intrinsic differences between minority and majority members. As discussed, our discovered effect seems somewhat robust, in the sense that it holds despite the alteration of various aspects of the game-theoretic model. That said, it is worth noting this effect can be reversed, resulting in a cultural red queen effect where the *majority* disadvantage equilibrium is the most likely outcome of an adaptive dynamic. Nonetheless, the results of this paper highlight a variety of interesting points and future avenues of research.

We have uncovered how relative group size determines how responsive a group is when dealing with out-group members. This is interesting for a variety of reasons. First, attaining a better understanding of the factors which influence responsiveness is vital given that social philosophers have recently begun to appreciate the important role responsiveness plays in the evolution of conventions and social norms. Brusse and Bruner (forthcoming), for example, find that in the David Lewis signaling game (Lewis 1969; Skyrms and Zollman 2010), communicative conventions are more likely to arise when the sender population is more responsive than the receiver population. Second, unlike prior work on the red king effect in which responsiveness is determined exogenously by a parameter in a game-theoretic model (Bergstrom and Lachmann 2003), the asymmetry of response between minority and majority groups is a natural consequence of the relative size of the groups involved in the strategic interaction. Other means by which group response rate can endogenously be determined are worth exploring further and will help shed light on the emergence of strategic behavior in a wide array of contexts.

The central result of this paper has important social implications as well. We have shown how it is possible for divisional norms disadvantaging minority groups to frequently take root in a community whenever individuals (i) can condition their behavior on group membership and (ii) learn to act in a way which promotes their self-interest in strategic contexts. These conditions both seem likely to hold, suggesting the cultural red king effect may crop up in many real-world situations.

Furthermore, the cultural red king effect also highlights how morally problematic arrangements can come about despite the fact those in the community harbor no animosity toward members of the minority. Instead, agents in our model simply act in their best interests when dealing with both in-group and out-group members. Furthermore, we do not assume there is any intrinsic difference between minority and majority members. In other words, members of both groups value the resource under contention equally, and neither group is inherently better at navigating strategic encounters. *Thus, it is possible for minorities to be at a real disadvantage in bargain-*

ing contexts by dint of their size alone. This is particularly significant for it suggests policies aimed at reducing prejudice and implicit bias may on their own not be enough to promote the interests of minority members. Even in the absence of prejudice, the dynamical effects outlined in this paper may nonetheless push the community toward arrangements which are to the detriment of minorities.

Acknowledgements Thanks to Brian Skyrms, Simon Huttegger, Jean-Paul Carvalo, Peter Vanderschraaf, Cailin O'Connor, Kim Sterelny and Bob Goodin for comments and suggestions. Additional thanks to audience members of the Social Dynamics Seminar at UC Irvine, Moral, Social and Political Theory Workshop at ANU, Formal Ethics 2015 at Bayreuth, the 2014 meeting of the Association for the Study of Religion, Economics and Culture at Chapman University and the Philosophy Seminar Series at Monash University.

References

- Alexander, J. M. (2007). *The structural evolution of morality*. Cambridge: Cambridge University Press.
- Axtell, R., Epstein, J., & Young, P. (1999). The emergence of classes in a multi-agent bargaining model. In S. Durlauf & H. P. Young (Eds.), *Social dynamics*. Oxford: Oxford University Press.
- Bergstrom, C., & Lachmann, M. (2003). The Red King effect: When the slowest runner wins the coevolutionary race. *Proceedings of the National Academy of Sciences of the USA*, 100, 593–598.
- Brown, G. W. (1951). *Iterative solution of games by fictitious play. Activity analysis of production and allocation*. New York: Wiley.
- Bruner, J. P. (2014). *Diversity and cooperation*. PhD Thesis.
- Bruner, J. P. (2015). Diversity, tolerance and the social contract. *Politics, Philosophy and Economics*, 14, 429–448.
- Bruner, J. P., & O'Connor, C. (forthcoming). Power, bargaining and collaboration. In Weisberg, M., & Mayo-Wilson, C. (Eds.), *Scientific collaboration and collective knowledge*. Oxford University Press.
- Brusse, C., & Bruner, J. P. (forthcoming). Responsiveness and robustness in David Lewis signaling games. *Philosophy of Science*.
- Ernst, Z. (2001). Explaining the social contract. *The British Journal for the Philosophy of Science*, 52, 1–24.
- Gallo, E. (2015). *Communication networks in markets*. Working paper.
- Hamilton, W. D. (1964). The genetical evolution of social behavior II. *Journal of Theoretical Biology*, 7, 17–52.
- Hammond, R., & Axelrod, R. (2006). The evolution of ethnocentrism. *Journal of Conflict Studies*, 50, 926–936.
- Lewis, D. (1969). *Convention: A philosophical study*. Cambridge: Harvard University Press.
- McPherson, M., Smith-Lovin, L., & Cook, J. (2001). Birds of a feather: Homophily in social networks. *Annual Review of Sociology*, 27, 415–444.
- O'Connor, C., & Bruner, J. P. (forthcoming). Diversity and dynamics in epistemic communities. *Erkenntnis*.
- Sandholm, W. (2010). *Population games and evolutionary dynamics*. Cambridge: MIT Press.
- Schelling, T. (1978). *Micromotives and macrobehavior*. New York: Norton.
- Skyrms, B. (1996). *The evolution of the social contract*. Cambridge: Cambridge University Press.
- Skyrms, B. (2004). *The stag hunt and the evolution of social structure*. Cambridge: Cambridge University Press.
- Skyrms, B., & Zollman, K. (2010). Evolutionary considerations in the framing of social norms. *Politics, Philosophy and Economics*, 9, 265–273.
- Smead, R. (2008). The evolution of cooperation in the centipede game with finite populations. *Philosophy of Science*, 75(2), 157–177.
- Sugden, R. (1986). *The economics of rights, cooperation and welfare*. Oxford: Blackwell.
- Van Valen, L. (1977). The red queen. *The American Naturalist*, 111, 809–810.
- Wagner, E. (2013). The explanatory relevance of Nash equilibrium: One-dimensional chaos in boundedly rational learning. *Philosophy of Science*, 80, 783–795.
- Young, P. (1993). An evolutionary model of bargaining. *Journal of Economic Theory*, 59, 145–168.
- Zollman, K. (2005). Talking to Neighbors: The evolution of regional meaning. *Philosophy of Science*, 72, 69–85.