



If Evolution Favours Fairness, Why Does Inequality Persist?

Cailin O'Connor on power and the emergence of bargaining norms

In 1950 the game theorist John Nash introduced an influential model of bargaining now known as the Nash demand game. Two actors would like to divide a resource. Each makes a demand for some portion of the resource—50%, say, or 22% or 91%. If the two individuals have compatible demands, each gets what they asked for. If, however, the demands are incompatible in that they sum to more than 100%, it is assumed that the actors cannot agree and that as a result they each get a low payoff, sometimes called the ‘disagreement point’. This model of bargaining has proven an extremely flexible and useful one over the years. It captures a basic scenario where two actors divide something (time, money, trade goods, household labour), where the division may take many different forms, some of which favour each individual to greater or lesser degrees, and where an inability to settle on a division is bad for all parties involved.

It was only a few years later that Nash **extended his model** to one where power plays a role. In these later models, the disagreement point is relabelled as a ‘threat point’. The idea is that each individual can issue threats about what will happen when bargaining breaks down. If one individual wields greater power—political, or economic, or personal—their threats may be more credible, leading to an asymmetry in the threat points of the two players. On traditional treatments of the game, whichever player can issue a more serious threat should benefit in the bargaining scenario. Since they stand to lose less when bargaining breaks down, they can wield this fact to their advantage.

Game theoretic models of bargaining rely on assumptions of rational choice—actors make decisions based on what they believe will be best for them—but bargaining behaviour in the wild is also highly dependent on ‘normative’ expectations. Human societies have norms for bargaining specifying who should get what and how much, and, at least to some degree, bargainers tend to follow these norms. For this reason, philosophers like **Brian Skyrms** and **Jason Alexander** have switched contexts to investigate the Nash demand game using cultural evolutionary models. These models assume a population of people, all of whom sometimes meet and bargain. The

models then explore what behaviours emerge on the population level as the group undergoes some cultural evolutionary process.

In particular, these philosophers have focused on **explaining norms of fairness**. Suppose that you are going to divide a small cake among yourself and three friends. It's a good guess that you'll do so by cutting the cake into four equal pieces. Many cultures have norms for equal divisions of resources, and in laboratory settings experimental subjects tend to follow these norms by demanding 50% in the Nash demand game. In cultural evolutionary models the 50% demand tends to emerge with high frequency, in part because this is the only symmetric demand that perfectly divides a resource. Any other set of demands that efficiently splits the resource—60% versus 40%, or 90% versus 10%—will be asymmetric, meaning that if a group adopts these demands they will sometimes mis-coordinate. When two actors who demand 60% meet, they will reach the disagreement point. When two actors who demand 40% meet, they waste part of the resource. The 50% demand is then a special one.

In recent work with **Justin Bruner**, I have instead been interested in explaining why, given these results, we often see unfair or discriminatory bargaining norms, especially between members of different social groups. Used car sales are a situation where buyers and sellers bargain. **Experiments find**, though, that car sellers make lower initial offers to men than to women, and lower initial offers to white people than to people of colour. When researchers send out otherwise identical job applications they find that applications with **white** and **male names** are more likely to receive offers, and to receive higher starting salaries. Why the discrepancies if fairness is so evolutionarily attractive?

Economists find that when tags—otherwise irrelevant observable features like gender, race, or class membership—are introduced into cultural evolutionary models of bargaining, the fair demand is no longer special. Instead, in these models, unequal bargaining regularly emerges between groups. Symmetry is no longer so important because actors can use tags to coordinate bargaining even when demands are asymmetric. Suppose that in a group with star-bellied sneetches and blank-bellied sneetches, the star bellies always demand 60% when they meet the other type, and the blank bellies always demand 40%. These actors always efficiently divide the resource, even though they make unequal demands, and this is possible because of their tags.

In a scenario like this, **Justin and I ask**, how does power influence the emergence of bargaining norms between those in different social groups? Economists like Nash have argued that power should improve an actor's bargaining prospects when two individuals bargain, but can it also impact society-wide patterns of bargaining?

To investigate this question, we have to start by saying something about what it means for one social group to have more power than another. In other words, we have to 'operationalize' power, or find a precise way to include it in a model of the evolution of bargaining. There are many ways to do this, so as a first stab we look at three types of power that a social group might have.

First, we consider social groups with different disagreement points. This difference could be induced by differential threats, as described above, or it could be a result of different cultural contexts. To illustrate how this might work, let's consider two examples. In order for a household to run properly, the members of the household have to decide who does what labour, and how much of it. This represents a scenario of either implicit or explicit bargaining. Because most households are formed by heterosexual couples, the majority of this bargaining happens between members of two different social groups—men and women. Furthermore, in many societies, different levels of economic empowerment mean that should household bargaining break down, the outcomes will typically be different for men and women. If women are unable to work outside the home, or can only earn low wages by doing so, they have lower disagreement points for the case of household bargaining.

Second, consider the example of collaboration across the academic hierarchy. (This is the case we explicitly consider in print.) Collaborating academics must bargain, implicitly or explicitly, to decide

who does what and how much, and also to decide how the credit for the project is divided by determining author order. Furthermore, collaborations often occur between academics at different stages of their careers. For an established academic, if collaborative bargaining breaks down, it typically will have little to no career impact, whereas a graduate student's entire career may ride on the success of one project. This induces a scenario with systematically different disagreement points for those at different levels of the hierarchy.

Justin and I find that in models where two social groups learn how to bargain, and where one has a higher disagreement point than the other, the group with the higher disagreement point will be more likely to end up at a norm where they get more and the other group gets less.

The second operationalization of power we consider assumes that actors in one social group have higher background payoffs than the other. This translates to, for example, a scenario where men bargaining in the household have more streams of income or resources than women, so that the household bargain matters less to them. Or to a case where established academics have more projects on the burner, so that each one is less important. Again, we find that as we raise the background payoff for one social group, that social group becomes more likely to get more at the resulting bargaining norm.

Lastly, we operationalize power by supposing that one group has a higher 'outside option' than the other. In these models, all actors have the option to do something besides bargaining. They can form a solo household instead of getting married. Or they can work on an individual research project instead of collaborating. If one social group has a higher outside option, this could be a better ability to live alone, because of economic empowerment or other social norms. Or it could translate to a higher likelihood of publishing individual research because of reputational effects. Again, in this scenario, we find that the group with a higher outside option becomes increasingly unlikely to end up getting poor bargaining outcomes as their outside option increases.

Importantly, these observations mean that power can multiply and compound for those in different social groups. Economic empowerment today can translate to a bargaining advantage tomorrow, which can translate to further economic empowerment. We think this is part of the story of why bargaining inequity between social groups is both ubiquitous, and sustaining.

Cailin O'Connor is Assistant Professor of Philosophy at University of California, Irvine, where she is also a member of the Institute for Mathematical Behavioral Science. This essay draws on her paper, written with Justin Bruner, entitled 'Power, Bargaining, and Collaboration', forthcoming in [Scientific Collaboration and Collective Knowledge](#), edited by Thomas Boyer, Conor Mayo-Wilson, and Michael Weisberg (Oxford University Press). Her research focuses on the philosophy of biology and the behavioural sciences, with a particular emphasis on evolutionary game theory. This work was completed under O'Connor's NSF grant, [Social Dynamics and Diversity in Epistemic Communities](#).

Image credit: Mike Leary, '[See-Saw](#)'

Monday, 7 November 2016

Share This Story, Choose Your Platform!

Related Posts

