

Fairness, Impartiality and the Self-Serving Bias

Abstract: Garrett Cullity claims that his formal account of fairness has practical application.¹ For example, appeals to fairness can limit the extreme moral demand proposed in Peter Singer's 1972 seminal paper, 'Famine, Affluence and Morality'. I suggest that Cullity's account lacks the specificity required for successful application in complex cases of this kind. A good theory of fairness must provide enough specific guidance about what fairness is for a conscientious application of the theory to sharply limit room for bias to influence conclusions about that fairness requires. However, Cullity's account lacks this specific guidance. I draw on scientific studies in psychology that explore the impact of the self-serving bias on judgements about fairness to argue that judgements made according to Cullity's account of fairness are *particularly* vulnerable to the self-serving bias.

Introduction

Garrett Cullity claims that all unfair actions share a common feature (2008: 3): 'Unfair actions are *failures of appropriate impartiality*; fair actions are those that are not unfair.' Although Cullity offers this statement as a formal, not substantive, account of fairness, he finds practical application for the formal account in two cases: (1) an appeal to fairness limits the amount of resources the affluent must commit to alleviating global poverty (2004: 167-187) and (2) fairness justifies the use of force to compel an agent to contribute towards the production of public goods, even if the public good is one that the agent has been given without asking for it or seeking it out.

It is reasonable to allow that what fairness requires is very often controversial and also that various biases are among the causes of people's coming to different conclusions about what would be fair in particular cases or classes of case. Acknowledging these two things, we would expect a good theory of fairness to provide enough specific guidance about what is fair, for conscientious application of the theory to sharply limit room for bias to influence conclusions about what fairness requires in a case or kind of case.

To say this is not to deny that a good theory of fairness will have to leave room for judgement, nor that even a very good theory of fairness may be difficult to apply in many

¹See Cullity (2004), chapters 8 and 10 and Cullity (2008).

cases. Moreover, conscientious applicers of the theory might reach different conclusions sometimes. But admitting those three things is entirely compatible with thinking that a good theory of fairness must provide enough specific guidance about what is fair for conscientious application of the theory to sharply limit room for bias to influence conclusions about what fairness requires. In this paper I claim that Cullity's account of fairness does not have enough specific guidance in it to inhibit the force of the self-serving bias.

A psychological mechanism that has received attention from experimental psychologists is the self-serving bias in judgements about fairness. This bias is a tendency to conflate what is fair with what benefits oneself (Babcock and Loewenstein, 1997: 109).² Recently psychologists have focused on the self-serving bias as a cause of negotiation impasses (B&L, 1997: 109-126) and as a factor in disagreements about the solution to climate change (Kriss et al., 2011: 601-615). In this paper, I wish to bring the problem of the self-serving bias into sharp focus.

Disputes about fairness often arise in complex cases because the parties to a judgment about fairness do not agree about which factors are most important. In more complex disputes, the self-serving bias is very likely to influence people's judgements and result in heartfelt disagreement. It is therefore very important not to add further layers of complexity to the process of establishing what fairness requires in any given situation, since doing so will further increase the likelihood of biased judgements and conflict.

In section one, I start with a brief description of Cullity's formal account of fairness. In section two, I provide a survey of the psychological literature on the self-serving bias. In section three, I return to Cullity's account and demonstrate how it lacks specific guidance and renders the process of judging what is fair far more complicated than necessary. I also show that, rather than limiting the room for biases to influence judgements about fairness, the unnecessary complexity instead introduces opportunities for the self-serving bias to operate on judgements about fairness.

1. Cullity's Account of Fairness

Cullity (2008: 3) proposes a general account of *unfairness* in actions: all unfair actions are failures of appropriate impartiality. I will argue that Cullity's account encounters a serious problem: the term 'appropriate impartiality' introduces so much scope for interpretation that the self-serving bias is very likely to infect judgements about fairness. Cullity offers a formal

² Hereafter, B&L.

account of fairness and applies that account to free-riding, to demonstrate that free-riding is unfair. Cullity then claims that refusal to pay for certain public goods can also be wrong, even when the good is compulsory.³ Cullity's arguments about free-riding are relevant to the demandingness and global poverty debate since the failure of others to contribute, which we might call widespread free-riding on the more generous natures of others, is a crucial factor in making morality so very demanding. Like Singer, Cullity holds that the moral failures of others do not limit the moral obligation of the contributors. However, unlike Singer, Cullity believes that extreme moral demandingness is mistaken because of reasons of impartiality. Since Cullity's account equates fairness with impartiality, his account implies that *fairness* cannot require the affluent to comply with 'Singer-style' demandingness.

We must now ask what fairness consists in. Cullity claims that unfair actions are failures of *appropriate impartiality* and that fair actions are those that are not unfair.⁴ He argues for his claim by pointing to the number and variety of different occasions on which judgements about unfairness correlate to failures of some kind of impartiality. Cullity's account draws on a familiar pre-theoretical view: fairness is tightly connected with impartiality.

We cannot plausibly hold that just any impartial way of making a decision would be a fair way of making that decision. For example, rolling a die is undeniably an *impartial* way to decide who receives the top grade on a paper, but it is not a *fair* way to allocate grades. Allocating grades to papers fairly requires more than simple impartiality; it requires the marker to be insensitive to students' aesthetic qualities, accent, origin, etc., but sensitive to differences in the qualities of the papers being graded. Rolling the die bypasses the evaluation appropriate for allocating grades. To employ Cullity's notion of unfairness one needs a nuanced understanding of the ways in which impartiality might be appropriate.

2. The self-serving bias

Some experiments demonstrate that subjects tend to rate themselves as 'above average' in a number of different areas, including ethical behaviour, managerial ability and productivity

³ Compulsory goods are public goods that, once produced, provide a benefit for everyone regardless of whether the beneficiaries wish to receive/consume the good. It is not usually possible to avoid consuming the good or to prevent people from consuming it. Examples of compulsory goods are law and order, clean air and street lighting.

⁴ For the purpose of this paper, I set aside the fact that many actions are neither fair nor unfair. I also note that partiality and impartiality are contraries, not contradictories. A failure to act impartially does not necessarily mean that one acts partially; partiality is irrelevant to certain actions.

(B&L, 1997, 110-111). The self-serving bias is also evident in other biases, such as the attribution error, whereby subjects tend to attribute their successes to their own efforts, skills or ability while attributing their failures to bad luck (B&L, 1997: 111). The self-serving bias affects many people's evaluations of not only themselves as individuals but also the groups of which they are members (B&L, 1997: 111).

The specific area of research into the self-serving bias that is important for my purposes here, relates to how people make systematic errors when making judgements about fairness, i.e., when people think that what is beneficial to themselves is also *fair*. Babstock and Loewenstein cite a 1979 study in which David Messick and Keith Sentis divided subjects into two groups and constructed an experiment using the following three stipulations: (1) Group 1 are told to imagine they each completed a task in 10 hours, (2) Group 2 are told to imagine that they completed the task in 7 hours, and (3) both groups are then told that the rate paid for completing the task in 7 hours was \$25. The findings were that (1) on average, those who were told to imagine they each completed the task in 10 hours suggested a rate of pay of \$35.24 for 10 hours work to complete the task, and (2) on average, those who were told to imagine they each completed the task in 7 hours suggested a rate of pay of \$30.29 for 10 hours work to complete the task. The empirical findings of the experiment suggest that the self-serving bias influenced people's judgements about what a fair rate of pay would be. I discuss the Messick and Sentis experiment further in Section Three in connection with Cullity's account.

I now turn my attention to two psychological studies with particular relevance to philosophy. The first, from Babcock and Loewenstein, suggests that the self-serving bias is a major cause of negotiating impasses. Imagine a dispute between two people where each party has identified a position that she thinks is a fair resolution of a dispute. If the self-serving bias is present, then each party's judgement about what is fair will be adversely influenced by what she assumes or infers would be most beneficial to herself. The self-serving bias can operate surreptitiously and its influence on judgements about fairness is often unnoticed. Especially where this happens the parties to the dispute will see themselves as driven by a conception of fairness, rather than self-interest. Nevertheless, if the self-serving bias is present and if the self-interests of the parties push in opposite directions, then the parties will be identifying different resolutions of the dispute as fair, and each will think the other's proposal is unfair.

Second, the research indicates that, since each believes she is offering a fair resolution, she may potentially view the other's proposal as not only unfair but also aggressive or exploitative. Finally, further research demonstrates that negotiators exhibit resistance to settling for anything less than what they think is fair. So, the bias to think that what is beneficial to oneself is also fair means that people are, understandably, resistant to settling for anything less than what is beneficial to themselves, not because they identify with the fact that the solution benefits them, but because they believe that they are being fair (B&L, 1979: 110). These findings are supported even when subjects share all the same information, suggesting that the self-serving bias, rather than uncertainty, is the primary cause of negotiation impasses.

A philosophically relevant observation is that, even when the self-serving bias is present, subjects tend to think not only that their proposed settlement is fair but also that (B&L, 1997: 120), 'their personal conception of fairness is impartial.' I will demonstrate in the next section how this poses a genuine difficulty for an account of fairness as appropriate impartiality. Cullity explicitly states (2008: 5) that he does not offer a substantive account of fairness, choosing to leave open both 'what ought to be done' and 'what forms of impartiality are appropriate to doing it.' The problem is that, while the practical application of such an account relies heavily on interpretation, the self-serving bias tends to make different people select different forms of impartiality and thus to interpret fairness differently and thus to feel that the positions they take are fair and the different positions of others are unfair.

Finally, Babstock and Loewenstein (1997: 120) report evidence of a different effect of the self-serving bias: this bias tends to lead one to think that what benefits oneself is *morally correct*. That the self-serving bias can be seen to affect both judgements about fairness and moral judgements shouldn't be surprising. The fairest action may not be the morally best action, all things considered, but we nonetheless tend to view fairness as something that counts in favour of an action. Moreover, if we combine the evidence that subjects with the self-serving bias think that their conception of fairness is impartial with another strongly held pre-theoretical view that impartiality and moral correctness are closely bound, it becomes clear that those who think they are acting fairly might also think that, in being impartial, they are also acting correctly.

Further investigation into the effects of the self-serving bias comes from Kriss et al. Climate change strategists face the challenge of determining what level of reduction in

emissions will be both effective and spread the emission reduction burden fairly. Kriss et al maintain that their research demonstrates that the self-serving bias ‘plays a major role in the difficulty of obtaining agreement on how to implement emissions reductions.’ Research revealed that subjects from the USA and China had very different ideas about what was ‘fair’ and that these differences reflected what was in the national interest of each country. How can we be sure that these subjects were really interested in a fair result rather than one that promoted their country’s interest? First, the subjects were not people in a position to determine or negotiate policy about climate change; they were simply asked to make a judgment about what fairness required. Second, the self-serving bias had an effect when the subjects in the experiment were affiliated to the countries under consideration, but the bias was not evident when there was no affiliation between the experimental subjects and the countries under consideration. When there was no affiliation between the experimental subjects and the countries under consideration, the subjects made the same judgements about what was fair. Hence this set of experiments showed that subjects reach similar judgements about fairness when the self-serving bias isn’t operating, but different fairness judgements when it is operating.

In the climate change debate, different judgements about fairness animate different policy proposals. One proposal is that how much CO₂ each country should be allowed to emit is proportional to the size of that country’s *population*, and thus the emissions per person in an economically developed country should be neither more nor less than the emissions per person in a economically developing country. Another proposal is that how much CO₂ each country should be allowed to emit is proportional to the size of that country’s *GDP*, and thus the emissions per person in an economically developed country should be allowed to be greater than the emissions per person in a economically developing country. Yet another proposal is that the emissions per person in a developing country should be allowed to be greater than the emissions per person in a developed country, because of the greater economic need in the developing country or because of the economic inequality between developing and developed countries. The relevance here of the dispute among such different views is that there is a correlation between the views of fair climate change policies that different countries espouse and the self-interest of the countries. In the global poverty debate, there is similar scope for a very wide interpretation of what is required from the affluent. The more varied the array of requirements of fairness or of morality, the more scope there is for distortion by the self-serving bias.

3. Cullity's vulnerability to the self-serving bias.

The psychological literature demonstrates that the bias can only operate when there are a number of standpoints from which a possible negotiation can be viewed. It is therefore important that a good account of fairness provides detailed and specific guidance, sufficient to rule out bogus standpoints. Cullity's account, which holds that fair actions are those in which the *appropriate* form of impartiality is used, can only limit the effect of the self-serving bias if it provides enough information for agents to identify the right kind of impartiality: the kind that will be appropriate in each different case.

But Cullity offers no general method by which to determine the type of impartiality that would be appropriate in particular instances of unfairness. He claims instead that what determines the appropriate kind of impartiality is an independent question answerable only within a more general normative theory (Cullity, 2008: 5). He offers only the following: different distributive practices are governed by their different aims. In each instance, the aim of the practice is to determine what kind of impartiality is appropriate (Cullity, 2008: 4).

Let us acknowledge that sometimes it is easy to identify the aim of the practice. In a 100m sprint race, the aim is to give people in a certain group the opportunity to prove that they can sprint faster over 100m than others in their group. Where identifying the aim of the practice is easy, identifying the appropriate form of impartiality is usually also straightforward. Again, consider a 100m sprint race. First, there are details about which facts about competitors are relevant in assessing their eligibility. If the race is the 'Boys U13 100m sprint', race officials are not flouting the requirements of impartiality when they exclude all females and anyone over 13. Once the qualification criteria are met, and race conditions and rules applied equally to all entrants, there is only one criterion for the judge to apply, namely the relative positions of the contestants as they pass the finishing line. Being appropriately impartial in the race example requires being uninfluenced by facts such as who put in the most effort or who would benefit most from a win or who is kindest or who would look most attractive on the podium. As Cullity suggests, there are other forms of impartiality that could be employed, such as pulling names out of a hat, but that form of impartiality is, quite obviously, inappropriate in cases such as the race example.

Exploring our intuitions about games may be a useful starting point for philosophical examples, but their use is limited. In sports related, rule-governed examples it is often clear what kind of impartiality is required and, in contrast, what kind of impartiality would be

inappropriate. And it is important to note that in practices where the appropriate form of impartiality is easy to identify, it is likely to be easy to identify whether judgments are fair or unfair. However, in more complex cases we need to know whether impartiality is a requirement at all and, if so, the kind of impartiality that applies.

Recall the Messick and Sentis experiment in which subjects in two groups were instructed to determine a fair rate of pay for working on a task. The aim for both groups was the same: reward people fairly for their work on a task. The scope for interpretation arose because their personal circumstances were different to the worker with whom they were to compare themselves in deciding what would be a fair reward. Group One had taken fewer hours to complete the task. It was therefore in their interests to believe that fairness requires a completion-related reward. Group Two had taken more hours to finish their task so it was in their interests to believe that fairness requires a time-related reward. But note that both groups think that they are acting fairly, not out of self-interest. This example illustrates a problem for Cullity's account of fairness: despite there being one aim, namely to reward people fairly for working on a task, the two groups did not converge on a single fair (or appropriately impartial) way in which to reward workers. There is a variety of different ways of determining pay rates. The different rates, either per hour or per job, are familiar and well respected. When groups disagree about which of these rates should be used, there is no reason to accuse either group of a failure of appropriate impartiality. Cullity's account can offer no guidance once such an impasse is reached.

I wish to stress that I am not setting out to reject Cullity's *formal* account of fairness. Cullity identifies something interesting concerning judgements about fairness, namely that they all involve some form of impartiality. The form of impartiality is likely to vary from case to case because, as Cullity says, there are indefinitely many different ways of being impartial (2008: 3). However, a significant problem for Cullity's account of fairness is that we want considerations of fairness to be effective in resolving conflict, and yet Cullity's account provides too much opportunity for the self-serving bias to induce people to focus on different forms of impartiality as the basis for judgements about fairness, and thus to reach different conclusions about what fairness requires. The application of Cullity's account of fairness can lead to conflict rather resolve it. Without guidance as to which kind of impartiality is appropriate in a contestable context, we cannot use Cullity's account to identify which judgements are unfair and which are fair.

Let us move away from simple game structures to more complicated judgements about fairness, using a richly detailed example. I will demonstrate that indeterminacy arises both in terms of what *form* of impartiality ought to be applied and, since impartiality is scalar, what *degree* of impartiality ought to be applied. However, only *one* form of impartiality will be appropriate, according to Cullity (2008: 3): ‘Judgements about fairness and unfairness, I claim, concern actions for which *one particular way* of being impartial is morally required.’⁵

3.1 *The photographer*

David is a keen amateur photographer and his good friends know that he is trying to break into professional photography in order to supplement his low income. David’s good friend Jenna asks him to photograph her wedding and David agrees. Unfortunately David is inexperienced in business. As a result, David and Jenna do not agree any details prior to Jenna’s wedding, including Jenna’s expectations for the day or a fair price for the job. Jenna simply says that she will pay him to take ‘candid shots’ throughout the day. At the wedding, David takes over five hundred photographs, which he spends many hours editing before he calls Jenna to set up a meeting to show her the photographs on his laptop. Claiming to be busy, Jenna asks David to put the photos on a disc and post it to her. David worries that Jenna may not pay him. Rather than post a disc, David holds out for a meeting and eventually shows Jenna the pictures, leaving her with an invoice for £400, less than one-third of the price that local professional photographers would charge. Jenna refuses to pay, claiming that the photographs were far more elaborate and plentiful than she wanted and that she thought a friend would not try to run up the cost.

Let’s set to one side the fact that prudentially there are many things that David and Jenna should have done. This kind of dispute, where friends have different expectations of each other and are reluctant to start out on a professional basis, is very common. This is also an example to which there is no obvious answer. We cannot simply say that this is a legal issue, since the two friends didn’t have the foresight to make an unambiguous contract. We can’t resort to whatever the market rate would be. In this type of case, there isn’t an obvious market rate for a part-time photographer who has skill but no portfolio, and there is no universally acceptable discount rate for friends and family.

⁵ My emphasis.

For the sake of argument, let's agree that the quality of photographs was good, even if the number and style was a point of disagreement. The following factors are relevant considerations of fairness. First, there was an informal agreement that David would do a job for Jenna and that Jenna would give him something in return for it. Despite the absence of a formal contract, we recognise that some degree of reciprocity exists between David and Jenna. Next, consider the obligations of friendship. David thought that Jenna was helping a friend out by employing him; he thought that he was helping her out by giving her so many photographs at a rate she would not be able to get elsewhere. Jenna thought that she was helping David by giving him experience and an opportunity to develop a portfolio, but she also expected that, as a friend, he would charge only a nominal sum or maybe even give her the photographs as a wedding gift. Finally there is a question of desert. David put much time and effort into taking and editing the photographs and the quality was good. Jenna would reply that he put more time and effort into taking the photographs than she requested.

One form of impartiality must be that each ignores their personal desires about the price of the photographs. That David wants to be paid £400 for the photographs and Jenna wants them for free are irrelevant to the question of what would be a fair price. So we turn to the requirements of reciprocity, friendship and desert. Even if David and Jenna agree that reciprocity, friendship and desert should be part of determining a fair price for the photographs, they may yet have different ideas about what reciprocity, friendship and desert require, and about which of these considerations, if any, are the most important.

Jenna thinks that if Friend A provides a service to another Friend B, then reciprocity requires that Friend B should, at the earliest opportunity, provide a service to Friend A. David thinks that if Friend A provides a service to Friend B, then reciprocity requires that Friend B should pay for the service. As we have seen, if David and Jenna are under the influence of the self-serving bias, they will both think that they have acted fairly and impartially. As readers, and privy to what is going on, we can see that one solution favours Jenna and one favours David. But without that information, we might be willing to allow that both appear to be reasonably appropriate forms of impartiality.

If David and Jenna are going to apply an appropriate form of impartiality to the degree of influence that friendship should have, they need to abstract their judgment from their own friendship and ask what friendship in general requires. But is that really fair? Doesn't the strength of friendship vary from case to case? Might fairness require one

judgment in a dispute between lifelong friends and another between recently acquainted but enthusiastic friends? Perhaps David and Jenna should ask what other, similar, friendships would require. I doubt there is a clear answer to this question but perhaps both could agree that friendship in general would require that both friends should make a degree of sacrifice. In the actual situation, this seems less likely. Often when friends have work done at ‘mates’ rates’, the labourer sacrifices some income and the person who is paying gets a lower price than the ‘arms-length’ rate.

Finally, considering desert, David and Jenna might agree that productivity and effort should be rewarded. In the actual case they are unlikely to agree about the degree to which productivity and effort should be rewarded. Perhaps, if they thought about their situation impartially, they could see that Jenna was disingenuous when she said she just wanted twenty snaps and David was over enthusiastic to take five hundred images. Maybe some research would show that, typically, wedding photographers provide n images of a wedding. Unfortunately, the Babcock and Loewenstein’s research suggests that neither is likely to identify a number of photographs that sounds fair or impartial to the other.

This example of David and Jenna demonstrates that there are various forms that appropriate impartiality might take. It might be more appropriate to reward desert and fulfil the requirements of reciprocity in one particular way than to comply with the requirements of friendship. Or it might be appropriate to comply with the requirements of friendship to some degree but to give similar weight to matters of desert. But the manner in which David and Jenna are left to establish what might be appropriate is fraught with difficulties and it is exactly cases like these in which the self-serving bias is most likely to be present yet not obvious to the parties involved, as I will demonstrate in §3.3.

Recall Cullity’s advice that the *aim of the practice* will determine the form of impartiality that is appropriate to apply. In the photographer example there is no clearly identifiable practice, certainly not with an aim. A significant number of disputes about fairness are like the photographer example in that they involve multiple kinds of considerations relevant to fairness considerations, attention to which can be more or less appropriate. One might think that this complication is limited to informal disputes and that the issue of fairness need not involve such indeterminacy when we consider established practices. However, established practices very often contain or develop indeterminacies. So the problem I have focused on has very wide scope.

3.2 *Oxfam*

Oxfam clearly states three aims: ‘We aim to save lives by responding quickly with aid and protection during emergencies, empower people to work their own way out of poverty, and campaign for lasting change.’⁶ Oxfam face the problem of limited resources, increasing the likelihood of disputes over the fair way to allocate resources. Which one of Oxfam’s aims dictates the form of impartiality appropriate to determining a fair distribution of aid resources? When is it fair to direct resources to campaigning for long-term changes that will benefit future generations at the expense of the suffering of existing generations? If Cullity’s account is to generate an unambiguous indication of what fairness requires, Oxfam must have a single overarching aim, or at least a strict ranking of the three aims previously identified.

To keep matters simple, let’s say Oxfam’s overarching aim is to satisfy the needs of the worst off. How do we get from this aim to the form of impartiality that is appropriate? The problem this time is not that the aim itself is hard to identify but rather that it’s unclear what kind of impartiality will lead to a fair judgment. For example, is it more important to be completely impartial regarding the location of the neediest people and seek out the neediest first and then the second neediest and so on? Doing so will involve enormous costs in terms of research and transportation. Is it more appropriate to identify an area where a large number of needy people co-exist and satisfy more needs, even though some people in more desperate need elsewhere will die as a result?

The problem for Cullity’s account of fairness is that it already raises more questions than it provides answers. Since the aims of many institutions may be unidentifiable, subject to debate or open to interpretation, there will be times when identifying the appropriate form of impartiality is an unworkable aim. If fairness is appropriate impartiality, our confidence in fairness will be undermined. This problem is exacerbated by the finding that self-serving bias has more influence when agents identify a range of potential solutions.

3.3 The self-serving bias and complex solutions

Psychological research demonstrates several philosophically relevant aspects of the self-serving bias. The self-serving bias is not common in very simple situations. If two subjects are asked to split a sum of money fairly between themselves and another candidate then, all things being equal, they are likely to make the same proposal, typically 50:50. However, the

⁶ <http://www.oxfam.org.uk/what-we-do/about-us/our-vision-values-and-goals/our-aims-and-approach>

self-serving bias becomes more evident as situations become more complex. Factors that increase complexity could be additional information about the situation or further criteria to be applied.⁷

For example, the subjects might be told to divide a sum of money fairly between themselves and a very rich person or between themselves and a poor but reckless person. The additional information about the other person provides opportunity for the subjects to introduce new criteria for their judgment of fairness. Perhaps they should take into account the other person's lack of need or irresponsible attitude.

Admittedly, subjects who think that a fair distribution of the sum requires taking need or desert into account *may* be employing an appropriate form of impartiality; on the other hand, they may simply be influenced by the self-serving bias. The problem is that in disputes there is usually some vagueness about which criteria are relevant or about how to interpret relevant criteria. The instruction to divide a sum of money fairly between oneself and a very rich person or oneself and a poor but reckless person is not clearly stipulated to take into account need in one case but deservingness in another. This vagueness, combined with layers of complexity, increases opportunities for the self-serving bias to affect judgements.

If Cullity's account offered clear guidance about the forms and degrees of impartiality that would be appropriate in various circumstances, he could reduce the incidence of the self-serving bias. But Cullity (2008; 5) does not set out to give such guidance:

[A] formal, not a substantive account of unfairness, ... does not itself tell us what ought, all things considered, to be done, nor what forms of impartiality are appropriate to doing it. And it leaves open what reasons ultimately settle those questions. Our survey suggests these reasons will be various. Although I have offered a simple and general unitary description of unfair action, I offer no simple and general unitary explanation of when and why different forms of impartiality are appropriate.

⁷ In the psychological literature, the kinds of situations I refer to as 'complex' are called 'asymmetrical'. However, they are not asymmetrical in the sense that one party has more information or more options than another. Rather they are cases in which there are a range of possible solutions, some of which favour one agent; some of which favour the other. The range of solutions generated is greater in complex cases such as those I have identified. For example, one agent focuses on desert, but another focuses on reciprocity. Where there is a greater range of solutions, the self-serving bias is likely to have a more damaging effect.

The only guidance offered, as I have mentioned, is that sometimes these appropriate forms of impartiality will be dictated by the aim of the relevant social practice. But some social practices don't have an aim at all, while other social practices have multiple aims with no clear priority ranking. And the aim or aims that some social practices have might need considerable interpretation. In all of these cases, the self-serving bias can worm its way in and influence the selection of a form of impartiality, and thus influence judgement about what is fair.

This must be a problem for Cullity, who claims that his account has practical application to redistributive practices—such as the highly disputed extent to which global poverty makes moral demands on the affluent. We have seen that complex judgements about fairness are more likely to be affected by the self-serving bias and that Cullity's procedure for making judgements about fairness is particularly complex. I will now demonstrate that when Cullity himself uses fairness to limit Singer's extreme moral demand, his conclusion appears to be affected by the self-serving bias.

Cullity aims to replace Singer's extreme demand with a less demanding principle. His strategy involves two stages. The first stage employs an argument of beneficence to demonstrate that morality cannot require us to live a purely altruistically focused life. The second stage aims to demonstrate that contributions towards global poverty are limited by the requirements of fairness. I maintain that the second stage of the argument is affected by the self-serving bias.

First, Cullity suggests that beneficence—a practical concern for the interests of others—requires us to save lives for reasons other than the simple benefit of being alive. He writes, (2004: 135):

However [the non-instrumental good of life] is not the only answer, nor, *for most of us*, the most important one. What is more important, *for most of us*, is that our lives are vehicles for the fulfilments that a well-lived life can contain.⁸

These fulfilments include friendship, achievements in worthwhile personal projects and participation in a culture. An important step in Cullity's argument is that reasons of beneficence presuppose the moral permissibility of life-enhancing goods for all agents

⁸ My emphases.

because, (Cullity, 2009: 23-24) Agent *A*'s interest in having her life saved can only ground a requirement on Agent *B* to save it as long as it is not something it is morally wrong for Agent *A* to have.⁹ The steps to Cullity's argument against Singer's extreme demand are as follows:

Premise 1: People who pursue worthwhile personal projects cannot also be living an altruistically focused life unless their personal project happens to be to live an altruistic life, since they will need to spend some time and money on the pursuit of these projects.

Premise 2: Even if one person's life goes better because they freely choose to live altruistically, fairness, as appropriate impartiality, dictates that it is not wrong for other people to pursue their own worthwhile personal projects.

Premise 3: Since I should keep donating to aid agencies until the point where making a further contribution would require a comparable moral sacrifice, it is immoral for me to pursue a non-altruistically focused life.

These premises lead to the contradiction. Either *it is wrong* to lead a non-altruistically focused life or fairness dictates that *it is not wrong* to do so. Cullity's argument proceeds:

Premise 4: My reason to save other people's lives is grounded in their (non-altruistically focused) interests.

Premise 5: Morality cannot require me to help any agent to do or obtain something that it is immoral for her to do or possess.

Conclusion: Either I am not morally required to pursue an altruistically focused life in order to save others, *or* both the person whose life is to be saved and I must be allowed to pursue non-altruistically focused lives.

But Cullity's argument is, I believe, an instance of the self-serving bias in action. As such, it must fail. The bias is more easily revealed if we consider Cullity's argument in a slightly different format.

⁹ Cullity G., (2009), 'Demandingness and Arguments form Presupposition', in Chappell T., (ed), *The Problem of Moral Demandingness. New Philosophical Essays*. Basingstoke: Palgrave MacMillan.

Premise 1: Since I should keep donating to aid agencies until the point where making a further contribution would require a moral sacrifice comparable to the suffering I am trying to alleviate, it is immoral for me to pursue a non-altruistically focused life.

Premise 2: Whenever I comply with Singer's extreme demand, I am (almost always) helping someone to live a non-altruistically focused life.

Conclusion: Whenever I comply with Singer's extreme demand, I am (almost always) helping someone to obtain something that it is immoral for her to possess.

Premise 3: Morality cannot require me to help any agent to do or obtain something that it is immoral for her to do or possess.

Conclusion 2: Morality cannot require my compliance with the extreme demand.

The second premise depends on Cullity's insistence that *for most of us* non-altruistically focused lives are the main source of value, rather than the simple benefit of being alive. In contrast, Singer's argument starts from the premise that 'suffering and death from lack of food, shelter, and medical care are bad' (Singer, 1972: 231). While, famously, Singer does not value any life at all over one of extreme misery and unbearable suffering, he does appear to value life for its own sake. I don't mean to suggest that a pleasant life is not more valuable than a very basic subsistence but I question Cullity's assumption about the kinds of life that might be available to '*most of us*'. As I will shortly demonstrate, this assumption guides Cullity in his choice of what kind of impartiality is appropriate in determining our moral obligation to donate to aid programmes. The assumption also leads Cullity to choose a form of impartiality that favours the affluent over the global poor.

The assumption that non-altruistically focused lives are the main source of value for most of us is foundational in the next step of his argument, in which Cullity turns to fairness as appropriate impartiality in order to determine what kinds or degrees of non-altruistically focused lives are morally permissible. Cullity notes that life-enhancing goods 'essentially involve *personal partiality*'.¹⁰ The suggestion is that any agent, whether affluent or poor, can only benefit from these life-enhancing goods if she nurtures *her* friendships, participates in *her* culture and pursues those worthwhile projects that are important to *her*. Despite this

¹⁰ My emphasis.

personal partiality, provided every agent is permitted to pursue such goods, then a form of complete impartiality obtains, since the pursuit of these goods is not restricted by sensitivity to any personal features of an agents such as their religion, age, location, race or gender; every agent is impartially permitted to pursue personally partial, life-enhancing goods. Therefore, Cullity claims that fairness (understood as appropriate impartiality) dictates that neither the poor nor the affluent are acting immorally in pursuing non-altruistically focused lives. Agents only act immorally if they fail to contribute to poverty alleviation whatever resources are not required for the pursuit of their own life-enhancing projects.

Cullity's solution is one that all moral agents should be able to apply to their own lives. He suggests that we periodically review how much we have spent on ourselves and how much we have given to poverty alleviation and adjust our future spending and aid contributions in the light of our findings. This exercise is one of making a judgment about fairness. It requires that each agent can judge what is fair for her to retain to pursue her own life-enhancing goods, knowing that fairness permits both the affluent and the global poor to pursue partial life-enhancing goods.

But this is too quick. Consider Cullity's discussion of the kinds of goods upon which that the affluent may permissibly spend their resources. These goods include: 'the kind of expensive tertiary education that is available to relatively few people, globally speaking. For it is almost always reasonable to believe that this will substantially advance a person's understanding achievements and participation in a culture, throughout her life.' Similarly, Cullity allows for expenditure on holiday travel, not justified by the pleasure that holidays typically bring, but because they are an opportunity to bond with family. Moreover, expensive music lessons for the talentless might be unjustifiable in terms of the music they end up producing yet justifiable in terms of their participation in a culture: 'So expensive music lessons (or instruments) for the talentless seem hard to justify. However, up to a point, it seems sensible to see tuition in music, drama, and other arts as participation in a culture'.

I don't dispute that the pursuit of tertiary education, holidays and means of participating in culture are valuable goods. Rather, I dispute the *fair and impartial status* of Cullity's conclusion that expenditure on these kinds of goods legitimately limits the contributions we ought to make to alleviate global poverty. Such a conclusion ignores the reality that many inhabitants in developing countries don't have the means to feed themselves adequately, let alone pursue meaningful personal projects. Suppose the life we can save is

‘merely’ one containing access to adequate nutrition, clean water, shelter, and protection from a range of basic threats such that the individual can work to provide for her family in future years. Suppose further that family bonding is restricted to the times spent preparing and eating the family meal and that participation in culture requires arduous work tending crops or digging wells. Suppose that education is basic. This life, devoid of holidays, rich cultural activities and tertiary education is nonetheless a life worth saving. Cullity spends the first half of his (2004) *The Moral Demands of Affluence* defending the moral obligation to save the lives of the global poor. He would agree that such a modest life *is* a life worth saving.

It seems likely that the self-serving bias affects Cullity’s arguments twice. First, I believe it leads Cullity to apply the wrong kind of impartiality to his method of assessing relevant factors for his argument. To 1.6 billion people living in absolute poverty, a non-altruistically focused life in which basic needs such as food, shelter, medical care and dignity are met will be an enormous source of value. That their lives might be better still with more goods does not legitimise those goods for the affluent. The global poverty crisis is so severe that lifting people above the poverty line is in itself an unachievable task, given the current level of donations. It is so severe that morality is extremely demanding, on Singer’s view, even if the project is limited to providing the world’s poorest people with basic subsistence goods. Until we reach the point where the purpose of aid is to lift people far enough above the poverty line that they can pursue similar projects and activities to the kinds pursued in developed countries, no form of impartiality that relies on the value of partial goods is appropriate in the sense that Cullity needs it to be for his account of fairness.

But a time may come when we do have a more ambitious aim for aid. In that case, Cullity’s appeal to the impartial moral status of the pursuit of partial goods would be legitimate. But the self-serving bias has a second opportunity to distort judgements about what fairness requires since Cullity leaves the decision about what level of spending ought to be justifiable in the hands of each individual agent. Since the self-serving bias is so hard to detect, even sincere agents are likely to be mistaken in their choices. An appeal to fairness or impartiality to justify a holiday for the purposes of family bonding may be an instance of a judgment of fairness that has been affected by the self-serving bias. The claim that expensive music lessons for the talentless may be justified because music enables participation in a culture may be another judgment of fairness affected by the self-serving bias. Cullity is

confident that we will be able to successfully assess the impact that different levels of spending will have on our lives. However, I do not share his optimism. The evidence concerning the incidence of the self-serving bias suggests that even if we sincerely try to make impartial judgements about the fairness of our spending levels, we are likely to make them in a self-interested manner.

5. Conclusion

I have noted that Cullity's formal account of fairness may be correct in the abstract but I have argued that it allows so much scope for interpretation that judgements about fairness made according to Cullity's account are likely to be adversely influenced by the self-serving bias. As a result, consensus about what fairness actually requires will be difficult to attain. I have cited Cullity's own application of his account as evidence of the self-serving bias in operation. Without some method of both identifying and eliminating the bias in operation, I maintain that an account of fairness as appropriate impartiality will be limited in application to only those situations in which it is clear both what form of impartiality is most appropriate and to what degree impartiality should apply. Those cases, I maintain, are scarce and unlikely to feature in disputes about fairness.

Bibliography

Babcock L, and Loewenstein G., (1997), 'Explaining Bargaining Impasse: The Role of Self-Serving Biases', *The Journal of Economic Perspectives*, 11:1, 109-126.

Cullity G., (2004), *The Moral Demands of Affluence*, Oxford: Oxford University Press.

Cullity G., (2008), 'Public goods and fairness', *Australasian Journal of Philosophy*, 86:1, 1-21.

Cullity G., (2009), 'Demandingness and Arguments form Presupposition', in Chappell T., (ed), *The Problem of Moral Demandingness. New Philosophical Essays*. Basingstoke: Palgrave MacMillan.

Kriss P, Loewenstein G, Wang X, and Weber R., (2011), 'Behind the veil of ignorance: Self-serving bias in climate change negotiations', *Judgment and Decision Making*, 6:7, 602-615.

Singer P., (1972) 'Famine, Affluence, and Morality', *Philosophy and Public Affairs*, 1, 229-43.