

PUBLIC GOODS AND FAIRNESS

Garrett Cullity

To what extent can we as a community legitimately require individuals to contribute to producing public goods? Most of us think that, at least sometimes, refusing to pay for a public good that you have enjoyed can involve a kind of ‘free riding’ that makes it wrong. But what is less clear is under exactly which circumstances this is wrong. To work out the answer to that, we need to know why it is wrong. I argue that when free riding is wrong, the reason is that it is unfair. That is not itself a very controversial claim. But spelling out *why* it is unfair allows us to see just which forms of free riding are wrong. Moreover, it supplies a basis from which some more controversial conclusions can be defended. Even if a public good is one that you have been given without asking for it or seeking it out, it can still be wrong not to be prepared to pay for it. It can be wrong not to be prepared to pay for public goods even when you do not receive them at all. And furthermore, it can be right to force you to do so.

When can we properly require individuals to contribute to our collective projects? This question can be read different ways. The most important version of it is also the most difficult to answer. Some goods—goods such as fresh air, national defence, and public health—are ‘compulsory’, in this sense: once they are produced, we cannot avoid receiving them without excessive cost. The question when people can be ‘required’ to pay for such goods might be read as asking when it is wrong for them not to pay. That is one issue I shall be addressing in what follows. But my overall interest is in this further, and harder, question:

When can someone for whom a collectively produced good is compulsory properly be compelled to contribute towards its production?

This question is important, since answering it might offer the key to explaining when the coercive power of a political state may legitimately be used against its citizens. There has been a recent revival of interest in the project of justifying political obligation and authority by appealing to a ‘principle of fairness’.¹ The overall aim of this paper is to give foundational support to that project.²

¹This is Rawls’s label for the following principle [1971: 111–12]:

a person is required to do his part as defined by the rules of an institution when two conditions are met: first, the institution is just (or fair), that is, it satisfies the two principles of justice; and second, one has voluntarily accepted the benefits of the arrangement or taken advantage of the opportunities it offers to further one’s interests.

I start by examining a relatively uncontroversial claim. It is widely accepted that ‘free riding’ by actively taking collectively produced goods without paying can be wrong. Moreover, many of us agree with Rawls [1971: 111–12] about why this is wrong (when it is): free riding is wrong because it is unfair. However, there is less agreement about exactly when this is true. To make progress, we need to dig deeper. If we could explain *what makes* free riding unfair, we would then be better placed to say when it is, and when it is not. This occupies Sections I–III.

Section IV turns to the more controversial case of compulsory goods—goods that are not actively taken. I argue that not paying for these can be unfair too, because it can share the feature that makes the more obvious cases of free riding unfair.³ Indeed, it can be wrong not to pay for some public goods (goods we possess jointly, in certain respects I shall explain) even when you do not receive them at all. Finally (in Sections V and VI), I shall show how this helps to justify *forcing* people to pay for public goods, even when they are compulsory.

My first task, then, is to explain what makes free riding unfair. What makes this unfair, I shall argue, is what makes any unfair action unfair. Free riding has the feature that is common to all unfair actions. I begin by addressing a surprisingly underdiscussed question: what is that feature?

This, of course, is ambitious. It is beyond the scope of a single paper to defend a general account of fairness—even the kind of formal account I shall be presenting—in full detail. However, the following discussion might at least help us to start filling a gap in recent moral philosophy. I hope to show two things: thinking about the nature of fairness can help to clarify our obligations concerning public goods, and thinking about public goods can help to illuminate the nature of fairness.

I. Fairness

Actions are not the only things we evaluate as fair or unfair: we do this for states of affairs, attitudes, persons, and institutions too. But let us focus here on asking what makes actions fair or unfair. In the current section, I present

Rawls presents this as a version of Hart’s [1955: 185] principle of ‘mutuality of restrictions’. For earlier expressions of similar ideas, see [Broad 1916: 384–90] and [Ewing 1953]. Many revised versions have been subsequently defended by others (including myself, in [Cullity 1995: 18–19]).

Rawls originally endorsed the project of justifying political obligation on these grounds [1964]. However, he abandoned this project in [Rawls 1971: 113–16], on the grounds that the principle of fairness generates obligations only in respect of voluntarily accepted benefits. This has become the standard objection to grounding political obligation in the principle of fairness: for some other influential statements of this objection, see Simmons [1979a; 1979b: chap.V; 1993: 256–60] and Dworkin [1986: 192–3]. For further criticism, see also Pateman [1979: 121–9].

²Recent supporters of the project include Klosko [1992; 2001]; Arneson [1982]; Dagger [1997: chap.5]; Davis [1987]; and (with qualifications) Wolff [1995].

These writers have largely been concerned to show that the principle of fairness can be formulated in a way that avoids intuitively compelling counterexamples without being restricted to voluntarily accepted benefits, and that the principle so formulated forms a coherent fit with intuitively appealing judgements about political obligation. The aim here is to complement those discussions by offering a deeper explanation of how refusing to pay for unsolicited benefits can be unfair.

³In Cullity [1995], I argued that *if* free riding is unfair, then it can be unfair not to be willing to pay for compulsory goods. The aim here is to show why free riding *is* unfair, and to use this to strengthen the argument for thinking that refusing to pay for public goods can be wrong, even when they are compulsory.

an answer to that question, show how it makes sense of a range of judgements about fairness and unfairness, and then make some refinements to the account. In the next section, I discuss a number of apparent objections.

The core of a good general account of the fairness and unfairness of action, I claim, is this. Unfair actions are *failures of appropriate impartiality*; fair actions are those that are not unfair.

Ways of treating the members of a given group impartially are ways of preserving neutrality between them—ways of abstracting from their individual points of view.⁴ There are many different ways of doing this. You might impartially respond to the merits, or the interests, or the needs, or the claims, or the abilities, or the feelings of the people in a given group—and you might do so in relation to one group but not others. Each of these would be a way of being impartial, as long as you were setting aside the *other* features of the individuals concerned. Indeed, in any situation, there will be indefinitely many ways of treating people impartially. Judgements about fairness and unfairness, I claim, concern actions for which one particular way of being impartial is morally required. Giving the prize in a running race to whoever tried hardest, or picking the name of one of the competitors out of a hat, would be impartial procedures. But usually, the *appropriate* kind of impartiality in awarding a competition prize is to give it to whoever came first within the rules; and that makes the alternatives unfair, along with procedures that are blatantly partial, as when the race official favours her own friends.

Impartiality is most familiar as a feature of attitudes. However, it can also be attributed to actions, independently of the attitudes of their agents. For example, I might distribute my goods impartially amongst my children, even if the eldest is the only one I care about: indeed, I might do so *because* it is best for him.⁵ Here, my action is impartial although my attitude is not. It is impartiality as an attitude-independent feature of *actions* that is relevant to the present discussion—what is being described is the fairness and unfairness of actions, not the attitudes from which they are performed.

The strength of this proposal comes from its success in making sense of the wide variety of contexts in which we evaluate actions as fair or unfair.⁶ Starting with judgements of distributive fairness and unfairness, it is noticeable that the contexts in which and the criteria by which we make such judgements vary significantly—thanks to the differing aims that properly govern the distribution of different goods. Thus, fairness requires awarding competition prizes to the highest achievers, distributing compensation payments in proportion to the losses people have suffered, giving welfare handouts to the neediest, and so on. However, judgements of fairness and unfairness are not confined to distributive contexts, in which a good is available for distribution to people with different claims to it.⁷ Requirements of procedural fairness need not fit that model: consider the following of

⁴I am not offering this explanation as an *analysis* of impartiality—as if the relevant concepts of neutrality or abstraction were more basic than that of impartiality itself.

⁵Impartiality, as I have explained it, is relative to a group. In this example, I treat my children impartially. But I am being partial towards my children in preference to other people's.

⁶I discuss this range of cases in slightly fuller detail in Cullity [2004: 115–17].

⁷Simmons [1992] suggests otherwise.

established rules of procedure in a criminal trial, and respecting conventions of queuing for service.⁸ We also make complaints of unfairness against those who exploit others' trust: a swindler is naturally described as treating his victims unfairly. Moreover, if swindling you leads me to prosper while you languish, we might judge that both my not having restored your loss and my not having been punished are unfair. Another kind of action, different again, which we evaluate as fair and unfair is one person's criticism or praise of another.

No doubt, this list of kinds of fairness and unfairness—distributive, procedural, exploitation-of-trust, restorative, retributive, judgemental—could be extended further. However, let us stop there, and ask how an account of unfairness as the failure of appropriate impartiality succeeds in covering this range of cases. We noticed that different distributive practices are properly governed by different aims. Those aims determine the kind of impartiality appropriate to the distribution. Musical bursaries are fairly distributed to those with the most talent; medical resources to those whose medical needs are greatest: in either case, fairness requires impartially setting aside any *other* characteristics of the recipients as irrelevant. In general, distributive unfairness occurs when a distributive practice fails to be governed by the kind of impartiality appropriate to the point of that practice.⁹ Next, procedural unfairness occurs when a rule has been established (formally or informally) for regulating recurring situations in which the interests of different parties conflict, and either that rule has itself not been framed or established with the kind of impartiality appropriate to regulating such conflicts, or it has but is being violated. Unfairness-as-exploitation, as practised by the swindler, involves a different failure of impartiality. Relationships of trust involve an expressed commitment to reciprocal self-constraint for the sake of cooperation. In abusing that trust, the swindler arrogates a privilege he relies on others to forgo. He exploits others' willingness to observe the forms of impartiality required for cooperative relationships. Restorative and retributive unfairness are also failures of appropriate impartiality: they are failures to ensure that losses and punishments fall where they ought, abstracting from the identities of the victim and perpetrator of an offence. And judgemental unfairness occurs when expressed judgements about people fail to reflect their true merits.

Thus, in each of the contexts mentioned above fairness does require an appropriate form of impartiality. It is tempting to ask: What determines

⁸The line I am drawing between 'distributive' and 'procedural' fairness does not correspond to everyone's use of these labels. For example, Rawls [1971: 86], would count the competition prize example as an instance of pure procedural justice.

⁹Saying this is consistent with John Broome's view [1991: 95] that distributive fairness satisfies each person's claim to a good in proportion to the strength of the claim. However, an objection to that view comes from cases in which it is fair to conduct a lottery to distribute an indivisible good to which several people have equal claims. The resulting distribution satisfies one claim completely; the others not at all. It therefore seems better to say that distributive fairness requires responding to claims to a good in a way that is appropriately proportional to the strength of each person's claim to that good. (On the distribution of indivisible goods, see Szaniawski [1991].) This suggestion also differs from Broome's in its implications concerning divisible goods. Suppose I have a cake to distribute, everyone with a claim to it has an equally strong claim, and I simply throw the cake away. On Broome's view, this is not unfair [1991: 97–8]; on the alternative just suggested, it is unfair, since I am responding to those who have a claim in the way that it is appropriate to respond to those who have no claim at all.

which kind of impartiality (if any) is appropriate to a given context? However, I am not going to offer here an answer to that deeper, independent question. As far as I can see, the only way to do that would be to present a much more general normative moral theory. Many different theories would be consistent with what I say here.

Failures of appropriate impartiality provide the core of an adequate account of fairness and unfairness. However, we need to add to this core. After all, not doing something that would have been appropriate does not ensure the kind of moral *failing* involved in acting unfairly. Unfairness requires not just that the impartiality you fail to display would have been appropriate, but that it is the appropriate way of doing what ought to be done, as it ought to be done. Taking this into account, I propose the following general description of what is common to unfair actions.

Not Φ -ing is unfair when:

- (i) something ought, all things considered, to be done;
- (ii) doing it as it ought to be done requires a form of impartiality;
- (iii) Φ -ing is the appropriate form for that impartiality to take; and
- (iv) the failure of appropriate impartiality can contribute to a non-instrumental explanation of the failure to do what ought to be done.

II. Impartiality and What Ought to be Done

My claim so far is that the variety we find in judgements of fairness and unfairness can be explained as the variety of forms of impartiality appropriate to different situations. Several features of the four-clause account just stated call for further explanation and defence.

One objection my account might seem to invite is this. In support of that account, I have claimed that it succeeds in identifying what is common to a broad range of judgements of fairness and unfairness. But that is surely dubious as a general method for answering normative questions. If I want to know what is right, and you tell me what is common to a range of actions that are widely thought of as right, that does not settle my question. For it leaves open whether the actions widely thought of as right really are right.

However, what has been offered above is a formal, not a substantive account of unfairness. That account 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. The reply to this first objection, then, is that the formal account of unfair action offered above does not claim to

settle any normative questions. Rather, it tells us what kind of normative question is raised by an issue of fairness: a question about the kind of impartiality it makes sense to require.

Conditions (i) and (ii) are intended to span all the ways in which something ought to be done, and all the kinds of impartiality that can be appropriate to doing it. What ought to be done may be an individual action or a collective one (an action whose agent is a group); it might or might not be something that ought *morally* to be done; and the form of impartiality appropriate to doing it might involve either individual or collective action. Sometimes, we ought for prudential reasons to do something together, and impartiality is required from individual contributors if we are collectively to do it—such cases will occupy us shortly. But the account also covers straightforward cases in which a form of impartiality itself constitutes what I ought morally to do—not swindle other people, for example.

This might seem to invite counterexamples of five main kinds. The first two question whether conditions (i)–(iv) are necessary for unfairness, the next two question their sufficiency, and the last is a circularity challenge.

One kind of counterexample would be a case in which what ought, all things considered, to be done is unfair. In an appendix, I show that cases of this kind can be allowed for by adding complications to condition (i), without undermining the rest of my argument. However, the discussion will be easier to follow if we use the simpler account set out above.

For a second source of counterexamples, we can ask: Cannot requirements of fairness arise in the course of doing what ought *not* to be done? The running-race or musical bursary scheme might be misguided, but that would not stop biased refereeing or bursary-distribution from being unfair. Double-crossing my fellow gangsters could be unfair, even though the collective action of the gang is immoral [Simmons 1979b: 110–11]. So it might seem unfairness cannot, after all, be explained by reference to what ought to be done.

However, these examples do involve things we ought to do: it is just that those things have to be described negatively. I ought: *not* to be deviating from the impartial procedure for judging the race; *not* to be distributing the bursaries in a biased way; and *not* to be exploiting the trust of my fellow-gangsters.¹⁰ In each case, there is a failure to do what I ought, and in each case it is a failure of appropriate impartiality.

This might seem to short-circuit my account. It might seem to make it irrelevant to what fairness requires of me whether I ought to be distributing the bursaries, or whether the gang ought to be doing what it does. But that is not true. For in each case, two different requirements must be distinguished. The requirement not to double-cross the gangsters is not a requirement to cooperate with them. And the requirement not to distribute the bursaries in a biased way is different from a requirement to distribute them at all. If the gangsters' protection racket is immoral, I ought not to be cooperating with

¹⁰But surely there's nothing wrong with exploiting the trust of a gangster'. I agree that there can be morally defensible exploitations of trust. However, it is no help to the current objection to maintain that this is always true about gangsters. For that either abandons the idea that this is unfair and hence a counterexample to my account, or it turns into the objection discussed in the appendix.

them. So my account of unfairness implies, correctly, that refusing to cooperate with them is not unfair. But it also implies, correctly, that double-crossing them could be unfair. Likewise, my account explains why, if I really ought to be distributing the bursaries, not distributing them could be unfair.

Thus, this apparent objection to my account actually reveals one of its important strengths. It not only explains how requirements of fairness can arise when what ought to be done is negative; it also explains the difference between these and the requirements of fairness that arise when what ought to be done is positive.

A third challenge is this: If ‘appropriate impartiality’ can stretch to cover the various kinds of fairness surveyed in Section I, what stops it from including the whole of morality? If it did, my account would fail, since surely actions can be morally wrong for reasons other than unfairness—they can be cruel, dishonest, disloyal, or wrong in many other ways. However, ‘impartiality’ does not stretch that far. If I treat someone cruelly, then no doubt I give insufficient attention to his interests, and too much to whatever interests of my own the cruel action serves. But this kind of moral failure is not a failure of *impartiality*—a failure to be neutral between others, or between my victim and myself. Rather, it is a failure to give his interests even minimally decent consideration. Similarly for dishonesty: lying to you is not a failure of impartiality. It is not a failure to be neutral between us, but a failure to treat you with minimal respect.¹¹

(Some radically revisionary normative moral theories—such as Smart’s utilitarianism—do require complete impartiality from all agents always. Such theories have no need for a concept of fairness, as distinct from their general concept of rightness. On the view I am presenting, the concept of fairness marks the distinction that the rest of us make between the contexts in which we *are* morally required to adopt a form of impartiality and those in which we are not.)¹²

A fourth potential counterexample explains the need for condition (iv). Suppose that all things considered I ought (for no moral reason) to go to a certain show. Doing this as I ought requires a form of impartiality: queuing for a ticket. But suppose I decide not to go to the show at all. My action may be foolish, but it is not *unfair*—even though it does involve not acting in the impartial way that is required in order to do what ought to be done. Condition (iv) handles this: what explains my failure to do what I ought is my laziness or passivity, perhaps, but not my failure of impartiality.

In unusual circumstances, that might be false. Suppose my sense of superiority makes me hate queuing, and *this* is why I decide not to go to the show. Then my failure of impartiality does explain why I have not done what I ought, but staying home is still not unfair. However, this is only an

¹¹I am not saying cruelty, dishonesty, or disloyalty can never be unfair. On the contrary: leaving my partner to bring up our young children on her own might be all of those things. My point is that not all instances of these other moral failings are also unfair, and my account explains why not.

¹²In saying this, I am not taking sides on the deeper question whether the content of morality as a whole is determined by a criterion of impartial acceptability. (For discussion of this further question, see, e.g., the symposium in *Ethics* 101 [1991], Nagel [1991: chap.2], and Wolf [1992].) Maybe what makes cruelty morally wrong is that it involves a standard of behaviour that could not be accepted from a completely impartial point of view. But that is not to say that we are criticizing the cruel person for failing to act impartially.

instrumental explanation—my sense of superiority instrumentally causes me to stay home—so condition (iv) excludes this too. In all the cases of unfairness surveyed earlier, it will be found that a lack of appropriate impartiality *non-instrumentally* explains the failure to do what ought to be done. When a judge presides over a legal case unfairly, for example, her unfairness is not a further effect instrumentally caused by her lack of impartiality.

Last, there is the following worry about circularity. Biased umpiring is something that (normally) ought not to be done. Why? Isn't it natural to answer that it is unfair? But that can seem to make my account circular: it answers 'Why is biased umpiring unfair?' by saying that it satisfies conditions (i)–(iv), so it cannot then answer 'Why does biased umpiring satisfy condition (i)?' by saying that it is unfair.

It does make sense to say that biased umpiring ought not to be done because it is unfair. However, that does not saddle me with any circularity. For remember: mine is a formal account of unfairness. It does not identify the substantive reasons that make actions unfair; rather, it tells us what category the reasons against an action must belong to in order for 'unfairness' to be the appropriate term to use. There would be circularity in claiming that the satisfaction of condition (i) is a substantive reason for unfairness, and unfairness is a substantive reason for the satisfaction of condition (i); but I am not claiming that.

An analogy illustrates the point. Consider the following claim: 'Disloyal actions are actions that are (a) morally wrong because (b) they fail to take proper account of reasons of personal connection'. This is a formal account, since it does not say what the reasons referred to in (b) are, nor what taking proper account of them involves. Still less is it saying that the wrongness of an action is a substantive reason for its disloyalty. So there is no circularity in making this claim while accepting that actions can be wrong because they are disloyal.

There is a further kind of objection that cannot be so neatly refuted. Some people's intuitive judgements about the varieties of fairness will not agree with mine, so they will not be impressed by my claims about the range of judgements my own account explains. This is a delicate matter: there is bound to be disagreement over which of the things people tend to say about fairness are correct, and which stretch the term 'fair' beyond its natural application or are simply mistaken. The only way to resolve such disagreements is by comparing the coherence of rival accounts of what the correct judgements about fairness and unfairness have in common. I accept that the account I have just given should be rejected if a better one can be produced that draws together and explains a different range of recognizable judgements of fairness and unfairness. But I do not think any such rival account has yet been produced.¹³

¹³The two usual suggestions seem to have obvious flaws. One suggestion is that what fair actions have in common is that they are ways of treating people as they deserve—giving them their due [Hinton 1990: 285]. This seems too broad: if you are uniformly aggressive to other people you are not giving them their due, but the moral objection against you is not that you are being unfair (for a memorable illustration of this point, see http://en.wikipedia.org/wiki/Sidney_Morgenbesser). It also seems too narrow: if the undeserving athlete

I turn next to showing how this account makes good sense of a further important group of judgements of unfairness: judgements that free riding on the production of public goods is unfair.

III. Public Goods and Impartiality

Public goods are goods we possess jointly, in at least one of two general respects.¹⁴ The first is ‘jointness in supply’—a family of ways in which, roughly speaking, supplying the good to anyone means supplying it to everyone.¹⁵ The second is ‘jointness in consumption’, which covers a range of ways in which one person’s consumption of a good does not impinge on others.¹⁶ These two kinds of jointness are independent. (Fish in the sea exhibit jointness in supply but not consumption; a concert in a private theatre displays jointness in consumption but not supply.) However, many goods display both kinds of jointness, and these provide most people’s paradigms of public goods: goods such as firework displays, street lighting, and law enforcement. Treating publicity as a matter of degree, we can say that a good is public to the extent that it possesses features of these two kinds.

With large-scale public goods, the willingness of any one person to contribute is usually unlikely to make a difference to whether they are produced. They will either be produced or not irrespective of whether I contribute, and if produced can be enjoyed by me irrespective of whether I contribute. The same is true of everyone else.¹⁷ In these circumstances, producing a good requires overcoming two kinds of incentives: incentives to receive it for free, and disincentives to incur a cost which makes no difference to whether the good is produced.

There are many *possible* ways of producing large-scale public goods despite these incentives. We could induce altruists to produce them for us, as

finishes first within the rules, it would be unfair to deprive her of the prize [Rawls 2001: 73–4]. At best, fair action would have to be described as one way of treating people as they deserve; and the task of giving an account of fairness becomes the task of saying what is distinctive of that way.

The other suggestion is that we might explain the concern for fairness as a concern for *equality*. This takes various forms: thus fairness is taken to require that ‘similar individuals should be treated similarly’ Klosko [1992: 34], ‘that no one should be advantaged or disadvantaged by arbitrary factors’ Wolff [1998: 106], or ‘equality of status’ Hinton [2001]. A similar view is suggested by Griffin [1985] and Dworkin [1981]. However, such suggestions struggle to account for several of the non-distributive kinds of fairness and unfairness I have identified—unless they are qualified in a way that makes them equivalent to the account I have offered. The same, I would argue, is true of a third kind of view, advocated by Carr [2000] and Rescher [2002], according to which fairness consists in fidelity to social practices that meet certain conditions.

¹⁴The literature on public goods presents several different and incompatible definitions; but they all cluster around these two broad characteristics. For a careful survey of those definitions, see Cullity [1995; 2001].

¹⁵This is rough because it spans two distinguishable features: the ‘compulsoriness’ of goods one cannot avoid receiving without excessive cost and the ‘nonexcludability’ of goods one cannot be prevented (without excessive cost) from taking for free, once produced.

¹⁶Discussions of ‘non-rival’ goods often run together two issues: whether one person’s consumption of a good diminishes the amount available for consumption by anyone else, and whether it diminishes the benefits available to anyone else from consuming it. For economists, the interesting feature of joint consumption goods is that total consumption of them is not equal to the sum of individual consumptions. See, e.g., Samuelson [1955: 350].

¹⁷One version of this problem will be a many-person ‘Prisoner’s Dilemma’: a situation in which non-contribution is in the interests of each *whatever* anyone else does. But there are other versions as well. For a helpful discussion of the variety of strategic situations surrounding the production of public goods, see Hampton [1987]; also Frohlich et al. [1975].

a favour. Or we could enslave people and make them do so. However, I take it that this is rarely how we *ought* to produce such goods. Usually, we ought to do so by cooperating to act collectively in producing the good ourselves. When the group is large, doing what we ought, as we ought, will require that individuals contribute to the collective action although no individual contribution makes a difference to whether the collective action is performed, and although the good we produce may be available to non-contributors.

Under these circumstances, contributing to the collective action of the group involves a kind of impartiality. It need not involve *motives* of impartiality. Such motives are indeed common: individuals commonly do take the fact that we ought collectively to produce a good as a reason for contributing to the collective action, and it is natural to say that someone who reasons in this way is impartially motivated. But individuals can contribute to groups from all sorts of motives.¹⁸ I might be contributing in order not to lose face with my neighbours, or to annoy them by appearing morally superior. However, whatever my motives, my *action* can still have an impartial form. In contributing, I am doing what must be done by individuals if the group is to do what it ought. The group can only do what it ought, as it ought, through enough individuals' not exempting themselves from contributing. In this way, my action of contributing to the group involves the appropriate kind of neutrality between my situation and that of every other member of the group. In doing what the group needs its members to do, I am doing what is required of me simply *as a member of the group*.¹⁹

Given this, we can argue as follows. Some public goods are very important: we ought, all things considered, to produce them. But producing them as they ought to be produced requires a certain sort of appropriate impartiality. It requires that individuals do not exempt themselves from contributing. So the free rider fails to observe a kind of appropriate impartiality that is required for the doing of what ought to be done, as it ought to be done. Moreover, the individual contributions *constitute* the performance of the collective action: they are not instrumental means to something else we ought to do. So individual failures of appropriate impartiality can contribute to a non-instrumental explanation of our collective failure to do what we ought.²⁰ If so, we can explain how free riding on the production of a public good is unfair. The conditions set out in Section I are all met. Free riding possesses the features that are common to all cases of unfairness.²¹

¹⁸On the variety of motives underpinning contributions to collective action, see Elster [1985]. For an empirical study of the conditions under which the uncoerced production of public goods actually happens, see Ostrom [1990]. For studies of the social psychology governing cooperating individuals, see Goldberg et al. [2005].

¹⁹Of course, my contribution also involves a form of partiality: partiality to this group rather than others. But that is no objection to my claim: see note 5.

²⁰My individual failure of impartiality might not *actually* contribute to explaining our collectively failing to do what we ought: if enough others join in, we might still succeed. But it remains true that such individual failures *can* contribute to explaining this.

²¹Most versions of the principle of fairness require that the costs and benefits of a cooperative scheme must be fairly distributed. Where this condition is justified, it will be a corollary of the account presented here, because it will bear on whether we ought to be operating the scheme as we are. However, I think it is wrong to claim that non-contribution is only unfair when costs and benefits are distributed with perfect fairness: I explain why in Section V.

Moreover, this can be true when we ought, all things considered, to be producing a good for the benefit of *others*, rather than for ourselves. Refusing to join in such an action is not naturally described as free riding. But it can still be unfair, as a failure of the kind of impartiality appropriate to our doing what ought to be done. Suppose a group of us are bystanders at an accident, when several people need to be helped. The bystanders ought collectively to help, and I ought to join in. When I refuse to do so, I might know that the others will work harder without me and everyone who needs help will get it. I might then deny that I wrong the accident victims. But I am still treating the rescuers unfairly.

I add two comments about the intended force of this argument.

First, a disclaimer. I am not claiming that whenever a group to which I belong ought collectively to be producing a good, I ought to be doing something towards producing it. That would have odd implications for collective actions to which no one is contributing, and collective actions to which others have contributed completely. If my neighbourhood really ought to be holding a street party but is not, fairness does not require me to hold up the traffic and drink in the street on my own. And if we ought to fill up a water tank, fairness does not require me after it has already been filled to pour in an amount that makes it overflow. However, my account accommodates such judgements. Where we ought collectively to be pursuing a goal, the *appropriate* kind of impartiality required of individuals is cooperating in a collectively agreed method for pursuing it. (More on this shortly.) So there is nothing unfair about my not unilaterally pursuing a collective goal when there is no such agreement. And there is nothing unfair about my not making an overflow ‘contribution’ towards a goal that has already been achieved, since that is not a form of *appropriate* impartiality either.

The other comment is this. I am not claiming that the conclusion that free riding is unfair can be deduced from premises that an opponent is committed to accepting. Those who deny that free riding is unfair will probably want to reject the Section I account. Indeed, they might add that no account of unfairness derived from a survey of judgements about unfairness could compel my conclusion without begging the question against a rival account—one that endorses all those judgements except for the unfairness of free riding. However, I have not been attempting a proof of the unfairness of free riding against such an opponent. Instead, the aim has been to show that we can make good sense of the widespread judgement that free riding is unfair, identifying a feature that it shares with other paradigm cases of unfairness—a feature it makes sense to see as morally important. Relations of justificational support run in two directions. The identification of this common feature tends to increase the credibility of the judgement that free riding is unfair. But also, our widespread disposition to see free riding as unfair supports the claim that the Section I account, which identifies a pattern common to all the cases of unfairness I have mentioned, gives a good general account of it. This does not show that a rival account is impossible; but avoiding my conclusion requires producing that account and arguing for its superiority.

Moreover, notice that my discussion relies on a particular reading of the Section I account of unfairness. Clause (ii) is satisfied when doing what ought to be done requires a form of impartiality. But while collectively producing large-scale public goods requires individuals to exercise a form of impartiality, it does not require *me* to do so. So a particular reading has been assumed: when doing what ought to be done requires a form of impartiality from people relevantly like me, and I fail to display that form of impartiality, this can be what makes my conduct unfair. The conviction that free riding is unfair, then, is itself guiding the interpretation of my account of unfairness.

The Section I account of unfairness, I claim, makes good sense of the relatively uncontroversial judgement that paradigm cases of free riding are unfair; and is in turn supported by those judgements. Now I turn to showing that this account also vindicates the more controversial judgements I mentioned at the outset.

IV. Compulsory Goods

Paradigm cases of free riding involve actively taking a good without having contributed to producing it: you refuse to help dig a new well, and then take water from it afterwards. But what if a good is conferred on you non-voluntarily? Suppose the village water supply is a river, and the water is unclean. Everyone else agrees to cooperate in cleaning it up, and you refuse. Can the others complain that you are acting unfairly if you continue drawing water from the river as you always have done? Or is your behaviour morally defensible, given that the benefit has been pressed upon you?

Clearly, not *all* unsolicited benefits are the subject of legitimate demands for payment. Nozick [1974: 90–5] is surely right that I am not morally required to pay for books that are thrown into my house with bills attached, nor to take a day off work to entertain the neighbourhood over a public address system after my neighbours have taken it in turns to do so. However, such examples do not show that refusing to pay for unsolicited benefits is always morally acceptable.²² Sometimes, it is wrong to refuse to contribute towards the cost of compulsory goods. We are now equipped to explain why.

Section III explained the conditions under which refusing to pay for a public good can be unfair. Those conditions can be met when the good is unsolicited. Some public goods are important enough to mean that we ought, all things considered, to produce them. This can be true of compulsory goods. But it is often also true of such goods that when we produce them as we ought, it is through individuals' contributing to the group, each meeting a share of the cost. This involves a kind of

²²Nozick's discussion does not commit him to this conclusion. However, he does argue that it is always wrong to *force* a person to pay for unsolicited goods.

impartiality: the impartiality of doing what the group needs its members to do if we are collectively to do what we ought. The failure to display this kind of impartiality satisfies the conditions that are common to cases of unfairness.

Opponents of this conclusion might object as follows. I cannot be accused of treating others *unfairly* if I refuse to pay for something I am forced to receive. A free rider is someone who exploits, or takes advantage of, other people's cooperative spirit. A recipient of unsolicited goods exploits no one: he is just minding his own business.

However, this relies on the assumption that unfairness requires exploitation. Section I showed us that that is unwarranted. Some forms of unfairness—the swindler's, for example—do involve exploitation. Free riding of the active sort involves a kind of exploitation. But this is not a general characteristic of unfairness. What unfair actions have in common is that they are failures to display the form of impartiality appropriate to doing something that ought to be done, as it ought to be done. Failures to contribute towards producing compulsory goods can have *this* feature. So such failures can be unfair.

We need to be careful about exactly which unsolicited benefits this argument applies to. It only applies to goods that really ought, all things considered, to be produced, and to our supplying them as they ought, all things considered, to be supplied. This has implications for unsolicited goods that *could* have been produced and supplied without making them compulsory. If, despite this, they are supplied to people compulsorily, there will usually be grounds for doubting that they are being supplied as they ought: it is important to give people opportunities for self-determination when we can.²³ And if a good is not being produced and supplied as it ought, that undermines the case for thinking that refusing to pay for it is unfair.²⁴ This applies to Nozick's examples. We should therefore agree with him about those examples, without generalizing the point to all unsolicited goods.

This gives us an argument that refusing to pay for compulsory goods can be unfair: it can be unfair in the same way that paradigm examples of free riding are unfair. However, to make this fully convincing, we must go further. For a deeper set of worries remains to be addressed.²⁵ The villager who won't help clean the river might respond: 'If others decide to perform some public-spirited actions, that may be laudable; but how does their decision bind *me*? I cannot be co-opted into collective projects, no matter how worthy, through *others'* deciding to pursue them'. This is the line of thought we need to address next.

²³Compare Arneson [1982: 622]. I give some reasons for thinking that this is only *usually* the case [1995: Section III].

²⁴As Section II showed, my actively taking a good without paying might be unfair even if the good ought not to have been produced. This is because it might be an exploitation of trust, procedurally unfair, or unfair in some other way. But refusing to pay for a *compulsory* good that ought not to have been produced is not unfair in any of those ways.

²⁵These worries seem central to Nozick's discussion [1974: 94–5].

V. Who Decides?

Two important questions remain unanswered by anything I have said so far. When is it right to force people to pay for public goods; and how are decisions about which goods to produce to be reached? So far, I have argued that when we ought, all things considered, to produce a public good, and producing it as it ought to be produced requires individual contributions, it is unfair to receive the benefit without contributing. Notice that this is not a claim about forcing people to do anything, nor about how we ought to decide what to do. However, these further questions are relevant to many real-world public goods. And they seem to be linked. If others decide which goods to produce without consulting me, and then force me to pay, surely I can complain about unfairness.

It is now time to turn to those two further questions. While the argument so far has not itself answered them, it provides materials that will help us.

Let us start with the question about collective decision-making. Groups of any size can contain differing opinions about what we ought to do; the larger the group, the more likely such disagreement is. Whose opinion should prevail?

Often, the solution is to establish a fair and efficient procedure for decision-making—one which, where practicable, gives a fair hearing to everyone's opinions (or, in large groups, the opinions of representatives), and then impartially arbitrates between those opinions. Often, but not always. Some collective actions are not important enough: if we disagree about how to carry the table into the next room, we need not sit down and devise a way of impartially arbitrating our disagreement. Others are too urgent: in a life-saving emergency, people should often just follow the instructions of the most competent-seeming person present, without wasting time on discussion.

Both kinds of case—where we should adopt a formal procedure for inclusive and impartial group decision-making, and where we shouldn't—fall under the scope of the account offered above. For notice that the question, 'How ought we to decide?' is also a question about a collective action: the action of making a collective decision. About actions of this kind, we can ask: What, all things considered, ought we to do? Sometimes, reasons of inclusiveness and efficiency decisively favour establishing a formal procedure for collective decision-making, so this is what, all things considered, we ought to do.²⁶

According to the Section I account of fairness, when doing what we ought to do, as we ought, requires certain appropriate forms of impartiality, those forms of impartiality are requirements of fairness. Which forms of impartiality are required in relation to formal procedures for collective decision-making? They fall into two groups. First, participants should respect the outcome, even when it is personally uncongenial. And secondly, a range of requirements governs the conduct of the discussion. Participants

²⁶I mention here some of the reasons that bear on what we ought to do, but have given nothing like a full account of them. My aim is not to offer full justifications for claims about what we ought to do; but to argue that when there *is* something we ought to do, requirements of fairness follow.

must respect properly impartial rules of procedure, and should advance only impartially acceptable reasons. It is an important question just what this rules out—we return to this in the next section. I shall not attempt a general account; but some examples are obvious enough. Any appropriately inclusive and impartial procedure will rule out using threats against those who disagree with you. Notice also that the requirement that we respect the outcome has implications for which reasons are impartially admissible. The fact that a certain policy would benefit free riders—that is, those who will *not* cooperate in implementing the outcome of our discussion—could hardly be admitted as an impartially acceptable reason in favour of that policy. Requirements of these different kinds are forms of impartiality required for doing what we ought, as we ought—namely, following an inclusive and impartial procedure for collective decision—so they are themselves requirements of fairness.

Thus, fairness will usually require me to respect both the procedure and its outcome, even if I disagree with the decision we reach [Klosko 1992: 72–7]. My disagreement will not affect the fact that we ought to reach collective decisions by using this sort of procedure, and that our doing so requires people like me to make an appropriately impartial contribution—a contribution which includes respecting the outcome. So when we ought collectively to be following an established decision procedure, fairness will usually require me to contribute towards what we have decided to do, even when I think we ought not to have decided to do it.²⁷ My failure to respect our collective decision is not free riding. But it has the more general characteristic (identified in Section I) that all forms of unfairness share in common.

(This might seem to generate contradictory requirements of fairness. Suppose we decide not to increase the assessment required of our students, when we ought to have decided otherwise. Doesn't that mean there are two things we ought to do: increase the assessment, and implement our decision not to? And doesn't my account of fairness therefore imply that fairness requires me both to contribute to increasing the assessment and to contribute to implementing our decision not to? No. For even if it is true that we ought to be doing what we ought to have decided to do,²⁸ doing it in opposition to a fairly reached contrary decision will not be doing it as it ought to be done. So clause (i) of our account is not satisfied.)

I am not claiming that we ought only to respect the outcome of a collective decision procedure when it is perfectly fair [Rawls 1971: 112; Klosko 1992: 66]. Human frailty being what it is, that is rarely true. But *that* makes an important difference to which decision procedures we ought, all things considered, to respect. Respecting only those that are perfect would be a disastrous limitation on collective action. *Therefore*, we ought to respect those that are good enough.²⁹

²⁷Usually: some collective decisions reached by fair procedures might be so bad that, all things considered, I ought not to respect them.

²⁸Is this true? The alternative is to say that although we ought not to have decided to perform this action, having decided to perform it, we ought to perform it. The present paper remains agnostic on this issue: my [unpublished] is not. (For relevant discussion, see Rawls [1964: Sections 3 and 4].)

²⁹For some suitable specification of 'good enough'. See note 26.

This leads to a regress. We started with the question: Which collective goods ought we to produce? Opinions about this can differ, so we faced a second question: How ought we to decide? We have seen that, if there are all-things-considered answers to the second question, they generate requirements of fairness on individuals. However, does this not simply produce the same problem, at a higher level? Opinions about the second question can differ too: how should we resolve *this*? That is a sensible question, and a practically important one. It is no objection to my discussion, though. The regress is not vicious: it is one of the practical realities of collective decision-making, and confronts all constitution-makers. How do we decide on the procedure for creating a constitution? This may itself be a question which ought to be openly debated following the disciplines of a fairly designed formal procedure: that will be appropriate in the most important cases. If so, the question arises again concerning the details of that procedure: how should we decide on the format for our constitutional convention? Plainly, if we are to reach substantive collective decisions, not every decision about decision-making procedures can be the subject of prior collective discussion. At some point, we must simply commit ourselves to a way of deciding how to decide.³⁰ So if we really ought to be making substantive collective decisions, we ought to make such a commitment.

VI. Coercion

The remaining question is this: When is it right to *force* someone to pay for a collectively produced good?

This clearly goes beyond the question when refusing to pay is unfair. For not all unfair actions can justifiably be prevented by force. The unfairness may not be serious enough: using force against a queue-jumper may be a disproportionate response. More importantly, the use of force may not be *my* prerogative. Given the importance of restraining the misuse of force and avoiding retaliatory escalation, we have good reason to vest the authorization of force in an impartial authority.³¹

These points suggest some important constraints on compelling payment for a public good. First, it must be the case that the good really ought, all things considered, to be produced, and is being produced as it ought. This often requires that the decision to produce it has been reached through the kind of inclusive and impartial procedure described above. Secondly, the use of force should usually be vested in an impartial authority.³² And thirdly, the good must be important enough to justify coercion.³³ I shall not attempt

³⁰With small groups, we can often solve this by unanimity. With large ones, we often cannot.

³¹This point—a point often made since Hobbes—invites another application of my argument, in support of the conclusion that vigilantism is a kind of unfairness. The rule of law is a very important public good which we ought, all things considered, to produce; the way it ought to be produced is through individuals' relinquishing the personal use of force; and doing so involves a kind of impartiality, doing what is required from a point of view that abstracts from one's own situation.

³²Usually: I leave open the possibility of emergency cases.

³³Compare Arneson [1982: 621].

an account of how important is ‘important enough’. Many such accounts will be compatible with the rest of my argument. However, it is plausible to think that the value of personal autonomy makes this third requirement non-trivial.

Perhaps a fourth condition should be added: coercion must really be necessary to feasibly producing the good. If it could feasibly be produced without coercive incentives, perhaps we should simply put up with a few free riders. Again, the *pro tanto* undesirability of coercion would seem to support this; but I shall not argue the point.

Rather than attempting to justify and explain more fully this list of constraints, let us ask this: How much further could it extend? Suppose you stand to be *compulsorily* benefited by a public good. And suppose it would be unfair not to contribute towards its production: the reasons spelt out in Section IV are in force. Might the compulsory nature of the good make it wrong to force you to pay? Should we add a fifth constraint: a general constraint against forcing people to pay for compulsory goods? The answer is No, and our earlier discussion helps explain why.

Suppose that, by a fair, inclusive procedure, we are deciding whether to produce a certain compulsory good. What decisions might we properly reach? Suppose the good is very important—collective security, for example—and we sensibly decide to produce it. Given this, Section V implies that fairness requires individuals to cooperate in implementing the decision.

However, suppose we confront a problem. Although some of us contribute, many do not. The group is large; the incentive not to contribute prevails. So suppose we ask (reinstating our fair procedure for collective deliberation) what to do about this. We see that the only practical way to produce the good is by using coercive incentives. It is proposed that we vest in an impartial authority the power to enforce contribution.

At this stage in our deliberations, all four conditions mentioned above are satisfied.³⁴ The question is whether any good objection could now be raised, as part of this fair collective procedure, against our deciding to use coercion. Some ways of pursuing socially desirable goals can be vetoed because they fail to respect the rights and liberties of some individuals, imposing unreasonable costs on them. Could an objection of this kind be raised on behalf of those who stand to be coerced? No: for non-contributors are acting *unfairly*. Asserting an entitlement not to be forced to contribute would amount to asserting an entitlement of non-contributors to treat everyone else unfairly. However, *this* cannot be an admissible reason within a fair collective decision-making procedure. In some contexts, there may be rights to treat other people unfairly.³⁵ But no such right could sensibly be invoked here. Remember: in Section V, we saw that two requirements of fairness—two forms of appropriate impartiality—must be met by participants in properly impartial and inclusive procedures for collective

³⁴This is not the *only* way the four conditions might be met. In particular, it is often unreasonable to have to try producing a good without coercion before resorting to coercion.

³⁵I might have a right to favour one of my children unfairly over the others, at least in the sense that no one else is morally entitled to prevent me.

decision-making. One is respecting the outcome. The other is observing the proper standards of impartiality that govern the conduct of the discussion—which includes restricting ourselves to advancing and accepting only impartially acceptable reasons. Asserting an entitlement to treat everyone else *unfairly*—an entitlement claimed on behalf of those who do *not* respect the outcome—could not meet that second condition. So it cannot ground a good objection to our fairly deciding to use coercion against non-contributors to the production of the compulsory goods we ought to produce.

Moreover, if this shows that people can legitimately be forced to pay for goods they receive compulsorily, it will also show that they can legitimately be forced to pay for goods that are produced for others' benefit. In Section III, we saw that failures to contribute to producing such goods can be unfair (without amounting to free riding). When that is so, the assertion of an entitlement not to contribute cannot be admissible as a reason within a fair collective decision-making procedure.

I have not argued that coercion can *always* properly be used to enforce contributions towards goals we collectively ought to pursue. We have already surveyed some reasons for doubting that. These reasons constrain the argument that has just been given. It provides no defence of the use of coercion to compel individuals to contribute to a collective action when either our action ought to have been decided through a fair procedure but has not, or the goal of the collective action can be attained without coercion, or it is not important enough to justify coercion.³⁶ These will be reasons for denying that, in using coercion to pursue the goal, it is being pursued as it ought. However, when goals are very important, and can only feasibly be achieved by using compulsion, the fact that non-contribution is unfair undermines non-contributors' complaints about coercion within an appropriately impartial collective decision procedure.

VII. Conclusions

I have argued for two main conclusions. First, non-contribution towards the cost of a collectively produced good is wrong, because it is unfair, when that good ought, all things considered, to be produced in a certain way, and producing it in that way requires individual contributions. This can be true of compulsory as well as non-compulsory goods. Questions about which goods we ought to produce and how we ought to produce them can be contentious, but there are ways in which we ought collectively to deliberate about those questions, and they generate requirements of fairness too. Secondly, a person can properly be compelled to pay for a collectively produced good when the collective action of producing it has been decided through a fair procedure, the good cannot be produced without coercion, it

³⁶My argument cannot do this; perhaps another can. Perhaps non-democratic societies as well as democratic ones can justifiably tax their citizens to produce important public goods. But a defence of this would have to come from elsewhere.

is important enough to justify coercion, and the coercion itself is vested in an impartial authority.

What I have shown is limited. To settle just which public goods we can properly force people to pay for, resources are needed that this paper has not supplied. The argument set out above does not itself tell us when a collective goal is important enough to justify coercion and when it is not. Further normative ethical argument is required to establish that. However, I have argued that forcing people to pay for public goods can be justified, even when those goods are compulsory. Even when you have not endorsed our collective goals, it can be right to force you to contribute towards them.

My argument has not amounted to a deductive demonstration of the correctness of these conclusions. Rather, my aim has been threefold: to present a plausible account of what unifies our attributions of fairness and unfairness; to show its application to public goods; and to present this in a light which shows why we should regard it as important—that is, why fairness and unfairness matter.³⁷

University of Adelaide

Received: November 2005

Revised: October 2006

References

- Arneson, Richard J. 1982. The Principle of Fairness and Free-Rider Problems, *Ethics* 92: 616–33.
- Broad, C. D. 1916. On the Function of False Hypotheses in Ethics, *International Journal of Ethics* 26: 384–90.
- Broome, John 1991. Fairness, *Proceedings of the Aristotelian Society* 91: 87–101.
- Carr, Craig L. 2000. *On Fairness*, Aldershot: Ashgate.
- Cullity, Garrett 1995. Moral Free Riding, *Philosophy and Public Affairs* 24: 3–34.
- Cullity, Garrett 2001. Public Goods, in *Encyclopedia of Ethics*, ed. Lawrence C. Becker and Charlotte B. Becker, 2nd edn, New York: Routledge: 1413–16.
- Cullity, Garrett 2004. *The Moral Demands of Affluence*, Oxford: Clarendon Press.
- Cullity, Garrett unpublished. Decisions, Reasons, and Rationality.
- Dagger, Richard 1997. *Civic Virtues*, New York: Oxford University Press.
- Davis, Michael 1987. Nozick's Argument for the Legitimacy of the Welfare State, *Ethics* 97: 576–94.
- Dworkin, Ronald 1981. What is Equality? Parts 1 and 2, *Philosophy and Public Affairs* 10: 185–246; 283–345.
- Dworkin, Ronald 1986. *Law's Empire*, Cambridge MA: Harvard University Press.
- Elster, Jon 1985. Morality and Collective Action, *Ethics* 96: 60–72.
- Ewing, A. C. 1953. What Would Happen if Everyone Acted Like Me?, *Philosophy* 28: 16–29.
- Frohlich, Norman, Thomas Hunt, Joe A. Oppenheimer, and R. Harrison Wagner 1975. Individual Contributions for Collective Goods: Alternative Models, *Journal of Conflict Resolution* 19: 310–29.
- Goldberg, Jeffrey et al. 2005. Symmetry and the Illusion of Control as Bases for Cooperative Behavior, *Rationality and Society* 17: 243–70.
- Griffin, James 1985. Some Problems of Fairness, *Ethics* 96: 100–18.
- Hampton, Jean 1987. Free-rider Problems in the Production of Collective Goods, *Economics and Philosophy* 3: 245–73.
- Hart, H. L. A. 1955. Are There Any Natural Rights?, *Philosophical Review* 64: 175–91.
- Hinton, J. M. 1990. Knowing and Valuing Fairness, *Philosophy* 65: 271–97.
- Hinton, Timothy 2001. Must Egalitarians Choose Between Fairness and Respect?, *Philosophy and Public Affairs* 30: 72–87.
- Klosko, George 1992. *The Principle of Fairness and Political Obligation*, Lanham Md: Rowman & Littlefield.
- Klosko, George 2001. The Natural Basis of Political Obligation, *Social Philosophy and Policy* 18: 93–114.
- Nagel, Thomas 1991. *Equality and Partiality*, New York: Oxford University Press.
- Nozick, Robert 1974. *Anarchy, State and Utopia*, New York: Basic Books.
- Ostrom, Elinor 1990. *Governing the Commons: The Evolution of Institutions for Collective Action*, Cambridge: Cambridge University Press.

³⁷This paper has been greatly improved by helpful comments from Maurice Goldsmith, Daniel Shapiro, Paul Tudico, two anonymous readers for the *Australasian Journal of Philosophy*, and audiences at the Universities of Edinburgh, Glasgow, St Andrews, CAPPE Canberra, the inaugural LSU Ethics Symposium and the 2003 AAP conference.

- Pateman, Carole 1979. *The Problem of Political Obligation*, New York: University of California Press.
- Rawls, John 1971. *A Theory of Justice*, Cambridge MA: Harvard University Press.
- Rawls, John 1964. Legal Obligation and the Duty of Fair Play, in *Law and Philosophy*, ed. Sidney Hook, New York: New York University Press: 3–18.
- Rawls, John 2001. *Justice as Fairness: A Restatement*, ed. Erin Kelly, Cambridge MA: Harvard University Press.
- Rescher, Nicholas 2002. *Fairness: Theory and Practice of Distributive Justice*, New Brunswick NJ: Transaction Publishers.
- Samuelson, Paul A. 1955. Diagrammatic Exposition of a Theory of Public Expenditure, *Review of Economics and Statistics* 37: 350–6.
- Simmons, A. John 1979a. The Principle of Fair Play, *Philosophy and Public Affairs* 8: 307–37.
- Simmons, A. John 1979b. *Moral Principles and Political Obligation*, Princeton: Princeton University Press.
- Simmons, A. John 1992. Fairness, in *Encyclopedia of Ethics*, ed. Lawrence C. Becker and Charlotte B. Becker, 1st edn, Chicago: St James Press.
- Simmons, A. John 1993. *On the Edge of Anarchy*, Princeton: Princeton University Press.
- Szaniawski, Klemens 1991. On Fair Distribution of Indivisible Goods, in *Logic and Ethics*, ed. Peter Geach, Boston: Kluwer: 275–88.
- Wolf, Susan 1992. Morality and Partiality, *Philosophical Perspectives* 6: 243–59.
- Wolff, Jonathan 1995. Political Obligation, Fairness, and Independence, *Ratio* 8: 87–99.
- Wolff, Jonathan 1998. Fairness, Respect, and the Egalitarian Ethos, *Philosophy and Public Affairs* 27: 97–122.

Appendix

Some people hold that sometimes what ought, all things considered, to be done is unfair. As it stands, the account of unfairness at the end of Section I is inconsistent with that view. I do not wish to debate here whether that view is correct. So instead, I shall show that the account can be easily modified to accommodate it.

There are two different grounds on which you might hold that what ought, all things considered, to be done can be unfair. You might think that sometimes it is morally wrong not to do what is unfair. Or you might think that an unfair action may be, all things considered, what a person ought to do even though it is morally wrong. Each of these can be accommodated by a disjunctive modification of clause (i) of my account.

Someone making the first claim faces a challenge: ‘What do you mean when you say that an action is unfair? If you just mean it involves not doing something that would be fair, you make unfairness uninteresting—there need be no moral defect in an action that is ‘unfair’ in *that* sense. But you cannot mean that it deviates far enough from fairness to be morally wrong. What then do you mean?’ As far as I can see, the most promising reply is this: an unfair action is one whose deviation from fairness would be enough to make it morally wrong were it not for the presence of countervailing moral reasons that count in its favour. Someone adopting that view should then modify clause (i) of the earlier account as follows:

- (i)* something either ought, all things considered, to be done, or it would have been the case that it ought to be done, had countervailing moral reasons not been present.

There may be other replies to the challenge that do not take this counterfactual form. They could be substituted in place of the counterfactual modification.

The second claim is that, even when unfair actions are wrong, they can still, all things considered, be what one ought to do, since non-moral reasons can defeat moral ones. Here, there is a simpler disjunctive modification:

(i)** something either ought morally or ought all things considered to be done.

Either modification could be adopted without affecting the argument of the paper. The cases I am concerned with, in which fairness requires individuals to contribute to our collectively producing a public good as we ought, will still fall under the first disjunct of (i)* and one of the two disjuncts of (i)**.