

McGill University

From the Selected Works of Jurgen De Wispelaere

December, 2002

Altruism, Impartiality and Moral Demands

Jurgen De Wispelaere



Available at: <https://works.bepress.com/dewispelaere/8/>

Altruism, Impartiality and Moral Demands

JURGEN DE WISPELAERE

Moral Demands and Moral Motivation

Morality typically imposes considerable demands on individuals. That is, any practical account of how people should behave towards one another – whether to establish a *modus vivendi* to avoid a Hobbesian war-of-all-against-all, or to promote a more inclusive Rawlsian well-ordered society – requires that individuals adhere to a set of rules, guidelines or maxims that pose *some* constraints on their decisions, choices and actions. The specific content of these demands obviously varies with the particular features of the moral theory in question, but even theories that are explicitly concerned with promoting agents' self-interest readily accept the idea of moral demands as justified constraints on personal (or group) agency.¹

The idea of moral demands as a constraint on agency immediately raises the question whether a moral theory needs an account of why people should acknowledge these demands and adjust their behaviour accordingly. Following Christine Korsgaard (1996a) we can discern reasons for thinking that we *ought* to abide by moral demands – the justificatory component – and reasons that agents *actually* employ when engaging in action or decision-making – the motivational component. There is undoubtedly considerable dispute amongst contemporary moral theorists about the question whether a motivational component should be an inclusive part of a moral theory (Mendus 2002: 9–18). But I think Korsgaard (1996a: 14) has it basically right when she claims that 'a theory of moral concepts which left the practical and psychological effects of moral ideas inexplicable could not even hope to justify it'. A moral theory that imposes overly demanding requirements – that is, requirements that would render it effectively impossible to motivate

individuals to behave in accordance with them² – is inherently problematic as a theory of practical reasoning (O'Neill 2000; Nagel 1970, 1991).

Motivations to respect moral demands come in two basic forms. Some agents will be directly motivated to adhere to the terms of the moral theory: they accept the demands – rules, regulations and maxims – dictated by the theory and consequently comply with them out of a *sense of morality*.³ For instance, I may refrain from cheating on my tax returns because I firmly believe that, as a matter of justice, the state has a right to require me to contribute towards the provision of public goods and the assistance of fellow citizens who face less favourable life prospects. My willingly allowing the state to tax my income or wealth here signifies my acceptance of the basic rules of justice, and the way they apply to society as a whole and me in particular.⁴ This motivation can be contrasted with that of agents who are motivated, not by acceptance of the requirement of justice, but by a *concern for others*. Concern for others reflects a distinctive moral disposition from the other-regarding outlook entailed in most common moral theories.⁵ The distinction between both is complex, but I believe it has important implications.

When I express a concern for others, and a willingness to respond positively to certain moral demands because I care about others, I do not in any way signal acceptance of any particular moral theory. Of course this is not to say that a disposition to care for others cannot *indirectly* support the moral demands associated with a moral theory. Suppose I happily donate to a poverty relief charity because I am deeply touched by the fate of those living in dire poverty. I may not think about justice at all, or even agree that I have a duty grounded in justice to do so, yet I may inadvertently discharge precisely the sort of obligations that a particular theory of justice would require of me. Adherents of that particular theory of justice might then have no qualms about 'enlisting' my caring disposition to bring poverty relief about.

However, because of this rather loose connection between my disposition to care for others and the requirements associated with a moral theory, it is also entirely possible that I fail to respond appropriately to moral demands. It is this latter aspect that is the prime topic of this paper, and I shall proceed by outlining a number of cases that appear to fail in this respect. Note that an 'inappropriate' response might mean that I fail to address certain moral demands, despite being

disposed to care about others; it could also mean that I take on more demands than I should to discharge my obligations. The problem has some similarity to the occurrence of Type I and II statistical errors in social policy. In social policy, rules often imperfectly track certain social objectives – say, poverty relief – and policy makers regularly have to make a choice between minimising errors that provide relief to non-deserving or non-eligible recipients and failing to provide relief to deserving or eligible recipients (Goodin 1985). Perhaps one could argue, following the analogy of welfare policy, that a disposition to care for others only tracks moral demands imperfectly, thus leading to the mentioned two types of failure to respond accurately to moral demands.⁶

To illustrate these issues I shall look at how altruistic dispositions fare in respect to the compliance with moral demands. Altruism is a theory about individuals' inherently beneficial and, therefore, morally desirable motivational dispositions towards furthering other people's good, and as such a prime example of a disposition to act out of a concern for others.⁷ To say that altruism is inherently beneficial is not to say that the theory provides us with decisive moral reasons for behaving altruistically at all times; this would clearly be too strong a requirement. The sort of altruistic theory I believe would be much more plausible – call it *pro tanto* altruism – generates only conditional reasons, which, it is accepted, at times may be overridden by moral or even non-moral considerations.⁸ But accepting *pro tanto* altruism implies an explicit endorsement that performing an altruistic action is *always* good, while allowing an altruistic act to be overridden by other pressing reasons, while justified, nevertheless implies some sort of moral loss. It is precisely this dual claim I want to argue against in this essay by showing how it fails to address moral demands that most common-sense accounts of morality would consider plausible.

Outlining Altruism

A first task of this essay is to outline unambiguously what altruistic dispositions are. Many writers outline altruism in vague terms, thereby inviting confusion and failing to point out what (if anything) is really distinctive about altruism. I propose to restrict altruism to a motivational disposition, according to which an intention or motivation to furthering other persons' good is both the *necessary* and *sufficient* condition for an act to classify as altruistic. This raises a number of important issues, which will be discussed further.

Altruism and Purposive Benefits

Intuitively, the idea that the altruist purposively benefits herself is hard to combine with the notion of genuinely altruistic motivations. Acts that appear altruistic while in fact being motivated by an ‘enlightened’ form of self-interest only *appear* altruistic at face value: they are best referred to as *quasi-altruistic*. A well-known example of quasi-altruism is the traditional custom of gift-giving in the Trobriand culture, famously described by Bronislaw Malinowski and others (Malinowski 1932; Mauss 1954). Anthropologists have discovered that the custom of the gift is in fact a social mechanism, rooted in the structure of kinship relations, aimed at producing socially beneficial reciprocal obligations. Richard Titmuss, author of a seminal book on the gift of blood donation, comments that

to give is to receive – to compel some return or create some obligation – either in the form of a similar or different material gift or in the overt expression of sentiment, pleasure or pain, manifested in physical acts of behaviour on the part of the recipient. No such gift is or can be utterly detached, disinterested or impersonal. Each carries messages and motives in its own language. (Titmuss 1997: 227; Arrow, 1972)

What is often described as ‘reciprocal altruism’ in evolutionary biology and game theory bears little resemblance to the moral disposition under review here (Trivers 1971). Quasi-altruistic social transfers are both very common and very valuable in society, but they do not fit the analytical category of an altruistic act precisely because their primary motivation is not a fundamental concern for the other but rather a Humean enlightened self-interest.¹⁰ Once we accept the distinction between altruism and quasi-altruism, many apparent cases of altruistic behaviour turn out not to properly fit the definition.

Altruism and Unintentional Benefits

The same arguments do not affect the case where an agent *unintentionally* benefits from an act that was entirely motivated by a concern for others. Provided the benefit is *genuinely* an unintentional side-effect, there is little reason to doubt the altruistic character of the act. If we were to impose a ‘no-benefit rule’ on altruistic actions we face a number of problems of a predominantly epistemic nature. First, since

individuals have only limited foresight into the effects of their actions, it is conceivable that what started out as an altruistic act – that is, an act entirely motivated to benefit others and not oneself – has long-term consequences that end up benefiting the altruist. Suppose I save a little child from being run over by a speeding car. The boy grows up to become a medical researcher and invents the cure for a disease I end up suffering in my old days. Does my initial altruistic act suddenly become any less altruistic because I happen to end up benefiting from it decades later? A related problem arises in cases where altruism generates an immediate benefit but the altruist is not aware of it, or could not in any way have foreseen it. Suppose I save another little child from a speeding car. This time the boy happens to have extremely rich and grateful parents who insist on depositing a large sum of money into my account. The split-second I decided to pull the boy away from of the car I could not have known that his rich parents would reward me handsomely – in fact, I would hardly have the time to deliberate costs and benefits associated with my action – and thus my action can be said to be genuinely motivated by my concern for the child’s welfare. Again, it would seem rather absurd to insist that because of my benefiting from my actions the act itself does not qualify as an example of altruistic behaviour.

The relation between altruism and unintended benefits is of considerable importance, for genuine altruism may, on balance, produce *systematic* beneficial side-effects for those who exhibit altruistic dispositions. The economist Robert Frank (1988) has advanced an original evolutionary argument to explain why genuine altruists overall do better in a competitive society than non-altruists. According to Frank, people are more likely to cooperate with genuine altruists since the latter are less likely to renege on promises or contracts or pursue adverse rent-seeking strategies. Frank’s argument of course depends crucially on individuals’ capacity to effectively distinguish genuine altruists from quasi-altruists, since the latter have every reason to want to be perceived as genuine by potential cooperators. The adequate perception of altruistic or at least cooperative dispositions is channelled through the public display of emotions, including minute signals sent out by our body language that are notoriously hard to fake.¹¹ In all these examples the key to determining the altruistic character of the respective acts is that *the benefits to the altruist are relevantly independent from the motivation underlying the act*. In each case, it is the failure to observe this

independence, and not the level of benefit to the altruist, that separates altruism from quasi-altruism.

Altruism and Self-Sacrifice

Let us now turn to perhaps the most controversial dimension of altruism. The need for self-sacrifice is often¹² thought of as closely associated with altruism, but I believe this thought rests on a serious confusion between conceptual analysis and valuation. It is of course trivially true that each act of altruism requires some sacrifice: each action has opportunity costs, so an action geared at advancing someone else's well-being implies foregoing options for pure personal gain. But this trivial notion of self-sacrifice is not what is at stake in the altruism debate.¹³ To see this, consider the question whether there is a positive correlation between the degree of self-sacrifice and the degree of altruism exhibited in a particular act. Is the fact that someone incurs additional costs necessarily a reason for thinking that act more altruistic? I believe not. Suppose a friend of yours has fallen on hard times and asks for your help. You are quite keen to help her and figure there are two ways of doing this. The first would make her better off in absolute terms; a second solution would help her out significantly but not quite as much as the first. Assume in addition that the first situation requires less in terms of self-sacrifice compared to the second. You now face a choice between maximising your friend's well-being or maximising the level of your personal effort in helping her. I think the act that maximises your friend's well-being is more altruistic than the one that implies more sacrifice on your part.¹⁴

So why is it widely held that an altruistic act deserves more appreciation when it involves some self-sacrifice than an equally beneficial act that was less demanding? I do not think such an appreciation is misconceived or irrelevant; on the contrary, I wholeheartedly agree that acts of personal sacrifice should be commended and merit social respect. But I simply do not think the sacrifice element renders the act itself any more *altruistic*. In my view, the degree of altruism is solely determined by how strong a person's commitment is to furthering other people's well-being *entirely independent* of regard of her own personal gain or loss.¹⁵ The latter qualification equally implies to the level of personal sacrifice involved.¹⁶ In other words, self-sacrifice is not central but rather epiphenomenal to altruism: while taking into account that certain sacrifices rightly

command admiration and commendation, to determine whether an act is more or less altruistic we only need look at the motivation for furthering other people's good.

Altruism 'Tracking' Moral Demands

In this section I question whether altruism properly 'tracks' moral demands by analysing to what extent it allows – approves or fails to disapprove of – actions that are on some plausible common-sense view considered immoral. Typical examples of altruism include instances of gift-giving – the voluntary¹⁷ transfer of goods and resources from oneself to others – and rescue cases – often involving personal action rather than material aid.¹⁸ In this section I will make ample use of such examples to elucidate my arguments.

Mixed Motivations

Let me begin by introducing a typical example of altruism:

Random Rescue – You are taking a walk along the neighbourhood lake and witness how two individuals fall out of a little boat. They are both clearly in distress and are at risk of drowning, shouting out for your help. You want to help but you face a hard choice: being only a moderately good swimmer with little experience, you guesstimate you can save one of them but not both. You make a random choice and swim towards one victim and pull her out: unfortunately by the time you return to save the other person, he has drowned.

In this crude example there is little doubt that your action is to be recommended. You have acted altruistically by saving one person, at some personal discomfort and possibly even risk, and should be commended for your action. The drowning of the second person is an unfortunate event, but surely there is no way you should be blamed for failing to save him under those conditions. To put it differently, the drowned person has no cause for *complaint* against you for making a random judgement call to save the other person first.

But let us now vary the example a little, by focusing a bit more on the choice of whom to save. Consider the following variation on the story:

Bigot – You are a bigot at heart and you notice that one of the people at risk of drowning is of your own race, while the other is clearly not. Again you jump in and save one person, but in making a choice between them you let yourself be guided by your racist inclinations.

How should we evaluate your action in this case? On the one hand, you deserve commendation for saving one person. Moreover, one could even argue that, on a purely consequentialist basis, your action produced exactly the same outcome as in our initial example: one person is saved, the other unfortunately drowned. Does the fact that your bigotry influenced your choice of whom to save have any further relevance? I think it does once we consider again whether the person that was not saved had a reason for a complaint. It appears she would have, since her being of a different race clearly meant she had less chances of being saved (perhaps no chance, if your bigotry runs deep). And since those reasons themselves would appear to be unacceptable in common-sense morality, that complaint carries serious moral weight.

To fully appreciate the ethical importance of the reasons underlying the choice of which to save, consider two further cases:

Prudence – You make a choice in favour of one person solely based on the grounds that she has drifted less far so the risks of saving her would be substantially smaller, while the chances of succeeding would be higher.

Parent – You see that one of the drowning children is your own child, while the other is that of a friend. You decide saving your own child first.

First compare the situation of Bigot with that of Prudence. Could the drowned person have a possible complaint against the saviour because of a decision weighted by non-moral properties? It would seem farfetched to think he does: it is true that his chances of being saved were effectively reduced, but prudential reasoning in this case seems fully acceptable.¹⁹ The cases of Parent and Bigot perhaps bear closer resemblance in that both judgements are based on non-prudential properties between those in need. The question whether personal attachments such as parent and child, or friends, or perhaps even fellow-nationals may provide sufficient reason for prioritising amongst the

needy is a matter of some controversy.²⁰ But let us assume, following conventional wisdom, that associative attachments and close relationship represent a different category of moral reasons than those based on ‘objective’ properties such as race, class or gender. Thus here too we might conclude that parent was fully justified to save her own child first, given strong maternal ties and the psychological difficulties in disassociating oneself from them, whereas Bigot’s reason for preferring to save the same-race person remain highly suspect.

What are the implications of this discussion for altruism? The first thing to note is that altruism suffers from a serious form of *moral myopia* in that it appears to have few resources to appreciate the problematic issues at hand. It is not just that Bigot has a deplorable disposition that mixes with his (commendable) altruistic outlook; the real issue is that altruism as a moral perspective has no way of judging what has been going on in Bigot’s ‘altruistic’ rescue act. For the altruist, the only thing that counts is a motivation to advance the well-being of (some) other(s): by jumping in and saving one person, Bigot has satisfied this requirement. Any further assessment of the larger picture – in particular, the mixing of ulterior reasons with one’s altruistic disposition – does not feature. The result may be reformulated as a failure to abide by the moral demand to respect a person irrespective of her race, an outcome that discredits the idea that *pro tanto* altruism is inherently desirable.

Mixed Outcomes

However, we do not necessarily have to rely on morally suspect ulterior motives to appreciate the limited scope of altruism. Consider yet another story:

Unbalanced Gifts – Suppose a society of altruists have a convention in which those who are rich regularly donate part of their wealth to some of those who are poor. But let us now assume that those altruists are only governed by the motivation to donate to someone who is less well off, and employ a completely random selection of whom to pass their wealth to. The outcome of this gift-giving is an unequal distribution in which some of the poor receive multiple gifts while others are left without a gift.

In this story there is no problem with unacceptable ulterior motivations. The resulting outcome is entirely caused by what might be called *systematic bad luck* (Dowding 1996). In reality many examples seem to

fit with the presence of systematic bad luck. For instance, the fact that few women are actively present in electoral politics may have nothing to do with preferences or actions from their male colleagues, but may be due to a social environment that systematically singles out properties that favour male politicians (Phillips 1995). Similarly, labour markets often discriminate against people who have no access to a network of 'weak ties', without there being any instance of wrong-doing on the part of employers or fellow job-seekers (Granovetter 1974).

Conventional wisdom and many moral and political theories, in particular those rooted in impartiality or substantial equality, would consider this a deplorable outcome despite the absence of bad actions. The fact that the rich are willing to share their wealth is undoubtedly good, but the actual distribution runs afoul of what a more egalitarian or impartial account would require. Altruism's focus on strict motivations again appears to lack the necessary resources to fully appreciate the situation. The supposedly inherently desirable disposition towards altruism may in fact be instrumental in maintaining an ethically suboptimal outcome, which in turn might be considered a failure to observe particular moral demands.²¹ This again contradicts the assumption that *pro tanto* altruism always contributes to the good.

The examples in this and the previous subsection establish ways in which altruists fail to fully observe moral demands. To put it differently, each case could be interpreted as grounds for some agents to lodge a complaint *against altruists*. Next I examine ways in which altruists would be entitled to lodge a complaint *against others* for not observing moral demands.²² Some readers may object to this because they believe that altruists can never be said to have cause for complaint since, by definition, they gain personal satisfaction from other agents' increased satisfaction.²³ I disagree, because altruists, like any other person, might be susceptible to perverse pressures to adapt their preferences. The hallmark of adaptive preferences is that they are often considered deplorable on an objective account precisely because they render an agent worse off than they would have been in the absence of that particular preference or disposition. An agent who has an altruistic disposition may not lament unfair demands but might still deplore the fact that she actually has an altruistic disposition. In fact, as I will show, it is merely another example of altruism's limited scope that altruists often cannot properly assess whether their action presents a case of taking on too many burdens.

Partial Compliance Dilemma

Let us consider a world in which two facts hold. First, fully satisfying moral demands requires collective action. Individually, moral demands impose restrictions on agents' actions – sometimes requiring them to act against their immediate self-interest, and modify or perhaps even abandon some personal projects – but since few (if any) moral demands can be fully satisfied by any single person, we need to think about moral demands as requiring moral agents to cooperate to collectively ensure demands are satisfied.²⁴ Second, far from constituting a one-motivational selfish or saint-like society, I assume a more plausible world that contains something like a normal distribution of self-centred 'knaves' and radically other-regarding 'knights' at either extreme, with a large number of mixed-motive moderate altruists (or 'moderates') in between. Moderates are individuals who respond altruistically or self-interestedly depending on certain properties of the situation (LeGrand 1997; Frey 1997).²⁵

Consider first an example in which agents deliberate non-strategically by assessing moral demands strictly in terms of the level of what needs to be done.

Poverty Relief – To ensure that not a single individual ends up living below the poverty level, a society with 100 rich individuals requires each of them to donate 1% of their wealth to the poor. If only the altruists (knights) among them comply with this demand, they would have to donate 20% of their wealth.

Given a fixed background distribution of motivational dispositions, it becomes relatively easy to predict that knaves will largely neglect moral demands, knights most likely will behave altruistically, and the large contingent of moderates acting will end up wavering between altruistic or self-interested action. So in the context of a normal distribution of moral dispositions, altruists or knights encounter the familiar problem of partial compliance, which might lead to the following moral dilemma. All those who act altruistically succeed in jointly satisfying all the moral demands – leaving no real 'moral deficit', as it were. This effectively requires that some agents take on more than their fair share of the burden in the face of others shirking their moral obligations. Alternatively, altruists abide by what Liam Murphy (2000: 76–82) has called the *compliance condition*. The compliance condition requires that

agents respecting moral demands under partial compliance should not be left worse off compared with a situation in which everyone would comply fully: it ‘gives expression to the idea that an agent should not be required to “take up the slack left by the non-compliance of others”’ (Murphy 2000: 77).²⁶ But remaining within the boundaries of the compliance condition of course means that not every moral demand will be satisfied, and a second type of unfairness ensues because now some legitimate moral claims are not addressed. The partial compliance dilemma forces altruists as it were to make a choice between ‘unfairness in the distribution of benefits’ and ‘unfairness in the distribution of burdens’.

Partial compliance with moral demands also threatens the stability of an altruistic morality to the extent that most agents are in fact ‘conditional co-operators’ who, roughly put, comply with norms and rules provided a sufficient number of others are doing so as well.²⁷ Significant failure to comply might then have a negative snowball effect in which, over time, all but the most extreme moral saints become non-compliers. This obviously represents the worst of possible outcomes, for it *combines* unfairness in the distribution of benefits *with* unfairness in the distribution of burdens.²⁸ In such a situation it is primarily altruists who can lodge a complaint against those who shirk their moral obligations and initiate the vicious circle of ever-decreasing compliance. But this depends on altruists being able to discern whether others have shirked their obligations, which would require some account of a fair distribution of burdens in the first place – an account that is lacking in the altruist’s moral outlook. And even if altruism could provide such a theory this would not mean that it could prevent altruists from taking on more than their fair share in virtue of being possessed with a disposition that attaches more importance to what happens to others than what happens to oneself. I shall return to this issue below.

Samaritan’s Dilemma

The partial compliance dilemma arises when agents deliberate parametrically, only looking at the demands to be fulfilled and not really taking into account the distribution of motivations of other agents. Next, consider a case where agents do not merely deliberate parametrically but also strategically. Once strategic interaction is introduced, altruistic morality becomes caught up in another famous social dilemma, introduced by James Buchanan in the 1970s (Buchanan 1975).

The samaritan's dilemma is a *moral hazard problem* rooted in the interaction of differentially motivated agents under conditions of imperfect information. The basic problem is as follows: agents who have a reasonable expectation of altruists assisting them in case of need may alter their own behaviour to shift the costs of preventing needs onto altruists.²⁹ Buchanan's analysis applies as much to one-to-one samaritan acts as to public systems of 'samaritan support', where the state takes on the role of a central altruist – typical examples include welfare assistance, health care, pension saving and so on. An often debated example would be the welfare parasite:

Welfare Parasite – I know that if I do not finish school and get a proper qualification I may seriously jeopardise my chances of getting a well-paid job. But I happen to live in a society with generous unemployment benefits and other social assistance, so I do not actively pursue a degree or qualification. When later in life I cannot get a job as predicted, I happily rely on society's generous assistance.

The samaritan's dilemma seems to take unfairness one step further by producing manifest exploitation. Free-riders who intentionally shift the costs of their actions onto those providing assistance do not merely fail to respond appropriately to moral obligations, but actively cause harm by externalising costs.³⁰

In addition to the prospect of manifest exploitation, the samaritan's dilemma too poses a considerable threat to the stability of altruistic morality. Where moderate altruists consider shifting towards non-compliance in cases where not enough other people comply because of a lack of solidarity, they are certainly likely to feel taken advantage of when realising that it is now the very individuals they are assisting that take advantage of their altruistic dispositions. Altruists, again, have little defence against the samaritan's dilemma: their disposition to care for others does not as such take strategic considerations into account. Once altruists start behaving strategically as well – a precondition for Buchanan to remedy the dilemma – they run the danger of actually having to modify their altruistic dispositions.

Embedding Altruism

In all the examples discussed above, the problem for altruism is clearly one of limited scope. The problems identified above are all symptoms of

altruism's common illness: a failure to perceive the full extent of moral demands. Altruism's significant blind spot when it comes to identifying objective moral demands has at least two important implications. Since the altruist's reasons for action are solely rooted in a concern for the well-being of others (without, it is important to keep in mind, affirming any moral theory about the need to further other agents' well-being), the more altruistically an agent is disposed the more she will try to further others' well-being. But the possession of an altruistic disposition is at the same time the only source of moral obligation; it therefore does not say anything about obligations for agents who are not similarly disposed. More worryingly, *pro tanto* altruists *assume* that their disposition is inherently desirable because it furthers the good; but the various examples above show that this is not self-evidently true. One key problem for altruists, then, is that they actually have no proper account of *why we should value altruism as a motivational disposition in the first place* if it turns out not always to deliver the good, understood as an appropriate response to moral demands.³¹ Altruistic dispositions produce a 'push' for agents to behave with concern for others, but this push remains essentially without direction – or rather, its direction might be determined by factors that form no part of a comprehensive moral theory – which explains why it fails to respond appropriately to a number of moral demands.³²

The solution, I propose, is to *embed* altruistic dispositions into a more comprehensive moral theory. Such a theory would need to meet two requirements. First, it needs to have the capacity to determine and justify a set of moral demands, in addition to detailing how these demands give rise to various duties and obligations that agents are required to meet. Second, such a theory ideally has room for altruistic dispositions to operate within the constraints set by the theory, and at the same time explain why altruistic dispositions are valuable and should be actively promoted – again, *within the confines of this theory*.

The Theory of Impartiality

One plausible candidate for the job, I think, is the political theory championed by Brian Barry (1995), justice as impartiality. Impartiality is a moral and political theory that captures a particular view of how people should relate to one another. Barry (1995: 7) writes: 'Principles of justice that satisfy its conditions are impartial because they capture a certain kind of equality: all those affected have to be able to feel that they

have done as well as they could reasonably hope to.’ Barry readily admits that this account presumes a number of substantive moral ideas, while resisting any consequentialist reading of his theory. Impartiality as a substantive moral idea clearly provides ‘direction’ in practical reason (although many complexities enter the picture once we try to figure out precisely what it is that impartiality requires of us). As such it appears to provide the sort of ‘pull’ that might ‘direct’ altruism to a more acceptable position (provided of course one accepts impartiality as a good moral theory).

However, at this point an important objection springs to mind: if the altruist is now motivated by impartiality, does this not render altruistic dispositions obsolete? If this were the case, it would amount to a failure of the second condition specified above and effectively render impartiality unsuitable as a candidate for embedding altruism. To counter the objection, we need to turn to a crucial distinction between *first-order* and *second-order* impartiality. In *Justice as Impartiality*, Barry (1995: 11) explains the distinction as follows:

What the theory of impartiality calls for are principles and rules that are capable of forming the basis of free agreement among people seeking agreement on reasonable terms. If we call impartiality in this context second-order impartiality, we can contrast it with first-order impartiality, which is a requirement of impartial behaviour incorporated in a precept. Roughly speaking, behaving impartially here means not being motivated by private considerations.

Barry furthermore argues that second-order impartiality does not require universal first-order impartiality, which opens up opportunities for combining partial considerations at the level of day-to-day decision making with the requirement that our background rules are governed by the principle of impartiality.³³ If we concede, with Barry, that impartiality only requires us to accept and promote impartial rules, regulations and procedures at the ‘constitutional’ level, this means altruists can go about their daily business being motivated by their altruistic dispositions. In other words, there is no conflict between altruism – which is essentially a first-order account of morality – and second-order impartiality.

Of course, once second-order impartial rules are in place and properly instituted, they will mediate and to some extent constrain altruists’ range of actions. But this is of course precisely what we are

looking for: second-order impartial mediation can be interpreted as a form of *self-binding* that prevents altruists and other agents from violating in their day-to-day engagements the very principles of impartiality they endorse at the level of abstract principle.³⁴ Once designed and implemented, impartial social institutions may significantly raise the costs of acting in a partial way, or even eradicating the option of doing so altogether, thus countering overly partialist outcomes of actions initiated by altruists.³⁵

How does an 'impartially embedded altruism' fare with respect to the moral and social dilemmas discussed before. First, in the case of 'mixed motivations', impartial background rules would clearly rule out the sort of special privilege of Bigot. One could doubt whether in that particular story an impartial morality could properly regulate an altruist's behaviour. To those critics I submit another example that is relevantly similar to the rescuing case. Suppose Bigot indicates that he is willing to donate his organs after death for transplant purposes, but specifies that they should only be used for same-race transplants.³⁶ Second-order impartiality could easily render such a decision legally invalid, offering the altruist a mere choice whether to donate his organs or not. Second, similar arguments apply to the case of 'mixed outcomes': even when no unacceptable motives are present, impartial regulations might be instituted to prevent the sort of hidden discrimination we still come across with respect to, say, the position of women or racial minorities in politics or the labour market (Barry 1995). In both cases, the altruist's dispositions to care for others are left intact while being simultaneously mediated by impartial second-order considerations.

Let us now consider the final two cases. First, in my view impartially embedded altruism provides a good answer to the problem of partial compliance. When impartial institutions properly enforce compliance with general rules, they make it harder for knaves to shirk their social obligations. This has not only implications concerning the prevention of knights being unfairly disadvantaged, but has also important consequences with respect to the stability the social system. A reliable system of rules provides just the sort of reassurance that moderate altruists require in order to keep complying with moral demands.³⁷ One might still wonder why we need impartiality to set enforceable rules and not simply argue the case for 'altruistic' regulation. The reason this option is a non-starter is precisely because the first-order/second-order distinction simply does not apply to altruism. There is something

distinctively odd about 'being forced to behave altruistically'; either one has an altruistic disposition or not.³⁸

Finally, in the case of the 'samaritan's dilemma', impartial rules provide the sort of solution Buchanan himself advocates: social institutions restrain the free-riding option of claimants or recipients, thus preventing altruists from being exploited and ensuring stability.³⁹ Now, nothing prevents one from adopting an altruistic perspective and, at the same time, recommending rules that prevent the sort of free-riding that produces the samaritan's dilemma.⁴⁰ After all, altruism is not incompatible with forcibly constraining free-riding recipients. Hence one might conclude first-order altruism is not really threatened by the samaritan's dilemma after all. But perhaps there remains a subtle but important distinction between altruistic and impartial theories of morality. An important question in the regulation of the samaritan's dilemma concerns the sort of rules or regulations one needs to introduce to achieve the objective. Impartialist moral theory has the resources to impose a set of principles that apply across the board by setting a common standard of morality for knaves, knights and (potentially free-riding) claimants alike. Each set of rules, even when applying to different agents, must be justified by the same impartial background considerations. Altruism, on the other hand, in effect needs to make up an entirely novel set of principles for regulating the behaviour of claimants. Impartiality has the advantage of being the more parsimonious of both theories in this respect. In a sense, the last point gets us back to an argument introduced earlier – to wit, altruism lacks an objective account of moral demands.

Why (and How) to Value Altruism?

One final aspect of impartially embedded altruism needs to be addressed. So far I have argued that embedding altruism resolves the sort of problems detailed before, but I have not yet said anything why an altruistic disposition should be valued within an impartial theory. One could after all argue that it is the second-order impartial considerations that seem to be doing all the ethical work in an impartially embedded altruism. Rejecting altruism clearly flies in the face of common-sense morality, which, across the board, deeply values altruistic acts, but this does not preclude the possibility that common-sense morality may be wrong. Is there a way to justify positively valuing altruism within a theory of impartial morality? I believe there is, and in the remaining section I shall briefly provide one possible argument.

Impartiality sets out providing a set of rules, principles and maxims that form a background against which agents can pursue their personal goals – whatever they are. Second-order impartial principles also require a set of social mechanisms that ensure that agents will operate within the constraints imposed by considerations of impartiality. Possible mechanisms could include a system of monitoring in combination with material penalties in case of transgression, but an equally plausible arrangement is to employ social carrots and sticks such as esteem and disesteem (Brennan & Pettit 2000). Unfortunately, institutions often track moral principles only imperfectly, leaving a niche for agents non-complying with moral demands. A possible solution might be to employ individual's intrinsic motivation to care for one another as a further 'proximate mechanism' to ensure impartiality (Sober and Wilson 2000: 109–122).

Under these conditions, altruism has a significant beneficial effect to the extent that the more people are altruistically motivated the less they will take advantage of niches for opportunistic behaviour. A good measure of altruism improves the effectiveness of impartial rules and regulations. A second reason why altruism might have a beneficial impact on second-order impartiality has to do with the efficient organisation of impartial rules. Minimising opportunities for free-riding or otherwise non-complying with moral demands might prove to be a task that becomes self-defeating in terms of implying extensive moral and organisational costs – for instance, immense monitoring is often associated with considerable intrusion in our private lives. Altruistic dispositions might contribute to the efficiency of impartial regulation by allowing comparatively less intrusive monitoring. We should not underestimate the extent to which moral motivation contributes to the co-ordination of behaviour (Frey 1997), even within social institutions that combine screening, monitoring and sanctioning to ensure compliance with moral demands as fully as possible, and should actively make use of those motivations wherever possible. I think this provides a good *instrumental* rationale for a theory of impartiality to value altruistic acts.

Clearly, if we accept that altruistic dispositions play an important role at the level of proximate mechanisms, it appears impartial regulators have every reason to actively promote altruism. Unfortunately it is not the case that altruism can simply be plugged into any kind of regulatory system. In recent years scholars have warned against the tendency of

policy makers to focus solely or even primarily on external incentives – principally financial incentives – to ensure compliance, arguing that a system that relies heavily on such external incentives runs a serious risk of *crowding out* intrinsic altruistic motivations (Frey 1997). This point was already made forcefully in Richard Titmuss’s (1997) study on blood donation, but the force of his insight is increasingly appreciated across many policy areas (Le Grand 1997). Most studies now accept that there are important *demonstration effects* to consider in the design of social institutions (e.g., Jones and Cullis 2000). A policy design that is sensitive to the effects of crowding out must clearly balance two related concerns. First, what sort of policies are immediately required as a matter of impartial justice, given a particular distribution of knaves, knights and moderates? Second, how do we ensure impartiality today with concerns regarding the stability of future arrangements, given the dynamics of crowding in/crowding out? Managing these potentially conflicting policy goals requires an approach that exhibits sensitivity to variation and contingency in background circumstances.⁴¹ It is not the purpose of this essay to further consider these issues, but moral theorists should be aware that there are important practical considerations at play when trying to embed altruism within impartiality.

Conclusion

This essay is informed by the conviction that a theory of morality built on altruistic foundations is a non-starter, and that contemporary political philosophy has rightly replaced a concern for altruism with a concern for impartiality.⁴² I have made two main claims. First, I argued *pro tanto* altruism, and its insistence on every altruistic act being inherently desirable, is untenable. Next I maintained the problems associated with *pro tanto* altruism could be mediated by embedding altruism within a more comprehensive moral theory, and I suggested Brian Barry’s (1995) theory of impartiality might be a suitable candidate for this job. While the upshot of this essay is to argue against the foundational status of *pro tanto* altruism, I do not in any way intend to deny the fact that altruism could play an important role as a proximate mechanism. Once embedded within a robust theory of impartiality, an agent’s altruistic dispositions are safeguarded against the problems and dilemmas identified in this essay. In addition, I believe altruism has a significant positive impact on the fairness and stability of a society governed by

impartialist principles, and for that reason deserves to be valued and promoted.

ACKNOWLEDGEMENTS

I am grateful to the Hoover Chair in Economic and Social Ethics, Catholic University of Louvain, Louvain-la-Neuve for providing a stimulating environment in which to write this essay, and to Els Compernelle, Cillian McBride, Jonathan Seglow, Lindsay Stirton, Philippe Van Parijs and the participants of a Hoover Chair seminar for useful comments on a previous draft. Finally, thanks are due to Jonathan Seglow for being an extremely patient editor.

NOTES

1. See e.g. David Gauthier (1986) on 'constrained maximization', or Nozick (1974) on rights as side-constraints on personal action.
2. There is a relevant distinction between requirements being *motivationally* demanding (the sense employed above) and requirements being *normatively* demanding. In the latter usage we are concerned about satisfying one set of demands to the detriment of another valuable set. An example here is Bertrand Williams' (1973) integrity objection against utilitarianism, or Railton's (1984) alienation objection against consequentialism.
3. Accepting a rule or maxim implies more than simple compliance. For discussion, see Hooker (2000: 75–80).
4. The same goes for people who deny part of their tax obligations precisely because they believe the goals on which the money will be spent to be in violation of the requirements of justice. Consider e.g. a citizen who refuses to pay that fraction of tax that would go towards defence spending.
5. This distinction has been invoked by a number of moral philosophers. David Schmidtz (1993: 52–3) distinguishes 'concern' and 'respect' for others, while Susan Wolf (1997: 80ff) talks about the 'Loving Saint' versus the 'Rational Saint'. In a review of Titmuss' *Gift Relationship*, economist Kenneth Arrow (1972: 348) differentiates between motives that cause a person to derive 'a utility from seeing someone else's satisfaction increased' from the quite different set of motives based on 'an implicit social contract such that each performs duties for the other in a way calculated to enhance the satisfaction of all'.
6. Note that in social policy Type I and II errors are typically linked in a way that means that they need to be traded-off against one another. I do not want to suggest a similar trade-off in the case of moral demands.
7. It could be argued that altruism is a *virtue*, since this rough-and-ready definition of altruism accords with Philippa Foot's (1997: 165) description of a virtue as a set of 'general beneficial characteristics' that 'a human being needs to have, for his own sake and that of his fellows'.
8. The association of *pro-tanto* altruism with reasons that constitute *genuine* but not necessarily *decisive* reasons is of course informed by David Ross's (2002: 16–47) famous discussion of *prima facie* duties as conditional reasons. Many commentators accept what Ross himself already hinted at, namely that the term *prima facie* is misleading in that it refers to an epistemological rather than a normative problem. For discussion, see Philip Straton-Lake's excellent introduction to the 2002 edition of Ross

(2002); also Kagan (1989 :17).

9. Evolutionary biologist E.O Wilson (1978) refers to 'soft-core altruism'.
10. In contemporary society a common form of quasi-altruism is the provision of commodities with public goods properties (they appear to represent a case of altruism when its provision is taken on by a limited number of individuals). But for cases where consumption is extended beyond those who produce the good, it is typically the supplier's inability to restrict such consumption, rather than genuine concern for the other's welfare, that motivates a unilateral provision of a public good – though, admittedly, it may be good political strategy to mask one's real self-interested intentions.
11. Cf. the 'costly-to-fake' principle (Frank 1988: 99–102, 109–13).
12. Often but not always: Jean Hampton (1993: 136), for instance, suggests that 'not all self-sacrifice is worthy of our respect or moral commendation'. This is because the act of persistently 'denying' oneself is not necessarily a morally worthy goal. See on this also Wolf (1997).
13. Schmidt (1993: 65) rightly points out that 'cost-bearing becomes self-sacrificial only when the agent prefers the value foregone to the value gained', which is evidently not the case with altruists. He furthermore suggests that altruism 'is necessarily self-sacrificial only for purely self-regarding agents', but I find this confuses the issues since clearly self-regarding agents must have a good ulterior motive to engage in actions that may be *perceived* as altruistic (cf. the discussion of quasi-altruism above).
14. There are obvious epistemic problems with determining what maximises someone's well-being. However, since altruism is concerned with motivation rather than outcome, we can allow room for genuine errors of judgement.
15. Jonathan Seglow suggested to me that self-sacrifice becomes relevant when we look at altruism from a comparative perspective. Consider the following example: a billionaire donates £1,000 to an orphanage and a homeless person offers his only blanket to a fellow homeless during a cold night. Both acts are altruistically motivated, but while the billionaire maximises overall well-being, the homeless person's act implies a higher sacrifice. Comparing both we would intuitively favour the homeless person as performing the most altruistic act, and this might cast doubt on my earlier insistence of the irrelevance of self-sacrifice. However, I do not think the example invalidates my argument. If we were really to insist on ranking altruistic acts across agents, we presumably would have to first weight them according to some sort of 'altruistic capacity scale'. My claim would be that, even when using such a weighted scale, maximising well-being remains the proper conceptual dimension along which to assess the altruistic nature of a particular act.
16. The argument could be put the other way around. Consider the case of someone with a masochistic nature who likes to assist people because it involves a sacrifice on his part; or the example of a thrill-seeker who loves to defy death in rescue operations because this is his preferred way of satisfying his craving for danger. Both cases score highly on the sacrifice dimension, but neither of these people are considered altruists precisely because the genuine commitment to furthering another person's good is lacking.
17. It seems counter-intuitive to argue that involuntary transfers – 'coercive altruism' – constitute a proper form of altruism. But it should be kept in mind that many first-order coercion acts are not coercive at all if they are agreed to at a higher level, and I shall use the term to refer to both first-order and second-order types of voluntariness.
18. One should be careful not to read too much into this distinction. Both 'transfer' and 'rescue' cases share many essential features, such as the possibility of opportunity costs or even self-sacrifice, and there are many examples that could fit both categories – e.g., private contributions to foreign aid.

19. One could perhaps even argue that the choice that maximises the probability of saving at least one person is not only acceptable, but in fact morally required. Prudential and moral reasons often intersect in this way.
20. The *locus classicus* is of course Godwin (1976).
21. If one accepts the idea of *negative responsibility*, altruists may be held 'passively' responsible for either not being more discriminate in how to spread their gifts or perhaps even not instituting an alternative arrangement which satisfies a more egalitarian distribution of gifts. On the notion of negative responsibility see in particular Nagel (1991: 99–102, 166–7).
22. So the next subsections focus on the other side of the Type I and II error analogy mentioned above.
23. The standard economists' interpretation of altruism presumes 'interdependent preferences' (Collard 1978: 18–29).
24. This point has been most fully analysed by Russell Hardin (1988; 1991).
25. The particular distribution of knaves, knights and moderates might differ considerably across time, across groups of individuals and even across particular contexts. But this variation does not detract from the main argument, which is based on the assumption of motivational variation itself.
26. See also Murphy (1993) and Mulgan (2001) for further discussion.
27. This social phenomenon is sometimes referred to as the Hume-Titmuss *contagion thesis*, after its classical sources (Collard 1978: 11). For further discussion, see also Miller (1988).
28. Altruism thus can be conceived as a 'public good'. See Collard (1978) for discussion.
29. In an extension of the reasoning behind the Samaritan's dilemma, known as rent-seeking theory, Buchanan's then collaborator Gordon Tullock relates how in some cases such behaviour does not only impose costs on others, but even on the agents themselves. Tullock (1980: 275) relates the anecdote of how Chinese beggars would 'deliberately and usually quite horribly mutilate themselves in order to increase their charitable take'.
30. I here refer to the definition of exploitation employed by Gauthier (1986) and further developed with respect to welfare policies by Van Donselaar (1997; 2003). Buchanan (1975: 82) also argues that altruists' behaviour may have further externalities towards the wider community by imposing potential costs on other potential samaritans.
31. This of course does not in itself imply that the altruistic act should not be performed *even if it does not fully track moral demands*. It would certainly go too far to suggest that Bigot should be actively prevented from jumping in the water: clearly, saving one person is always better than saving none. But it is equally conceivable that society institutes a set of rules that may cause some altruistic acts to be prevented, based on the view that the short-run costs of restricting altruism are outweighed by the long-term benefits of countering racial discrimination.
32. I believe my claims are closely related to Thomas Nagel's view in *Equality and Partiality*. 'Ethics does not license an *unmediated* universal altruism, precisely because that leads in ordinary circumstances of disagreement over the nature of the good to inevitable conflict rather than possible unanimity. Mere altruism, surprisingly enough, does not provide a common standpoint from which everyone can reach the same conclusions' (Nagel 1991: 164, italics added).
33. For discussion of the relation between the first and second level, see also Mendus (2002), and on combining partiality and impartiality, Nagel (1991).
34. On this see in particular Elster (1979). An example would be the vegetarian who abhors the industrial breeding and killing of cattle for human consumption, but is often tempted by the smell of sausages to give into his cravings. In this example a self-binding mechanism might be to tell all your friends about your views so that sheer social

- control raises the consequences of giving in to temptation. To suggest that those who accept impartiality as a valid moral account would have little reason not to abide by them would be to seriously misjudge human psychology. Thomas Nagel (1991: 86) has rightly stressed that 'the motivational problems connected with acceptance of a general social framework as legitimate are different from the motivational problems that arise for individuals acting within it'. Individuals are often motivationally ill-equipped to withstand the many easy ways out of hard situations.
35. Rational choice theory provides us with one powerful mechanism of altering choice sets by raising the relative costs of adopting one option over another. Another possibility is taken from law: in cases where rules create options, changing a rule might effectively eliminate that option from one's choice set.
 36. A few years ago debate on this issue raged in the UK. In 2001 the Council of Europe adopted a Recommendation (Rec 2001/5) on the management of organ transplant waiting lists and waiting times, specifying that the criteria for inclusion 'should ensure there is no discrimination on the grounds of race, religion, disability or any other non-medical factor. Priority on the waiting list such as 'urgent' or 'very-urgent' categories should be based solely on medical factors relating to the severity of risk for the individual patient.' Contrast this recommendation with the case of California State Senator Jeff Denham, who earlier this year proposed a bill to give California organ donors the freedom to decide that their organs not go to inmates. Denham argued that donors would like to know that their deaths have brought lives to those who will contribute to society, and not to those behind bars who have not only failed at that task but have instead brought harm.
 37. See Collard (1978) on modelling altruism as an Assurance Game.
 38. For further discussion, see Jonathan Seglow, 'Altruism and freedom', pp.145–63, this volume.
 39. 'There may be no escape from the generalised samaritan's dilemma, in its public form, except through the collective adoption and enforcement of rules that will govern individual situational responses. As they are applied, such rules must be coercive, and they must act to limit individual freedom of action' (Buchanan 1975: 81).
 40. A further reason for countering free-riding behaviour is rooted in paternalistic concerns. Tyler Cowen (1993: 238) writes that a 'true altruist may sometimes precommit to not helping an individual he or she cares about' because this would maximise that person's well-being in the long run. Paternalist arguments of this type are often advanced at the level of public policy, chiefly by thinkers associated with the New Right.
 41. Amongst political theorists, Philip Pettit expresses an extreme sensitivity to these issues. Compare his recent discussion of 'motivating' versus 'managing' strategies in Pettit (2002).
 42. This is *inter alia* demonstrated by the relative lack of scholarly work on altruism as compared to impartiality.

REFERENCES

- Arrow, K. 1972. 'Gifts and exchanges', *Philosophy and Public Affairs*, 1/4, 343–62.
- Barry, B. 1995. *Justice as Impartiality*. Oxford: Oxford University Press.
- Brennan, G. & P. Pettit. 2000. 'The hidden economy of esteem'. *Economics and Philosophy*, 16/1, 77–98.
- Buchanan, J., R. Tollison & G. Tullock, eds. 1980. *Toward a Theory of the Rent-Seeking Society*. College Station: Texas A&M University Press.

- Buchanan, J. 1975. 'The Samaritan's Dilemma'. Phelps 1975: 71–85.
- Collard, D. 1978. *Altruism and Economy: A Study in Non-Selfish Economics*. Oxford: Martin Robertson.
- Crisp, R. & M. Slote, eds. 1997. *Virtue Ethics*. Oxford: Oxford University Press.
- Cowen, Tyler. 1993. 'Altruism and the argument from offsetting transfers'. Paul et al. 1993: 211–15.
- Dowding, K. 1996. *Power*. Maidenhead: Open University Press.
- Dowding, K., J. De Wispelaere & S. White, eds. 2003. *The Ethics of Stakeholding*. Basingstoke: Palgrave/Macmillan.
- Elster, J. 1979. *Ulysses and the Sirens. Studies in Rationality and Irrationality*. Cambridge: Cambridge University Press.
- Foot, P. 1997. 'Virtues and vices'. Crisp & Slote 1997: 163–77.
- Frank, R. 1988. *Passions within Reason: The Strategic Role of Emotions*. New York: W.W. Norton.
- Paul, E., F. Miller & J. Paul, eds. 1993. *Altruism*. Cambridge: Cambridge University Press.
- Frey, B. 1997. *Not Just For the Money: An Economic Theory of Personal Motivation*. Cheltenham: Edward Elgar.
- Gauthier, D. 1986. *Morals by Agreement*. Oxford: Oxford University Press.
- Godwin, W. 1776. *An Inquiry Concerning Political Justice*. Harmondsworth: Penguin.
- Goodin, R.E. 1985. 'Erring on the side of kindness in social welfare policy'. *Policy Sciences*, 18/3, 141–56.
- Granovetter, M. 1974. *Getting a Job: A Study of Contacts and Careers*. Cambridge, MA: Harvard University Press.
- Hampton, J. 1993. 'Selflessness and the loss of self'. Paul et al. 1993: 135–65.
- Hardin, R. 1988. *Morality Within the Limits of Reason*. Chicago: University of Chicago Press.
1991. 'Acting together, contributing together'. *Rationality and Society*, 3/3, 365–80.
- Hooker, B. 2000. *Ideal Code, Real World. A Rule-consequentialist Theory of Morality*. Oxford: Clarendon Press.
- Jones, P & J. Cullis. 2000. 'Individual failure and the analytics of social policy'. *Journal of Social Policy*, 29/1, 73–93.
- Kagan, K. 1989. *The Limits of Morality*. Oxford: Clarendon Press.
- Korsgaard, C. 1996a. 'The normative question'. Korsgaard 1996b: 7–48.
- 1996b. *The Sources of Normativity*. Cambridge: Cambridge University Press.
- LeGrand, J. 1997. 'Knights, knaves or pawns? Human behaviour and social policy'. *Journal of Social Policy*, 26/2, 149–69.
- Mauss, M. 1954. *The Gift*. London: Cohen & West.
- Malinowski, B. 1932. *Argonauts of the Western Pacific*. London: Routledge & Kegan Paul.
- Mendus, S. 2002. *Impartiality in Moral and Political Philosophy*. New York: Oxford University Press.
- Miller, D. 1988. 'Altruism and the welfare state'. Moon 1988: 163–88.
- Moon, J.D., ed. 1988. *Rights, Responsibilities, and Welfare*. Boulder: Westview Press.
- Mulgan, T. 2001. *The Demands of Consequentialism*. Oxford: Clarendon Press.
- Murphy, L. 1993. 'The demands of beneficence', *Philosophy and Public Affairs*, 22/4, 267–92.
2000. *Moral Demands in Nonideal Theory*. New York: Oxford University Press.
- Nagel, T. 1970. *The Possibility of Altruism*. Princeton: Princeton University Press.
1991. *Equality and Partiality*. New York: Oxford University Press.
- Nozick, R. 1974. *Anarchy, State, and Utopia*. Oxford: Blackwell.
- O'Neill, O. 2000. *Bounds of Justice*. Cambridge: Cambridge University Press.
- Pettit, P. 2002. *Rules, Reasons, and Norms*. Oxford: Oxford University Press.
- Phelps, E., ed. 1975 *Altruism, Morality and Economic Theory*. New York: Russell Sage.

- Phillips, A. 1995. *The Politics of Presence. The Political Representation of Gender, Ethnicity, and Race*. Oxford: Clarendon Press.
- Railton, P. 1984 'Alienation, consequentialism, and the demands of morality'. *Philosophy and Public Affairs*, 13/2, 174–231.
- Rawls, J. 1999. *A Theory of Justice*. Rev. ed. New York: Oxford University Press.
- Ross, D. 2002. *The Right and the Good*, 2nd ed., ed. P. Stratton-Lake. Oxford: Clarendon Press.
- Schmidtz, D. 1993. 'Reasons for altruism'. Paul et al. 1993: 52–68.
- Smart, J.C.C. & B. Williams. 1973. *Utilitarianism: For and Against*. Cambridge: Cambridge University Press.
- Titmuss, R. 1997. *The Gift Relationship: From Human Blood to Social Policy*, Ed. A. Oakley & J. Ashton. New York: New Press.
- Trivers, R. 1971. 'The evolution of reciprocal altruism'. *Quarterly Review of Biology*, 46, 35–57.
- Tulloch, G. 1980. 'The cost of transfers'. Buchanan et al. 1980: 269–82.
- Van Donselaar, G. 1997. *The Benefit of Another's Pains*. Amsterdam: University of Amsterdam.
2003. 'The stake and exploitation'. Dowding et al. 2003: 94–113.
- Williams, B. 1973. 'A critique of utilitarianism'. Smart & Williams 1973: 77–150.
- Wilson, E.O. 1978. *On Human Nature*. Cambridge, MA: Harvard University Press.
- Wolf, S. 1997. 'Moral saints'. Crisp & Slote 1997: 79–98 .