UC Berkeley

ISUS-X, Tenth Conference of the International Society for Utilitarian Studies

Title

Consequentialism and Moral Responsibility

Permalink

<https://escholarship.org/uc/item/3r95842n>

Author

Mason, Elinor

Publication Date

2008-09-04

Consequentialism and Moral Responsibility

Elinor Mason, University of Edinburgh

Abstract

I aim to vindicate the thought that there is something very unsatisfactory about objective consequentialism – moral theories just can't be like that. The root of this thought is to be found in our intuitions about the nature of moral responsibility, so in order to vindicate it we have to investigate moral responsibility. I argue that moral responsibility and normative theory are inextricably linked. R.J. Wallace has defended a related view, but Wallace claims that investigation of moral responsibility must proceed by investigating our normative notion of *fairness*. I argue that we should not assume that fairness – only one particular aspect of our normative moral theory - is what should inform our account of moral responsibility. I show that our questions about responsibility are really questions about morality, and correspondingly, our questions about morality are also questions about responsibility. Thus plausible answers to each of these questions must be plausible answers to both, and this provides interesting constraints on our theories.

I thus show that objective consequentialism does not have its own corresponding theory of responsibility, and that is why it seems unsatisfactory. Subjective consequentialism is in fact a variant of a basic moral outlook that is Kantian. In Kantian terms, what matters is that the will is good – according to a subjective consequentialist, the good will is one that forms consequentialist intentions.

This paper is part of a larger project in which I try to show that theories of moral responsibility are grounded in normative moral theories. That project was sparked by thinking about a problem for consequentialism – whether consequentialists should think in terms of actual consequences (so called objective consequentialism) or expected consequences (subjective consequentialism). In this paper I try to motivate the larger project and show how it will settle the debate between objective and subjective consequentialism.

Objective and subjective consequentialists disagree about how rightness should be determined. On the one hand, objective consequentialists claim that rightness attaches the consequences that would actually be best. Of course, one must be a little careful with the counterfactuals in formulating this claim –the so-called ‘actual best consequences’ may not be actualised. Peter Railton expresses the view as follows: "Objective Consequentialism is the view that the criterion of the rightness of an act or course of action is whether it in fact would most promote the good of those acts available to the agent."¹ Proponents of the alternative view argue that there are two problems with this. First, we do not always know what the right act is on this criterion, and so the theory is not usefully action guiding.² Second, it might even be intuitively reckless and bad to do the right act – Frank Jackson’s well known example illustrates the point by setting up a situation where a doctor must make a choice between three drugs to prescribe to a patient with a non trivial but not life threatening skin complaint. The first drug, A, will alleviate the symptoms but will not cure the condition completely. The doctor knows that one of B and C will kill the patient, and the other will cure her, but she does not which is which. Clearly, the doctor ought not to prescribe B or C, even though prescribing one of them is surely the right action according to objective consequentialism. This leads to the absurd result that one should not do the right action. Jackson takes this absurdity as sufficient grounds to say that we should be probable consequences consequentialists, or decision theoretic consequentialists, as he puts it.

¹ Railton 1984.

² For examples of this sort of objection to objective consequentialism see Bergstrom 1996, Howard Snyder, Mason 2003, Wiland 2005, Miller 2004

These criticisms of objective consequentialism have not convinced its supporters, who tend to insist that the basic point of consequentialism requires that right action be defined in terms of actual consequences, even if we have to make a separation between criterion of right and decision procedures.³ The idea here is that even if our decision making procedures are imperfect, the primary notion is still objective consequentialism – objective consequentialism somehow provides a grounding for subjective consequentialism.

We might think that there is a sense in which the dispute is only verbal. Perhaps we can all agree that there is something important to be said about the action that actually would have the best outcome. We might want to call that action ‘right’, but we might think that it is clearer to call it the ‘best’ action. In any case, we should also all agree that what agents ought to do may not always be to pursue the best action – sometimes, because of lack of time or knowledge, they should use an alternative decision procedure, such as pursuing highest expected utility.⁴ So objective and subjective consequentialism can coexist quite happily – they are not competing for the same logical space. The only competition is over the use of the word ‘right’.⁵

This is all very well, but it has seemed to many consequentialists that there must be a deeper issue here, something concerning the very nature of moral theory. Somehow, objective consequentialism seems to violate an important stricture on the sorts of things that theories can say about obligations. According to Julia Driver, who is defending objective consequentialism, objective consequentialism is a form of ‘evaluational externalism’, in that “the moral quality of a person’s action or character is determined by factors external to agency, such as actual (rather than expected) consequences.” (Driver, p.68). There is something very suspicious about evaluational externalism, but theorists (including myself) have struggled to articulate this thought. Two avenues have been pursued. The first, which goes back to Pritchard, puts the

³ See for example Feldman (2006). Others who count themselves as objective consequentialists include G E Moore, Peter Railton, Julia Driver, The original suggestion about a separation of criteria of right from decision making procedures goes back at least as far as Sidgwick and Mill. The contemporary way of putting it is due to Eugene Bales (1971).

⁴ There are also disputes about whether expected utility is any more practical a notion than actual utility (see e.g. Lenman (2000) and Feldman (2006)). However, expected utility is only problematic if it is understood as itself an objective notion. Some philosophers have understood it along these lines (Oddie and Menzies (1992)). I will not enter into that debate here, except to say that my preferred understanding of expected utility is as an entirely subjective notion.

⁵ For example, this is more or less the view espoused by Ingmar Persson in a recent paper. (Persson, 2008)

issue in terms of knowledge and obligation, suggesting that obligations cannot be unknowable.⁶ The second, closely related but more recent, frames the issue in terms of the ‘ought implies can’ principle.⁷ In both cases the thought is the same: that obligations must be, in some sense, *accessible* to the agent. Or, to put it in terms closer to Driver’s formulation, evaluational externalism is not plausible because moral qualities must be determined by factors *internal* to agency.

The problem with these approaches is that they hover at the edge of another vastly complicated philosophical question – the nature of agency and moral responsibility – and so they often beg more questions than they answer. Why do we think that ought implies can? What is the relevant sense of ‘can’ in the ought implies can principle?⁸ Why would knowledge matter for obligation? What is behind the rather vague thought that obligations have to be accessible to the agent? Why is agency important? In order to flesh out the suspicion that evaluational externalism is a non-starter, we have to investigate the relevance of theories of responsibility to moral theories.

Moral Responsibility

In this paper I unapologetically set aside the possibility that we have free will in any metaphysical sense, and concentrate on naturalist accounts of responsibility. The problem for naturalist accounts of course, is how to be compatibilist – how to show that we have responsibility without free will. My aim is to provide a sketch of a solution to the compatibilist’s problem. The solution I suggest does not contradict the existing literature on moral responsibility. In fact, I think that the right account of moral responsibility is probably already out there. Rather my aim is to give a meta account of responsibility - an account of what would make one theory of responsibility rather than another the correct one.⁹ If there are no metaphysical

⁶ Pritchard...see also Frankena and Sorenson’s arguments against Pritchard.

⁷ See Bergstrom 1996, Howard Snyder, Mason, McConnell.

⁸ See Frances Howard-Snyder, ...and my ‘Moral Responsibility and Moral Theory’. Some philosophers reject the ought implies can principle on the grounds that there are many uses of ought in ordinary language that clearly do not imply can. (E.g. I ought to be taking my daughter to a party but I am stuck in traffic). See for example, Sinnot Armstrong and Michael Stocker. Such arguments are not relevant to the argument here.

⁹ Obviously the best candidates will be internally coherent and properly naturalist – i.e. will not contain any question begging terms and so on. In this paper I suggest another criterion – an overarching rationale for the theory. On top of that there remains another meta project –

standards determining responsibility, we need to look somewhere else. In this paper I make a suggestion about where we should look.

My view is in some ways close to Jay Wallace's. Wallace has argued that our interest in responsibility is primarily a moral interest, and so we can find out about responsibility by looking at moral theory. In brief, Wallace argues that that we ought to hold people responsible when it is morally appropriate to do so; it is morally appropriate when it is fair to do so, and it is fair to do so when the person has behaved autonomously. Wallace's view is very complex, but the basic point, that we should hold people responsible when it is morally appropriate to do so, is clearly unsatisfactory. When we ask about whether we should praise or blame someone we are not asking whether we *morally* ought to engage in the act of praising or blaming him or her. Of course, some consequentialists have made the mistake of thinking that that is the only question,¹⁰ but as has been pointed out frequently, and is pointed out by Wallace himself, a consequentialist account of responsibility may succeed in being naturalist, but it misses the point of responsibility entirely.¹¹

Ideally, a naturalist account of responsibility would be able to preserve the commonsense thought that the 'should' in 'should we hold P responsible?' is a theoretical should rather than a practical one. I argue for a view according to which we should hold P responsible when it is *correct* to do so, not when it is expedient, or has the best consequences, or is morally right on some other normative view. I will show that we can give an account of facts about responsibility without appealing to a realm of independent facts about responsibility. On my view, facts about responsibility are facts about morality, more specifically, about the applicability of moral norms. There remains the meta ethical project of validating facts about morality. My point is that facts about responsibility and facts about morality stand or fall together – there is just one hard question here, not two.

Critique of Wallace

the meta ethical project – giving an account of how the standards of correctness are anything other than relative.

¹⁰ I shall argue that Wallace makes this mistake too. In what follows I ignore the possibility that the consequences of acts of praise and blame should be evaluated independently of attributions of responsibility.

¹¹ Philosophers who give a consequentialist account of moral responsibility include JJC Smart (1961), Daniel Dennett (1984).

Wallace's argument proceeds by focussing on the incompatibilist's suggestion that it would be unfair to hold the agent responsible in the absence of free will. He says:

Philosophers have often suggested that freedom matters to responsibility because it would be unfair to hold people responsible in the absence of freedom of will. I will take this thought much more seriously than it has been taken heretofore, and work it up into an interpretation of what is at stake in the debate between compatibilists and incompatibilists. This normative interpretation has two important consequences. First, we cannot establish what it is to be a morally responsible agent unless we first understand the stance of holding someone responsible—the stance of the moral judge, rather than of the agent who is judged. Second, determining what the conditions of moral responsibility are will require an excursion into normative moral theory; we will need to investigate our principles of fairness, to see what they entail about the conditions under which it would be fair to hold people responsible.¹²

Wallace is doing various things in his book, but I will focus on his attempt to make sense of the standards of correctness for holding people responsible. As Wallace points out, 'should we hold P responsible?' is a normative question, and we need to know which norms are in play. Wallace argues that we cannot appeal to theoretical norms, as that would amount to assuming that there is a fact of the matter about responsibility independently of our stance of holding people responsible.¹³ Wallace is unwilling to grant that it is even coherent to talk about a prior and independent realm of facts about responsibility. So, Wallace argues, if the 'should' in 'should we hold P responsible' is not a theoretical should, it must be a practical should. Thus, the argument goes, to determine which substantive norm is relevant here we should look at practical norms – *norms that subserve our practical commitments*, as Wallace puts it.¹⁴ Wallace understands the question, 'should we hold P responsible?' as a moral question, and in particular, the moral question, '*is it fair to hold P responsible?*'.

¹² P.5.

¹³ Hard-nosed incompatibilists would of course insist that our attributions of responsibility must be judged with reference to theoretical norms. Wallace has various responses to the incompatibilist on this point, but I will not canvass them here.

¹⁴ P.92. Note that Wallace classes moral norms as practical norms, which is of course, of a piece with his Strawsonian motivational internalism about attributions of responsibility.

Wallace defends this interpretation of the question by pointing out that the incompatibilist complaint is often that it would be unfair to hold people responsible if determinism were true. So, we should hold someone responsible only when it would be fair for us to do so. Hence Wallace proceeds by examining our concept of fairness in order to get a grip on the standards of correctness for attributions of responsibility.

This is a very neat argument, but it is flawed. It is flawed in two ways. First, Wallace is wrong to conclude that the only alternative to thinking that there is a realm of independent facts about responsibility is to say that the ‘should’ in ‘should we hold P responsible’ is a practical should. I shall leave that aside for the moment and come back to my own account in the next section. Second, even if we grant for the moment that the should is a practical should, Wallace does not have the right sort of grounds to conclude that the practical norm in question must be the norm of fairness.

Wallace sees that there are other possible candidates for the practical norm, For example, we might argue (as some have) that we should hold people responsible when doing so would have the best consequences. But, as Wallace points out, a consequentialist approach to responsibility seems very implausible:

“Of course it would be possible to bring other practical norms to bear in assessing the stance of holding people responsible. Such norms might even support a verdict at odds with the conclusion that is grounded in moral considerations of fairness. For instance, it might turn out that there are good strategic reasons to hold responsible a class of people whom it would clearly be unfair to treat in that way – perhaps such treatment would materially advance our interests in security, or happiness, or the eventual attainment of a more egalitarian social world.”¹⁵

The lesson Wallace draws from this is that “norms of fairness have a privileged position in determining what it is to be a responsible agent”¹⁶, but that is not the right lesson to draw. Wallace’s conclusion, that ‘we should hold P responsible’ amounts to

¹⁵ p.94. As Wallace goes on to say, we may act in accordance with these pragmatic norms, but we wouldn’t think that the agents we treat as responsible really are responsible. Wallace provides a more detailed rebuttal of consequentialist accounts on pp. 55-60.

¹⁶ P.94

saying that it is fair to hold P responsible sounds convincing because we are inclined to a very thin reading of fair. Wallace gets a lot of mileage out of his claim that the incompatibilist contends that it would be unfair to hold people morally responsible in the absence of free will. That is what justifies Wallace's claim that his account of moral responsibility is perfectly general – we should all accept it, no matter what our prior views. However, when the incompatibilist says that it would be unfair to hold P responsible in the absence of free will, she is using the word 'fair' in a thin sense.

Imagine a conversation between two incompatibilists:

X: It would be unfair to hold people responsible if they lack freedom of the will.

Y: What do you mean unfair? I agree of course that it would be incorrect, but why do you say it would be unfair?

X: Just because you shouldn't do it – you shouldn't hold people responsible when they are not.

Y: I see, and I agree – but that doesn't mean that it would be *unfair* to hold people responsible in the absence of free will. On my account of fairness, fairness is a matter of participating in cooperative practices from which we benefit. Holding someone responsible erroneously has nothing to do with that.¹⁷

X: Good point. I forgot that there are so many different substantive account of fairness. All I meant was that you shouldn't hold people responsible when they are not.

Y: Then we do not disagree.

The point is that there is an uncontroversial claim that it would be unfair to hold an agent responsible in the absence of free will, but it does not amount to much. It is just the claim that you *should not* hold an agent responsible if she is not. We are not going to learn anything about *when* an agent is not responsible by examining the claim that we should not hold her responsible if she is not responsible. The claim that it would be unfair to hold an agent responsible in the absence of free will can be given a thicker interpretation, according to which by fairness we mean something connected to 'reasonableness or desert.'¹⁸ But if we take the thick reading of fair (as Wallace

¹⁷ Wallace himself points out that if fairness is participation in cooperative practices, then it doesn't seem particularly relevant to responsibility. (p. 106).

¹⁸ P.109.

needs to in order to learn anything about what responsibility is), then it is no longer obvious that the claim would be generally accepted.

My criticisms of Wallace are not intended to show that his account of moral responsibility is implausible. In fact, I think it is a very plausible account of moral responsibility, and one that has a long tradition, from Kant to Fischer and Ravizza. My criticisms are intended to show that Wallace's account is not a *general* account of moral responsibility. Wallace uses a very particular account of what morality is, and so inevitably reaches a certain account of moral responsibility.

A general account of the correlation between responsibility and moral theory

I agree with Wallace that our interest in responsibility is closely related to our interest in morality – we are interested in whether people are responsible because we want to know whether we should praise or blame them. However, it is not that we want to know whether we *morally* ought to engage in the act of praising or blaming them, we want to know when it would be *correct* to praise or blame someone.

Our problem is that, in the absence of free will, there seems to be nothing theoretical available to anchor our account of moral responsibility. We might pick any section of the causal chain that leads to an action, and claim that that section is the part that makes an act responsible. But of course that would be arbitrary. Wallace attempts to show that norms of fairness anchor our attributions of responsibility. I argued in the previous section that Wallace's argument begs the question – it requires that we *already accept* Wallace's account of what makes an attribution of responsibility unfair, in which case we have already accepted an account of responsibility.

Moral responsibility seems to be one of the conditions of praise and blame being appropriate. So another way to get at the question of what it is to be responsible is to ask, what conditions must hold for an act to be apt for normative appraisal? There are different sorts of appraisal – aesthetic, moral, and so on. The conditions of aptness for these different sorts of appraisal include conditions of content as well as conditions that correspond to our notion of responsibility. For example, only certain sorts of content make an act apt for moral appraisal as opposed to aesthetic appraisal, and one of the differences between different conceptions of morality is the differing accounts of the content that they give. Most obviously – some moral theories take purely self-regarding acts to be moral, while others do not. We should only praise or

blame P for her act if it is the right sort of act in terms of content, and we should only praise or blame P if she is responsible for her act. I will give an account of how we ought to think about responsibility that is analogous to what I take to be a quite uncontroversial view about how we ought to think about these questions of content.

Just as Wallace asks, ‘when should we hold P responsible?’, we ask, ‘what sort of content must an act have for it to count as a moral act?’ As with the original question about responsibility, this is a normative question, and so we need to know which norms are in play. Wallace’s argument was that that we cannot appeal to theoretical norms, as that would amount to assuming that there is an independent fact of the matter. We certainly should not assume that there is an independent realm of facts making it true that certain content renders an act moral and certain content renders an act aesthetic. So what should we say? Wallace argues that the ‘should’ in ‘should we hold P responsible’ must be a practical should. But though that may seem plausible in the case of responsibility (though I shall argue that it is not), it is surely the wrong reading of the question about content. It seems ridiculous to say that we should use moral or other practical grounds to determine whether an act has moral content. It is fairly obvious that we should see the question as a conceptual question – it is not a metaphysical question and it is not a normative question. Whether or not an act’s content is moral depends on particular moral theories. So, if you are a Kantian, you will think that merely self-regarding actions are apt for moral appraisal, but if you are of some other stripe, you may think that only other regarding actions are apt for moral appraisal. A moral theory seems plausible when it says plausible (and internally coherent) things about which actions are right and wrong. A theory that tells us that it is wrong to step on cracked paving stones is implausible – that doesn’t strike us as the right sort of thing to be wrong. The relationship between your moral theory and the view you take on what counts as moral content is not one with a clear direction of priority – in fact, the question about what sorts of content are apt for moral appraisal is almost the same question as the question of what sort of moral theory you have at all – taking a particular view about what counts as moral content isn’t simply dictated by the moral theory you have, it is a large part of having that moral theory.

I suggest that we should think of moral responsibility in a similar way. Having a normative theory involves thinking that certain sorts of content are apt for moral appraisal. It also involves thinking that certain sorts of human behaviour – of agency –

are apt for moral appraisal, and my claim is that this is all there is to having a view about moral responsibility. The existing compatibilist literature on moral responsibility argues over which part of the causal chain leading up to action is important – which matters – which should count as responsibility. I am not trying to arbitrate between different substantive views about what matters, rather, the question I am trying to answer here is a question about *the way in which* some part of the causal chain matters. It is not that there really is free will, so what is it? On my view, we should understand the literature on moral responsibility as an argument over which part of the causal chain matters *morally*.¹⁹ Different ethical outlooks place importance on different aspects of agency, just as they also differ in their account of what contents of actions are morally important. As with content, the direction of priority is not clear – each informs the other.

We should thus see a particular ethical theory as being a package, in which a view about moral responsibility is inseparable from the view about morality. So for each broad theory, the appropriateness of praise and blame, rightness and wrongness, and moral responsibility are interdependent. An agent is morally responsible, and the attitudes of praise and blame are appropriate exactly when the theory bestows a verdict of rightness of wrongness – *because that is all there is to responsibility*.²⁰ This account of moral responsibility may not match all our pre philosophical intuitions about responsibility.²¹ In fact, it is bound not to. The reason that compatibilists are so wedded to our existing concept of moral responsibility is that they don't have anything else to go on – they don't have an overarching rationale for an account of moral responsibility. That is what I aim to provide in this paper.²² The overarching rationale for a particular view about moral responsibility is just the same as it is for the rest of the corresponding moral theory.

So what are the different packages? The clearest sort of view is the Kantian view articulated by Wallace, according to which reasoning is crucial. This family of views is united by the claim that there is some sort of intellectual process that the

¹⁹ This is not something that existing writers on moral responsibility should be hostile to – indeed some (notably Jay Wallace of course) already realise that their view is not independent of a moral outlook. Elsewhere I argue this point in more detail.

²⁰ As usual I am ignoring uses of praise and blame which evaluate them as actions distinct from the cognitive attitude they embody.

²¹ I agree with Manuel Vargas that compatibilist theories may have to be revisionist.

²² I argue this in more detail in my 'Moral Theory and Moral Responsibility'.

agent must have gone through in acting in order to count as responsible.²³ It is not hard to see this as part of a package according to which actions are right or wrong when the agent has gone through some sort of intellectual process, and correspondingly, agents merit praise or blame only when they have gone through the relevant intellectual process. I entirely agree with Julia Driver when she says that subjective consequentialism is a view like Kantianism – an extremely evaluationally internalist view, and is a quite different sort of view to objective consequentialism.

Objective consequentialism is at the other end of the scale, in that it places all the weight on something that is entirely outside of the agent. Rightness, and thus the appropriateness of praise and blame and moral responsibility, all depend on outcomes. I will return to this shortly.

Driver claims that Aristotle's view is a mixed view, as Aristotle thinks that rightness is partly determined by the goodness of outcomes, and partly by the goodness of intentions. This is of course true of the historical Aristotle. However, when we look at the literature in moral responsibility it is possible to identify a strand of thinking about moral responsibility that corresponds to an ethical outlook that has a lot in common with Aristotle's. The traditional landscape in theories of moral responsibility divides up into two sorts of view. On the one hand are what I call 'intellectual process views', and on the other hand are 'real self views' – that locate moral responsibility in the agent's real self. Real self views try to show that we hold (or at least, should hold) an agent morally responsible when her action issues from her 'real self'.²⁴ The hard question for this sort of view is 'what is the real self?' – what makes an action one that really is the agent's?²⁵

²³ These are often referred to as 'historicist' views, because they use the history of the act to determine whether the agent is morally responsible. Fischer and Ravizza classify them as 'externalist views' (Fischer and Ravizza, 1998). Elsewhere Fischer refers to these as 'reasons responsiveness accounts' ('Recent Work' p. 127). Susan Wolf, Hilary Bok and John Fischer and Mark Ravizza, and of course Jay Wallace have argued for intellectual process views about responsibility

²⁴ These views are also known as 'structuralist views' (see Fischer and Ravizza, 1998), 'internalist' views (see Fischer and Ravizza, 1998)), and 'Self-disclosure views' (see Gary Watson, 2004).

²⁵ For a more detailed overview of recent work in moral responsibility see my XXX. In early work, Frankfurt argues that we are responsible for our acts when we act according to our second order volitions. In later work he argues that an agent is morally responsible when her action is 'wholehearted', and changes his mind again in 'The Faintest Passion' where he suggests that the crucial notion is that an agent be *satisfied* with her higher order desire.²⁵

A survey of the literature on moral responsibility suggests that there are two packages. The Kantian package, which we are very familiar with, according to which rightness and moral responsibility depend on reasoning, and the broadly Aristotelian picture, according to which rightness and moral responsibility depend on deep features of the agent's character. Subjective consequentialism fits neatly into the Kantian category – according to subjective consequentialists what matters is that the agent has certain intentions. Objective consequentialism does not fit into either of these packages, and so we should add another package - the objective consequentialist outlook, according to which moral responsibility and rightness depend on outcomes. The first two are evaluationally internalist in Driver's terms; the last one is evaluationally externalist.

Consequentialism

I am now in a position to explain why objective consequentialism seems so unsatisfactory. It is because the package is unsatisfactory. The account of rightness can seem plausible when considered in isolation, but once we look at the whole package, and in particular, the corresponding account of moral responsibility, we see that the theory is crazy, just as any evaluatively external theory will be. Such theories do not provide the right sort of connection between the agent and the things he is responsible for - a merely causal connection is not enough to capture the 'deep sense of responsibility' that even compatibilists are entitled to look for. It is no surprise that in the literature on moral responsibility we find theories that correspond very neatly to Kantian and to broadly Aristotelian theories of morality, but we do not find objective consequentialist theories of responsibility.²⁶

Frankfurt's progression here represents his struggle with the issue of what is involved in an agent identifying with her motivation – or, to put it another way, which actions really do come from the agent's real self? This problem is at the heart of real self views, and the theories of Gary Watson and Michael Bratman attempt to develop a satisfactory account of the real self.

²⁶ Note that the so-called consequentialist account of moral responsibility is not consequentialist in the sense I am talking about here. Here I am comparing an account of responsibility that corresponds to the moral theory of consequentialism the way that reasons responsiveness accounts correspond to intention-focussed accounts of morality and real self views correspond to character-focussed accounts of morality. The corresponding consequentialist account would say that we ought to hold people responsible for all the causal consequences of their actions. This is not the same as the proposal that is commonly known as the consequentialist account of responsibility – that we should hold people responsible when doing so has good consequences. Neither account has much to recommend it.

Of course, the objective consequentialist will reply that she is not trying to give an account of moral responsibility. My response is that she should have been, or perhaps even that she must have been without realizing it. The thought that subjective consequentialists have been trying to express for years is that objective consequentialism is defective because the view about moral responsibility that it embodies is absurd. What I have done in this paper is vindicated that thought. I have argued that the best compatibilist account of moral responsibility makes it inseparable from moral theory. This leaves most of our theorizing and arguments untouched – we still have to decide between a Kantian or broadly Aristotelian approach, and within each we have to settle large disagreements. The progress I have made is to establish that in deciding between different views of ethics and different views of responsibility we should look at the whole package. Part of what makes a view about responsibility plausible is the corresponding view about morality, and vice versa. Once we see this, we see that objective consequentialism is a bad package.

Does Subjective Consequentialism need Objective Consequentialism?

As I have said, objective consequentialists often point out that there is a distinction between criteria of rightness and decision making procedures. So, they might argue, objective consequentialism does not purport to be a complete package – the combination of the decision procedure and the criterion of right is what makes the complete package. In fact, the objector might go on, subjective consequentialism is lost without objective consequentialism – objective consequentialism is the spirit of the theory, without it subjective consequentialism makes no sense.

Before answering this objection I should first say a little about subjective consequentialism. On my view, subjective consequentialism belongs in the same family as Kantian moral theories, because it shares a basic view according to which (roughly) what matters morally are the agent's intentions – it matters that the agent aims for the right. Subjective consequentialism agrees with Kant's dictum that the only thing that is good without qualification is the good will. However, subjective consequentialism differs from Kant's own ethical theory, in that it has a quite different account of what the agent should aim for. Kant thinks the agent should aim to act in such a way that everyone else can act that way too without contradiction. Subjective consequentialists think that the agent should aim to bring about the best possible outcome. Neither of these quick descriptions is very illuminating as they

stand, and of course both need and have received numerous elucidations. I will leave aside the many other complications in the formulation of subjective consequentialism - the question here is whether the objective consequentialist standard of rightness is somehow prior to the subjective standard.²⁷

Why might we think objective consequentialism is prior? The objectivist is probably thinking of cases where what one subjectively ought to do is not in fact what would be best. For example, imagine that Frank Jackson's doctor decides to prescribe A, the drug that will ameliorate the symptoms without curing the patient. However, just as she is writing the prescription, her colleague comes in with the results of a new study. A simple blood test will show whether B or C is the drug that will cure the patient completely. The colleague should obviously tell the doctor about the test, and the doctor should alter her prescription accordingly. But, according to subjective consequentialism, the doctor was doing the right thing in prescribing A. If she was doing the right thing, why should she change her plan? It seems to the objectivist that there is something missing from subjective consequentialism. The objectivist argues that the subjective consequentialist must be guided by what is objectively right when that is possible, and so objective rightness is the prior notion.

The subjective consequentialist answers this by saying that what should guide her, what makes her decisions right, is information about what is *best*, not what is right. So, we do all agree that there is something important about the option that actually would be best – the subjectivist agrees that if we know about it, it becomes the right option. But that doesn't mean that it must have been the right option before. What is prior to the subjective consequentialist account of rightness is the consequentialist account of bestness, not the objective consequentialist account of rightness. This may sound like a merely verbal dispute. But there is more to it than that. I have argued that the concept of rightness has conditions of moral responsibility built into it, and so objective consequentialism is clearly misusing the concept of rightness.

²⁷ The answer I give below is roughly the same answer that I gave in my 2003. See also Howard Snyder 2005.

Bibliography

- Bales, Eugene** (1971) 'Act-Utilitarianism: Account of Right-making characteristics or Decision-making procedure?', *American Philosophical Quarterly* 8 (1971).
- Bergstrom, Lars** (1996) 'Reflections on Consequentialism', *Theoria*, Vol. 62, Part 1-2, 1996, pp.74-94.
- Bok, Hilary**, (1998) *Freedom and Responsibility* Princeton University Press.
- Bratman, Michael** (1999) 'Identification, Decision and Treating as a Reason', reprinted in *Faces of Intention* (Cambridge: Cambridge University Press, 1999).
- Dennet, Daniel C.** (1984), *Elbow Room: The Varieties of Free Will Worth Wanting* (Cambridge, Mass.: MIT Press, 1984).
- Driver, Julia** (2001) *Uneasy Virtue* CUP
- Fischer, John Martin and Ravizza**, (1998), *Mark Responsibility and Control: An Essay on Moral Responsibility*, (Cambridge: Cambridge University Press)
- Frankena, William** (1950) 'Obligation and Ability', in Max Black, (ed.) *Philosophical Analysis: A Collection of Essays* (London: Prentice-Hall, 1950).
- Gruzalski, Bart** (1981) 'Foreseeable Consequence Utilitarianism', *Australasian Journal of Philosophy* 59 (1981), 163-176.
- Howard-Snyder, Frances** (1997) 'The Rejection of Objective Consequentialism', *Utilitas* 9 (1997), 241-248.
- Howard Snyder, Frances** (2005) 'It's the Thought that Counts' *Utilitas* 17 (2005), 265-281.
- Jackson, Frank** (1991) 'Decision-theoretic Consequentialism and the Nearest and Dearest Objection', *Ethics* 101 (1991), 461-482.
- Lenman, James** (2000). 'Consequentialism and Cluelessness.' *Philosophy and Public Affairs*, 29: 342-370.
- Mason, Elinor** (2004) 'Consequentialism and the Principle of Indifference', *Utilitas*. N 04. 16(3), 316-321.
- Mason, Elinor** (2003) 'Consequentialism and the "Ought Implies Can" Principle'. *American Philosophical Quarterly*. O 03. 40(4), 319-331.
- McConnell, Terrance** (1989) "'Ought' Implies 'Can' and the Scope of Moral Requirements", *Philosophia* 19 (1989), 437-454.

- Miller, Dale E., (2004)** ‘Actual–Consequence Act Utilitarianism and the Best Possible Humans’, *Ratio* 16, (2004) pp.49-62.
- Oddie, Graham & Menzies, Peter** (1992) ‘An Objectivist’s Guide to Subjective Value’, *Ethics* 102 (1992), 512-533.
- Persson, Ingmar,** (2008) ‘A Consequentialist Distinction between What We Ought to Do and Ought to Try’, *Utilitas* 2008
- Prichard, H.A.** *Moral Obligation* (Oxford: Clarendon Press, 1949).
- Railton, Peter** (1984) ‘Alienation, Consequentialism, and the Demands of Morality,’. **Philosophy and Public Affairs** 13 (1984).
- Sinnott-Armstrong, Walter** (1984) ‘‘Ought’ Conversationally Implies ‘Can’’, *The Philosophical Review* 93 (1984), 249-261.
- Sorenson, Roy,** (1995) ‘Unknowable Obligations’ *Utilitas* 7 (1995),
- Smart, J.J.C.,** (1961), ‘Freewill, Praise and Blame’ *Mind* 70, pp. 291-306.
- Stocker, Michael** [1] ‘‘Ought’ and ‘Can’’, *Australasian Journal of Philosophy* 49 (1971), 303-316.
- Vargas, Manuel,** (2004) ‘Responsibility and the Aims of Theory: Strawson and Revisionism’, *Pacific Philosophical Quarterly* 85 (2004), pp.218-241.
- Vargas, Manual,** (2005) ‘The Revisionists Guide to Responsibility’, *Philosophical Studies* 125 (2005), 399-429.
- Wallace, R.J.** (1996) *Responsibility and the Moral Sentiments*, Harvard University Press.
- Wiland, Eric** (2005) ‘Monkeys, Typewriters, and Objective Consequentialism’, *Ratio* S 2005. 18(3), 352-360.
- Wolf, Susan,** (1990), *Freedom within Reason* (New York: Oxford University Press, 1990).
- Zimmerman, Michael** (1997) ‘Moral Responsibility and Ignorance’ *Ethics* 107 (1997), pp. 410-426