

# Utilitarianism, Altruism, and Consent

Christopher J. G. Meacham

*Forthcoming in the Journal of Ethics and Social Philosophy*

## Abstract

A number of criticisms of Utilitarianism – such as “nearest and dearest” objections, “demandingness” objections, and “altruistic” objections – arise because Utilitarianism doesn’t permit partially or wholly disregarding the utility of certain subjects. A number of authors, including Sider (1993), Portmore (2008), and Vessel (2010), have responded to these objections by suggesting we adopt “dual-maximizing” theories which provide a way to incorporate disregarding. And in response to “altruistic” objections in particular – objections noting that it seems permissible to make utility-decreasing sacrifices – these authors have suggested adopting a dual-maximizing theory that permits disregarding one’s own utility.

In this paper I’ll defend two claims. First, I’ll argue that dual-maximizing theories are a poor way to incorporate disregarding. Instead, I’ll suggest that “variable-disregarding” theories provide a more attractive way to incorporate disregarding. Second, I’ll argue that the right way to handle these “altruistic” objections isn’t to permit disregarding one’s own utility, it’s to permit disregarding the utility of those who consent. Together, these two claims entail that the best way to modify Utilitarianism to handle “altruistic” objections is to adopt a variable-disregarding view that disregards the utility of those who consent.

## 1 Introduction

Consider the following case:

**Parental Sacrifice:** Your child’s birthday is coming up. You have the option of working overtime in order to buy your child a gift. Doing so would significantly decrease your utility, a decrease greater than the utility your child would receive from the gift. No one else’s utility would be affected.

According to *objective act Utilitarianism*, an act is morally permissible *iff* it maximizes total utility. Since working overtime would not maximize total utility, Utilitarianism entails that it’s impermissible to work overtime to buy your child the gift. But this

strikes many as the wrong verdict. For while it's not obligatory to work overtime, it's surely permissible to do so.<sup>1</sup>

The standard diagnosis of this problem is that Utilitarianism provides implausible prescriptions in cases like Parental Sacrifice because it doesn't accommodate an important "self-other asymmetry". In particular, it doesn't accommodate the fact that it's permissible to disregard one's own utility when making decisions about how to act.

The standard remedy is to modify Utilitarianism by adopting an altruistic version of what I'll call a "dual-maximizing" theory. Dual-maximizing theories provide a particular way of allowing agents to (partially or wholly) disregard the utility of some group of subjects. And such views have been employed to address a number of objections to Utilitarianism. In response to "nearest and dearest" objections to Utilitarianism, some have suggested adopting dual-maximizing views that permit partially disregarding the utility of those one isn't close to, and thus privileging the utility of one's friends and family. In response to "demandingness" objections to Utilitarianism, some have suggested adopting dual-maximizing views that permit partially disregarding the utility of others, and thus privileging one's own utility. And in response to "altruistic" objections to Utilitarianism like Parental Sacrifice, some have suggested adopting dual-maximizing views that permit disregarding one's own utility, and thus privileging the utility of others.<sup>2</sup>

I'll argue that both the standard diagnosis and the standard remedy are mistaken. I'll argue that the standard diagnosis is mistaken because the cases like Parental Sacrifice do not demonstrate that it's permissible to disregard one's own utility. Rather, what such cases show us is that it's permissible to disregard the utility of those who *consent*.

And I'll argue the standard remedy is mistaken because the way in which dual-maximizing theories permit disregarding the utility of groups of subjects is problematic. I'll argue that if we want to modify utilitarianism to permit disregarding, we should do so in a different way, employing what I'll call a "variable-disregarding" theory. And this is true not just in the context of "altruistic" objections like Parental Sacrifice, but in every case in which one might want to permit disregarding.

Note that these two theses are independent of one another. One can accept that it's permissible to disregard the utility of those who consent, and deny that the best way to permit disregarding is via a variable-disregarding theory. And one can accept that the best way to permit disregarding is via a variable-disregarding theory, and

---

<sup>1</sup>I've made your child the beneficiary of your sacrifice in order to make the case easy to imagine. But there's nothing important about the beneficiary being your child. If you wanted to make such a sacrifice to benefit a stranger, most would still feel it's permissible to do so.

<sup>2</sup>See Smart and Williams (1973) for a classic presentation of the "nearest and dearest" objection, see Scheffler (1994) for a prominent discussion of "demandingness" objections, and see Slote (1984) for an influential discussion of "altruistic" objections. See Sider (1993), Portmore (2008) and Vessel (2010) for suggestions on how to use dual-maximizing views (or multi-maximizing views; cf. footnote 8) to address these objections.

deny that it's permissible to disregard the utility of those who consent. However, these two theses are independently attractive, and together they entail that cases like Parental Sacrifice are best accommodated by adopting a variable-disregarding view that permits disregarding those who consent.

Two preliminary comments before we proceed. First, the issues I discuss in this paper are relevant to accounts of both *subjective obligation* (where an agent's beliefs mediate what they ought to do) and accounts of *objective obligation* (where they don't).<sup>3</sup> But for simplicity, I focus my attention here on theories of objective obligation.

Second, the issues discussed in this paper aren't just relevant to Utilitarianism, but to many theories that take utility into account.<sup>4</sup> For example, consider a hybrid rights theory that takes rights to be side constraints, but otherwise tells you to maximize utility. In cases in which no one's rights are being threatened, this theory will yield the same prescriptions as Utilitarianism. Thus one can raise "nearest and dearest", "demandingness", and "altruistic" worries for this kind of rights theory too. And it's natural to deal with these worries by allowing agents to disregard the utility of certain groups of subjects. In a similar vein, a typical person-affecting view will yield the same verdicts as Utilitarianism in cases in which all of the same individuals exist regardless of what one does. Thus one can raise "nearest and dearest", "demandingness", and "altruistic" worries for these person-affecting views. And again it's natural to address these worries by allowing agents to disregard the utility of certain groups.

That said, it's easiest to see our way through these issues when our underlying theory is as simple and straightforward as possible. So in what follows I'll focus my attention on how one might apply these kinds of modifications to Utilitarianism.

The rest of this paper will go as follows. In the second section I present dual-maximizing theories, show how one can use them to handle cases like Parental Sacrifice, and then raise some worries for such theories. In the third section I present variable-disregarding theories, show one can use them to handle cases like Parental Sacrifice, and then raise some worries for the standard diagnosis of these cases in terms of some kind of self-other asymmetry. In the fourth section I make the case for thinking that the morally important issue that cases like Parental Sacrifice raise is not the distinction between self and other, but the distinction between giving and withholding consent. I then present a variable-disregarding theory that accommodates the role of consent, show how it handles all the cases discussed so far, and discuss the conditions required for morally relevant consent. In the fifth section I consider several objections. In the sixth section I briefly summarize these results.

---

<sup>3</sup>For some recent discussions of these issues, see Zimmerman (2006), Sobel (2009) and Graham (2010).

<sup>4</sup>For some recent examples of other views that one could modify in this way, see Roberts (1998), Preda (2011) and Meacham (2012).

## 2 Dual-Maximizing Theories

### 2.1 The Structure of Dual-Maximizing Theories

Utilitarianism says that we should treat all subjects equally. But various objections to Utilitarianism, such as the “nearest and dearest” objections, “demandingness” objections, and “altruistic” objections, might make one reconsider this tenet. For while it’s presumably permissible to treat everyone equally, one might also want to permit partially or wholly disregarding the utility of some subjects, whether it’s other people, those one doesn’t feel close to, or oneself.

If we want to modify Utilitarianism in order to permit (but not require) disregarding some subjects, one popular option is to adopt a dual-maximizing theory.<sup>5</sup> We can formulate dual-maximizing theories as follows, with respect to some condition  $\phi$  that picks out the disregarded subjects, and some weight  $w \in [0, 1]$  that corresponds to the degree to which disregarded subjects are taken into account (with  $w = 0$  if we give the utility of disregarded subjects no weight, and  $w = 1$  if we give the utility of disregarded subjects full weight):

**Dual-Maximizing Theories:** An act  $a$  is permissible *iff* either (i)  $a$  maximizes utility, or (ii)  $a$  maximizes the sum of (the utility of those who don’t satisfy  $\phi$ ) plus ( $w$  times the utility of those who do satisfy  $\phi$ ).

We’ll spend most of our time focusing on cases in which  $w = 0$ , and the utility of disregarded subjects is completely discounted. In these cases, dual-maximizing theories consider two things: whether an act maximizes utility overall, and whether an act maximizes utility for the subjects we’re not disregarding. If the answer to either question is “yes”, then the act is permissible. If the answer to both questions is “no”, then the act is impermissible.

As Portmore (2008) notes, dual-maximizing theories have several attractive features.<sup>6</sup> As we saw above, we can use dual-maximizing theories to accommodate self-centered options, special consideration for one’s friends and family, and self-sacrificing options. And we can employ dual-maximizing theories to make sense of a broad range of supererogatory acts. E.g., given a self-centered dual-maximizing theory, there will often be a number of permissible acts which are worse for the agent, and better for others, than other permissible options.<sup>7</sup>

---

<sup>5</sup>If one wants to modify Utilitarianism to *require* agents to disregard a group of subjects, then this is easy to do: one can either change one’s characterization of utility to simply exclude the group in question, or (equivalently) adopt a theory which tells you to ignore the utility of such subjects when evaluating acts. Since the question of how to require disregarding is not contentious (unlike the question of how to *permit* disregarding), I won’t address it in the text.

<sup>6</sup>Strictly speaking, Portmore (2008) makes these claims regarding a broader class of theories that he calls “dual-ranking theories”, which consist of any theory with a pair of conditions (i) and (ii) such that an act is permissible *iff* it satisfies one of these conditions. But all of the particular theories Portmore considers are either dual-maximizing theories or multi-maximizing theories (see footnote 8).

<sup>7</sup>Of course, how to understand supererogation is controversial, and one might want to cash this no-

Dual-maximizing theories are a special case of a more general kind of theory we might call “multi-maximizing theories”, which allow for any number of different disregarded groups with different weights. But to keep things simple, I’ll bracket these complications and focus on dual-maximizing theories in what follows.<sup>8</sup>

## 2.2 Self-Other Utilitarianism

Cases like Parental Sacrifice invoke the intuition that there’s a distinction between how one takes one’s own utility into consideration and how one takes the utility of others into consideration. In particular, the intuition is that it’s permissible to disregard one’s own utility when making decisions, but not permissible to disregard the utility of others. One way to modify Utilitarianism in order to accommodate this intuition is to adopt a dual-maximizing theory, where the condition is being oneself, and the associated weight  $w$  is 0. Doing so yields the following theory, proposed by Sider (1993):

**Self-Other Utilitarianism (SOU):** An act  $a$  is permissible *iff* (i)  $a$  maximizes the utility of all subjects, or (ii)  $a$  maximizes the utility of others.

SOU allows agents to perform acts that decrease overall utility, as long as by doing so they don’t decrease the utility of anyone but themselves. Thus this theory allows for permissible self-sacrifice.

SOU yields the desired prescriptions in Parental Sacrifice. Suppose we represent your options as follows (where the numbers indicate the utility of the relevant subject given that act):

Options	Self	Child	Total
Don’t Work Overtime:	10	10	20
Work Overtime:	8	11	19

tion in a number of different ways. And as Vessel (2010) notes, on this understanding of supererogation, even Utilitarians will admit to some supererogatory acts (since there can be pairs of acts which both maximize overall utility, one which is worse for the agent and better for others, and one which is better for the agent and worse for others). Nevertheless, the point remains that there are natural ways of understanding supererogation according to which self-centered dual-maximizing theories can recognize a much broader range of supererogatory acts than typical Utilitarian theories can.

<sup>8</sup>We can formulate multi-maximizing theories as follows, with respect to  $m$  conditions  $\phi_i$  with associated weights  $w_i$ :

**Multi-Maximizing Theories:** An act  $a$  is permissible *iff* for some  $is$ ,  $a$  maximizes the sum of (the utility of those who don’t satisfy  $\phi_i$ ) plus ( $w_i$  times the utility of those who do satisfy  $\phi_i$ ).

Multi-maximizing theories offer an appealing generalization of dual-maximizing theories, as they can simultaneously permit disregarding others, those one isn’t close to, and oneself. See Portmore (2008) and Vessel (2010) for formulations of multi-maximizing theories that incorporate multiple kinds of disregarding.

According to Utilitarianism, it's impermissible for you to work overtime, since this option fails to maximize total utility. But according to SOU, both options are permissible. It's permissible for you to not work overtime because this option maximizes total utility. And it's permissible for you to work overtime because this option maximizes the utility of others.

### 2.3 Problems For Dual-Maximizing

Although dual-maximizing theories like SOU are attractive in some respects, they have implausible consequences. To see this, let's consider a counterexample to SOU inspired by Splawn (2001) that's helpful in both identifying the problem with dual-maximizing, and identifying how to remedy this defect.

Consider a version of Parental Sacrifice in which you have the option of making both small and large sacrifices:

**Variable Parental Sacrifice:** Your child's birthday is coming up. You have the option of working overtime for a week in order to buy your child a gift, or for two weeks in order to buy your child an even better gift. Both options would significantly decrease your utility – in both cases, this decrease would be greater than the utility your child would receive from the gift. No one else's utility would be affected.

In this case we might represent your options as follows:

Options	Self	Child	Total
Don't Work Overtime:	10	10	20
Work 1 Week of Overtime:	8	11	19
Work 2 Weeks of Overtime:	7	12	19

According to SOU, it's permissible to not work overtime, since this option maximizes total utility. Likewise, it's permissible to work two weeks of overtime, since this option maximizes the utility of others. But it's impermissible to work one week of overtime, since this option maximizes neither total utility nor the utility of others.

This is an implausible result.<sup>9</sup> Given the range of options usually available to agents, it would follow that agents could only make total-utility-decreasing self-sacrifices

---

<sup>9</sup>Splawn (2001) offers a criticism of SOU along these lines. Portmore (2008) offers a response, although he assesses the worry with respect to a kind of multi-maximizing theory he calls "Schefflerian Utilitarianism" instead of SOU, where Schefflerian Utilitarianism takes an act to be permissible *iff* it either (i) maximizes the utility of others, or (ii) maximizes the sum of (the utility of others) plus (10 times one's own utility). Like SOU, Schefflerian Utilitarianism has the implausible consequence that only the first and third options in Variable Parental Sacrifice are permissible. Portmore responds to this objection by arguing that working only one week of overtime *should* be impermissible, because you would be "unreasonably selfish" to sacrifice your utility in this inefficient way. If you work two weeks of overtime then others get two-thirds of the utility you sacrifice, whereas if you work one week of overtime then others get only one half of the utility you sacrifice.

I think it's false that non-maximizing self-sacrifices are only permissible when they're maximally efficient.

of the most extreme kind. Modest gifts and sacrifices for loved ones would not be permissible. Only the most extreme and lavish gifts would be permitted.

This points to a general problem with dual-maximizing theories. The problem is that dual-maximizing theories are fundamentally disjunctive. They treat maximizing everyone's utility and maximizing the appropriately weighted utility of privileged and disregarded subjects as distinct goals, and take acts to be permissible *iff* they're the best at achieving either of these goals. But by treating these goals as distinct, they're unable to allow for natural trade-offs between them.

How might one modify dual-maximizing theories in order to allow for such trade-offs? Here is one natural thought. The difference between the first and second goals of SOU can be seen as a difference in the value of the weight  $w$ . The first goal – maximizing utility – is a case where  $w = 1$ , and the utility of disregarded subjects isn't discounted, while the second goal – maximizing the utility of the privileged subjects – is a case where  $w = 0$  and the utility of disregarded subjects is completely discounted. So if we want to allow for tradeoffs between these two goals, one natural thought is to permit any act which maximizes weighted utility for any value of  $w$  between these two extremes.

Unfortunately, this natural thought doesn't pan out. On this proposal, the total weighted utility of each act will be:

Options	Self	Child	Weighted Total
Don't Work Overtime:	10w	10	10 + 10w
Work 1 Week of Overtime:	8w	11	11 + 8w
Work 2 Weeks of Overtime:	7w	12	12 + 7w

But no value of  $w$  will make working one week of overtime maximize weighted utility. If  $w = 1$ , then not working overtime will yield a higher weighted utility (20) than working one week of overtime (19). And if  $w < 1$ , then working two weeks of overtime will yield a higher weighted utility than working one week of overtime. This is because working one week of overtime is just like working two weeks of overtime, except it shifts one unit of utility from your child to yourself. And since we're discounting the utility assigned to you (because  $w < 1$ ), this shift will decrease weighted utility.

Why does working one week of overtime seem permissible? It seems permissible because it seems you should be free to determine how much of your utility you're willing to sacrifice. You can choose not to sacrifice any of your utility and not work overtime, you can choose to sacrifice a lot of utility, and work two weeks of overtime, or you can choose to just sacrifice some of your utility – up to two utility, say – and work one week of overtime. The problem with the natural thought sketched above is that what we want to vary is not the *degree* to which you discount your utility, but the *amount* of utility which you're willing to discount.

---

That stance would entail that if we could make a sacrifice to work a week of overtime to buy our child a present, but could also make a more efficient sacrifice by selling our leg in order to allow our child to go to college, then only selling our leg (or making no sacrifice at all) would be permissible. This is implausible.

Let's turn to theories which allow us to do that.

## 3 Variable-Disregarding Theories

### 3.1 The Structure of Variable-Disregarding Theories

In this section I'll introduce *variable-disregarding theories*. In the interest of accessibility, I'll introduce them in two stages. In section 3.1.1 I'll describe these theories informally. Then in section 3.1.2 I'll describe these theories more formally, in a manner mirroring my presentation of dual-maximizing theories in section 2.1.

#### 3.1.1 Variable-Disregarding Theories, Take I

We want to modify Utilitarianism to permit giving less weight to the utility of some group of subjects. And, as we saw in the last section, we also want our theory to consider various amounts of utility that we might discount. Let's look at one way to do this.

There are some acts that will bring about more utility for a group than any other act; i.e. will bring about the maximum utility for that group.<sup>10</sup> And there are other acts that will lead to drops in utility for this group, relative to this maximum. One way to give less weight to the utility of this group is to assign less importance to these drops – that is, to treat their *effective* utility as closer to their maximum than it really is. So suppose we're giving no weight to the utility of a subject. If their maximum utility is 10, and their utility given a certain act is 6, we can disregard this drop in their utility by treating their utility as 10 for the purposes of evaluating this act. Alternatively, suppose we're giving half weight to the utility of a subject. If their maximum utility is 10, and their utility given a certain act is 6, we can give half weight to this drop in their utility by treating their utility as 8 for the purposes of evaluating this act.

Now, we also want the theory to permit discounting various *amounts* of utility. Here is one way to do that. Consider different magnitudes of drops in utility. Some acts will maximize effective utility if we discount drops in a group's utility of up to some positive finite size. Some acts will maximize effective utility if we discount *all* drops in a group's utility (equivalently: will maximize effective utility if we discount all drops in a group's utility of up to size  $\infty$ ). And some acts will maximize utility *simpliciter* (equivalently: will maximize effective utility if we discount all drops in a group's utility of up to size 0). We can ensure that our theory permits discounting various amounts of utility by taking all of these different acts to be permissible.

Putting these thoughts together, we can say that an act is permissible *iff*, for some amount of utility, discounting drops in a group's utility of up to that amount makes

---

<sup>10</sup>In certain infinity cases, there will be no act that maximizes utility; I'm bracketing such complications here (cf. footnote 12).

the act maximize effective utility.<sup>11</sup> I'll call a theory of permissible action of this form a *variable-disregarding theory*.

Let's work through an example. Suppose we want a variable-disregarding theory that allows you to give *no* weight to the utility of those you don't feel close to. Then we can take an act to be permissible *iff* there's some amount of utility of those you're not close to that you can disregard which makes that act maximize effective utility.

So consider the following case:

**Dividing Goods:** You're considering whether to split some good evenly between a friend of yours and a stranger, or give all of the good to your friend. If you split the good evenly both will get a fair amount of utility, whereas if you give all of it to your friend, your friend's utility will be slightly higher, but the stranger's utility will be much lower.

We might represent your options in Dividing Goods as follows:

Options	Friend	Stranger	Total
Split Evenly:	9	9	18
Favor Friend:	10	5	15

Which of these acts is permissible?

Let's start with Split Evenly. Is there some amount of utility of the stranger which you can discount which makes Split Evenly maximize effective utility? Here the answer is trivially yes – if we don't discount *any* amount of utility, then Split Evenly will maximize utility. So Split Evenly is permissible.

Now consider Favor Friend. Is there some amount of utility of the stranger which you can discount which makes Favor Friend maximize effective utility? Well, suppose we discount an arbitrarily large amount of the stranger's utility. Favor Friend drops the utility of the stranger by 4, relative to Split Evenly. And if we give that drop in utility no weight — we give it none of the importance we normally would — then we can treat the stranger's effective utility as 9. So the total effective utility of Favor Friend will be 19. By contrast, disregarding drops in the stranger's utility won't change the effective utility of Split Evenly, since in Split Evenly there's no drop in the stranger's utility to disregard. Thus the effective utility of Split Evenly will still be 18. Since Favor Friend has a higher effective utility, it follows that Favor Friend is also permissible.

Options	Friend	Stranger	Effective Total
Split Evenly:	9	9	18
Favor Friend:	10	5 (9)	19

(Of course, we didn't have to discount an *arbitrarily large* amount of the stranger's utility to get this result. Discounting drops in the stranger's utility of up size 4 would yield the same result.)

---

<sup>11</sup>I'm simplifying slightly here by talking in terms of a group's utility instead of the utility of individuals in a group; for a more precise presentation, see section 3.1.2.

Now consider a different theory. Suppose we wanted a theory that only allows you to *partially* disregard – say, give half weight to – the utility of those you’re not close to. What acts would be permissible in Dividing Goods?

As before, Split Evenly will be permissible, since if we don’t disregard *any* utility, Split Evenly will maximize effective utility. What about Favor Friend? Is there some amount of utility of the stranger you can partially disregard (give half weight to) according to which Favor Friend maximizes effective utility?

Well, suppose we discount an arbitrarily large amount of the stranger’s utility. Favor Friend decreases the stranger’s utility by 4. And if we give that drop in utility half weight — we give it half of the importance we normally would — then we can treat the stranger’s effective utility as 7. So the total effective utility of Favor Friend will be 17, while the total effective utility of Split Evenly will remain 18. So Favor Friend will not maximize effective utility in this case either.

Options	Friend	Stranger	Effective Total
Split Evenly:	9	9	18
Favor Friend:	10	5 (7)	17

More generally, a little thought reveals that there’s no amount of utility we can discount which will make Favor Friend maximize effective utility. So if we’re required to give the stranger’s utility at least half weight, giving all of the good to your friend is impermissible.

### 3.1.2 Variable-Disregarding Theories, Take II

Now let’s formulate variable-disregarding theories more precisely.

Let’s begin by introducing some terminology. As in section 2.1, let  $\phi$  be the condition that picks out the subjects we’re disregarding, and let  $w \in [0, 1]$  be the weight assigned to the utility of the disregarded subjects. Let  $u_a(s)$  be the utility of subject  $s$  given act  $a$ . And let  $MAX_s$  be the maximum utility that subject  $s$  could end up with in a given decision problem.<sup>12</sup> (Thus in Dividing Goods,  $MAX_{\text{friend}}$  is 10, and  $MAX_{\text{stranger}}$  is 9.)

Now let’s introduce the *effective utility* $_{(\phi,w,v)}(a)$  of an act  $a$ , where  $v$  is a sequence of values  $v = \{v_{s_1}, v_{s_2}, \dots\}$  assigned to the  $\phi$ -subjects that corresponds to the amount of their utility we’re discounting. The *effective utility* $_{(\phi,w,v)}(a)$  of an act  $a$  is the sum for each subject  $s$  of (1)  $MAX_s - w(MAX_s - u_a(s))$ , if (given  $a$ )  $s$  satisfies condition  $\phi$  and  $s$ ’s utility is within  $v_s$  of  $MAX_s$ , or (2)  $u_a(s)$ , otherwise.

So when  $w = 0$ , the effective utility $_{(\phi,0,v)}$  of an act is just the sum of the utilities of each subject, with the following exception: if a subject satisfies  $\phi$ , and has a utility

---

<sup>12</sup>In certain infinity cases there might not be a maximum utility a subject could have. (E.g., a case in which you can choose any natural number  $n$ , and receive that much utility.) Since such cases pose problems for both standard-maximizing theories and variable-disregarding theories, I’ll ignore these complications here.

within  $v_s$  of their maximum, then we treat their utility as that maximum. So the effective utility  $u_{(\phi,0,v)}$  of an act ignores drops in utility of up to size  $v_s$  for each subject  $s$  who satisfies  $\phi$ .

When  $w > 0$ , these drops in utility aren't ignored, just given less weight. So when  $w = 0.5$ , the effective utility  $u_{(\phi,0.5,v)}$  of an act is the sum of the utilities of each subject, except that subjects who satisfy  $\phi$  and have a utility within  $v_s$  of their maximum are treated as having a utility that's only half as far from their maximum as it actually is; e.g., if their utility is 6 and their maximum is 10, they're treated as having a utility of 8. Thus the effective utility  $u_{(\phi,0.5,v)}$  of an act gives half weight to drops in utility of up to size  $v_s$  for each subject  $s$  who satisfies  $\phi$ .

With this terminology in hand, we can formulate variable-disregarding theories as follows, with respect to some condition  $\phi$  that picks out the disregarded subjects, and some weight  $w \in [0, 1]$  that corresponds to the degree to which disregarded subjects are taken into account:

**Variable-Disregarding Theories:** An act  $a$  is permissible *iff*, for some assignment  $v$  of values to subjects who satisfy  $\phi$ ,  $a$  maximizes effective utility  $u_{(\phi,w,v)}$ .

So when  $w = 0$ , variable-disregarding theories tell us that  $a$  is permissible *iff* there's some assignment of values  $v$  to the  $\phi$ -subjects such that disregarding drops in utility of up to those amounts makes  $a$  maximize effective utility. And when  $w > 0$ , variable-disregarding theories tell us that  $a$  is permissible *iff* there's some assignment of values  $v$  to  $\phi$ -subjects such that partially disregarding (to a degree determined by  $w$ ) drops in utility of up to those amounts makes  $a$  maximize effective utility.

Let me pause for a moment to say a bit more about the roles of  $\phi$ ,  $w$  and  $v$ . The first two are parameters that are fixed by a variable-disregarding theory. In spelling out what variable-disregarding theory we're using, we have to specify a particular condition  $\phi$  and a particular value  $w$ . By contrast, the  $v$ s – the amounts we're disregarding – are not fixed by a variable-disregarding theory. For any given variable-disregarding theory, we consider *all* of the different possible  $v$ s – all of the different possible amounts of utility we might disregard. And we take an act to be permissible *iff*, for at least one of these  $v$ s, that act maximizes effective utility.

To get a feel for this terminology, let's work through the Dividing Goods case again. Suppose we adopt a variable-disregarding theory that allows you to disregard the utility of those you don't feel close to. So  $\phi$  is the condition of being someone you don't feel close to; call this condition  $nc$ . And let's suppose the theory allows you to completely disregard the utility of those you're not close to, so that  $w = 0$ . Recall that your options in the Dividing Goods case are as follows:

Options	Friend	Stranger	Total
Split Evenly:	9	9	18
Favor Friend:	10	5	15

If we set  $v_{\text{stranger}} = 0$ , and so don't disregard any drops in the stranger's utility, then Split Evenly maximizes effective utility  $u_{(nc,0,0)}$ :

Options	Friend	Stranger	Effective Total ( $v = 0$ )
Split Evenly:	9	9	18
Favor Friend:	10	5	15

If we set  $v_{\text{stranger}} = \infty$ , and so disregard drops in the stranger's utility of any size, then Favor Friend maximizes effective utility $_{(nc,0,\infty)}$ :

Options	Friend	Stranger	Effective Total ( $v = \infty$ )
Split Evenly:	9	9	18
Favor Friend:	10	5 (9)	19

Since each act maximizes effective utility for some value of  $v$ , this theory takes both acts to be permissible.

Suppose instead we adopt a more plausible theory on which  $w = 0.5$ , so that we can only *partially* disregard drops in utility for those we're not close to. If we set  $v_{\text{stranger}} = \infty$ , then we'll give half weight to drops in the stranger's utility (because  $w = 0.5$ ), and we'll do so for drops in utility of any size (because  $v_{\text{stranger}} = \infty$ ). Thus Split Evenly will maximize effective utility $_{(nc,0.5,\infty)}$ :

Options	Friend	Stranger	Effective Total ( $v = \infty$ )
Split Evenly:	9	9	18
Favor Friend:	10	5 (7)	17

More generally, one can show that Split Evenly will maximize effective utility for any value of  $v$ . Thus if  $w = 0.5$ , we'll get the result that only Split Evenly is permissible.

In what follows, we'll mostly focus on cases in which  $w = 0$  and the utility of disregarded subjects is completely discounted. To simplify notation, I'll leave the  $w$  index implicit when it's clear that  $w = 0$ .

Variable-disregarding theories share all of the advantages of dual-maximizing theories. We can use them to permit self-centered options, options which favor one's friends and family, and options which favor others. And variable-disregarding theories will also allow for a broad range of supererogatory acts.

Just as dual-maximizing theories are a special case of multi-maximizing theories, variable-disregarding theories are a special case of what we might call "multivariable-disregarding theories", which allow for any number of different disregarded groups with different weights. But for simplicity, I'll focus on variable-disregarding theories in what follows.<sup>13</sup>

<sup>13</sup>Let  $\phi_1, \dots, \phi_m$  be a sequence of conditions, and let  $v_i = \{v_i^{s1}, v_i^{s2}, \dots\}$  be a sequence of values assigned to the  $\phi_i$ -subjects. Let the *effective utility* $_{(\phi_1, \dots, \phi_m, w_1, \dots, w_m, v_1, \dots, v_m)}(a)$  of an act  $a$  in a given decision problem be the sum over each subject  $s$  of (1)  $\text{MAX}_s - w_i(\text{MAX}_s - u_a(s))$ , if (given  $a$ ) for some  $i, s$  satisfies  $\phi_i$  and  $s$ 's utility is within  $v_i^s$  of  $\text{MAX}_s$ , or (2)  $u_s(a)$ , otherwise.

We can formulate multivariable-disregarding theories, with respect to  $m$  conditions  $\phi_i$  with associated weights  $w_i$ , as follows:

**Multivariable-disregarding Theories:** An act  $a$  is permissible *iff*, for some  $m$  assignments  $v_1, \dots, v_m$  of values to subjects,  $a$  maximizes effective utility $_{(\phi_1, \dots, \phi_m, w_1, \dots, w_m, v_1, \dots, v_m)}$ .

### 3.2 Self-Discounting Utilitarianism

Suppose we want to handle Parental Sacrifice by modifying Utilitarianism to allow us to disregard ourselves. As we've seen, we can do this by adopting a dual-maximizing theory, SOU. But we can also do this by adopting a variable-disregarding theory. Taking the condition  $\phi$  to be being oneself (call this condition  $s$ ), and taking the associated weight to be  $w = 0$  yields the following theory:

**Self-Discounting Utilitarianism (SDU):** An act  $a$  is permissible *iff*, for some value  $v$ ,  $a$  maximizes effective utility $_{(s,v)}$ .

So SDU tells us that  $a$  is permissible *iff* there's some value  $v$  such that disregarding drops in one's own utility of up to that amount makes  $a$  maximize utility.

SDU yields the desired prescriptions in Parental Sacrifice. If we set  $v = 0$ , and so don't disregard any drops in our utility, then not working overtime maximizes effective utility $_{(s,0)}$ :

Options	Self	Child	Effective Total ( $v = 0$ )
Don't Work Overtime:	10	10	20
Work Overtime:	8	11	19

If we set  $v = \infty$ , and so disregard drops of any amount to our utility, then working overtime maximizes effective utility $_{(s,\infty)}$ :

Options	Self	Child	Effective Total ( $v = \infty$ )
Don't Work Overtime:	10	10	20
Work Overtime:	8 (10)	11	21

Thus both options are permissible.

Unlike SOU, SDU also yields the desired prescriptions in Variable Parental Sacrifice. If we set  $v = 0$ , and so don't disregard drops of any amount to our utility, then not working overtime maximizes effective utility $_{(s,0)}$ :

Options	Self	Child	Effective Total ( $v = 0$ )
Don't Work Overtime:	10	10	20
Work 1 Week of Overtime:	8	11	19
Work 2 Weeks of Overtime:	7	12	19

If we set  $v = 2$ , and so disregard drops of up to 2 units to our utility, then working one week of overtime maximizes effective utility $_{(s,2)}$ :

Options	Self	Child	Effective Total ( $v = 2$ )
Don't Work Overtime:	10	10	20
Work 1 Week of Overtime:	8 (10)	11	21
Work 2 Weeks of Overtime:	7	12	19

If we set  $v = \infty$ , and so disregard drops of any amount to our utility, then working two weeks of overtime maximizes effective utility $_{(s,\infty)}$ :

Options	Self	Child	Effective Total ( $v = \infty$ )
Don't Work Overtime:	10	10	20
Work 1 Week of Overtime:	8 (10)	11	21
Work 2 Weeks of Overtime:	7 (10)	12	22

So SDU yields the desired result that all three options are permissible.

### 3.3 Further Problems

SOU yields plausible verdicts in Parental Sacrifice, and SDU yields plausible verdicts in both Parental Sacrifice and Variable Parental Sacrifice. But both SOU and SDU yield strange results in cases where other subjects are willing to make sacrifices as well, such as the following:

**Two Parental Sacrifices:** Your child's birthday is coming up. You have the option of working overtime in order to buy your child a gift. Doing so would significantly decrease the utility of both you and your partner – the decrease for each of you would be greater than the utility your child would receive from the gift. But your partner is willing to make the sacrifice, and no one else's utility would be affected.

In this case we might represent your options as follows:

Options	Self	Partner, Child	Total
Don't Work Overtime:	10	10, 10	30
Work Overtime:	8	8, 11	27

According to SOU it will be impermissible to work overtime, since doing so maximizes neither total utility nor the utility of everyone else. Likewise, according to SDU it will be impermissible to work overtime, since regardless of how much of your utility we disregard, not working overtime will maximize effective utility $_{(s,v)}$ . If  $v < 2$ , then the effective utility $_{(s,v)}$  of not working overtime will be 30 instead of 27:

Options	Self	Partner, Child	Effective Total ( $v < 2$ )
Don't Work Overtime:	10	10, 10	30
Work Overtime:	8	8, 11	27

And if  $v \geq 2$ , then the effective utility $_{(s,v)}$  of not working overtime will be 30 instead of 29:

Options	Self	Partner, Child	Effective Total ( $v \geq 2$ )
Don't Work Overtime:	10	10, 10	30
Work Overtime:	8 (10)	8, 11	29

So no matter what we set  $v$  to, not working overtime maximizes effective utility. Thus SDU takes not working overtime to be obligatory.

These are implausible verdicts, especially when combined with the verdicts SOU and SDU make in Parental Sacrifice. If it's permissible for you to willingly sacrifice some of your well-being for your child, then it should be permissible for your partner to do so as well. The fact that in Two Parental Sacrifices your partner cannot themselves act in order to make this sacrifice doesn't seem morally relevant.

Moreover, if we accept SOU or SDU, we have to accept an odd conflict between our permissions and those of others. To see this, return to the Parental Sacrifice case. When it's the parent making the decision, SOU and SDU both yield the verdict that it's permissible for the parent to work overtime. But suppose that working overtime also requires the parent's manager to sign off. Then it will be the manager making the decision. And since the manager isn't one of the parties involved, SOU and SDU will both yield the same verdicts as standard Utilitarianism, namely, that it's impermissible for the parent work overtime. So while SOU and SDU allow the parent to make this self-sacrifice, it won't allow the manager to let this self-sacrifice take place.

This yields an uncomfortable tension between our permissions and those of others. While SOU and SDU will permit you to make self-sacrifices, they'll also require everyone else to prevent you from doing so. This has the consequence that it would be virtually impossible to carry out these self-sacrifices in a society where everyone did the right thing. This is because most of our potential sacrifices require the acquiescence of others to go through, and others would be forbidden from allowing these sacrifices to take place. Banking agents would be obligated to not allow the relevant checks to clear, store owners would be obligated to not sell the relevant products to the relevant people, bosses would be obligated to refuse to allow the relevant people to work overtime, and so on. This would be a strange state of affairs.

Let's take a step back. Two Parental Sacrifices seems similar in morally relevant respects to Parental Sacrifice and Variable Parental Sacrifice, and it seems the explanation of why it's permissible to work overtime in each case should be the same. But while a self-other asymmetry could potentially explain the permissibility of working overtime in Parental Sacrifice and Variable Parental Sacrifice, it can't explain the permissibility of working overtime in Two Parental Sacrifices. For in Two Parental Sacrifices, working overtime doesn't just lower your utility (which a self-other asymmetry could allow you to ignore), it also lowers the net utility of others. Thus we have reason to be skeptical that a self-other asymmetry is what's at the heart of these cases.

More generally, this suggests that the various authors who have taken Utilitarianism to have a special problem with "altruistic" objections have missed the forest for the trees. They've identified a particular batch of problem cases, and have taken those problem cases to encapsulate the whole problem. What cases like Two Paternal Sacrifices show is that these "altruistic" objections are really just an instance of a broader interpersonal phenomena. And in order to address this broader phenomenon, we need to do more than to just permit agents to disregard their own utility.

## 4 Consent

### 4.1 Consent-Discounting Utilitarianism

In his discussion of these “altruistic” objections, Slote raises the following thought:

“It has been suggested to me that the reason we are allowed to harm ourselves or avoid some benefit, where we should not be permitted to harm another person or prevent her from receiving a similar benefit, lies in the consent implicit in actions we do to ourselves. If I harm myself to avoid a benefit, I presumably do this willingly, whereas the agent whom I refuse to benefit does not consent to this neglect (and when she does there is nothing wrong with what I do). It might be then thought that the moral asymmetry we have noted is not a deep feature of morality, but rather derivative from and justifiable in terms of the moral importance of consent.”<sup>14</sup>

Although Slote goes on to reject this suggestion, I think this is exactly right. The key moral distinction in these cases is not the distinction between self and others, it’s the distinction between those who consent and those who do not.

If this thought is correct, then the modification of Utilitarianism suggested in section 3.2 is too cautious. Instead of modifying standard Utilitarianism to permit disregarding oneself, we should modify it to permit disregarding those who consent. That is, we should change the condition  $\phi$  picking out who we can disregard from “oneself” to “those who consent to the act”. Then we can explain the appearance of a self-other asymmetry in such cases as deriving from the asymmetry in what subjects typically consent to.<sup>15</sup>

In some of the literature, the term “consent” is used to mean something like waiving a right, or releasing someone from a duty to you.<sup>16</sup> That is, consent is used to mean something which is by definition morally significant. This is not how I’m using the term “consent”. I’m using the term “consent” in its more colloquial sense, to denote something like agreement or acquiescence. And while one might plausibly take agreement or acquiescence to have moral significance, we can use these notions without presupposing that they have moral significance.

It’s plausible that consent (so understood) must satisfy certain conditions in order to be morally relevant. For example, we might want to require the consenting subject to be informed, competent, free from coercion, and so on, in order for their consent to count. For now, I’ll simply take for granted that there are some conditions of this

---

<sup>14</sup>See Slote (1984), p.190-191.

<sup>15</sup>The suggestion is that *this* apparent self-other asymmetry can be explained in terms of consent. It is not that *all* self-other asymmetries can be explained in terms of consent. For example, if one chooses to deal with “demandingness” objections by allowing agents to partially disregard the utility of others, then one is positing a kind of self-other asymmetry. But this self-other asymmetry cannot be explained in terms of consent.

<sup>16</sup>E.g., see Liberto (2017).

kind, and I'll call consent that satisfies these conditions "morally relevant consent", or "consent<sub>m</sub>" for short. We'll return to examine the question of what consent<sub>m</sub> is in section 4.2.

Now let's return to the task at hand: formulating a theory that permits disregarding those who consent. We can construct a variable-disregarding theory that does this by taking the condition  $\phi$  to be consent<sub>m</sub> to the act in question (call this condition  $c$ ), and taking the associated weight to be  $w = 0$ :

**Consent-Discounting Utilitarianism (CDU):** An act  $a$  is permissible *iff*, for some assignment of values  $v$  to subjects who satisfy  $c$ ,  $a$  maximizes effective utility <sub>$c, v$</sub> .

So CDU tells us that  $a$  is permissible *iff* there's some assignment of values  $v$  to subjects such that disregarding drops in utility of up to those amounts for subjects who consent<sub>m</sub> to  $a$  makes  $a$  maximize utility.<sup>17</sup>

Now let's see how CDU handles the three cases we've discussed. In all of these cases you consent<sub>m</sub> to whichever action you perform, and for concreteness I'll assume your child also consents<sub>m</sub>, though we get the same results regardless of whether your child consents<sub>m</sub> or not. Let's start with Parental Sacrifice. If we set  $v = \{0, 0\}$  (i.e., set  $v_{\text{self}} = 0$  and  $v_{\text{child}} = 0$ ), then not working overtime maximizes effective utility <sub>$c, \{0, 0\}$</sub> :

Options	Self	Child	Effective Total ( $v = 0, 0$ )
Don't Work Overtime:	10	10	20
Work Overtime:	8	11	19

And if we set  $v = \{\infty, 0\}$ , then working overtime will maximize effective utility <sub>$c, \{\infty, 0\}$</sub> :

Options	Self	Child	Effective Total ( $v = \infty, 0$ )
Don't Work Overtime:	10	10	20
Work Overtime:	8 (10)	11	21

So CDU will take both options to be permissible.

Now consider Variable Parental Sacrifice. If we set  $v = \{0, 0\}$  (i.e., set  $v_{\text{self}} = 0$  and  $v_{\text{child}} = 0$ ), then not working overtime maximizes effective utility <sub>$c, \{0, 0\}$</sub> :

Options	Self	Child	Effective Total ( $v = 0, 0$ )
Don't Work Overtime:	10	10	20
Work 1 Week of Overtime:	8	11	19
Work 2 Weeks of Overtime:	7	12	19

If we set  $v = \{2, 0\}$ , then working one week of overtime maximizes effective utility <sub>$c, \{2, 0\}$</sub> :

---

<sup>17</sup>Note that although I've described CDU as disregarding *drops* in utility of those who consent, it also effectively disregards *increases* in utility of those who don't consent to the increase. E.g., consider a subject who doesn't consent to the option  $a$  that maximizes their utility, and instead prefers an option  $b$  that brings them a lower utility instead. Since the subject consents to  $b$ , they'll be treated as if they had the utility they would have given  $a$  for the purposes of evaluating  $b$ , and thus the increase in utility that  $a$  brings them won't end up telling in favor of  $a$ .

Options	Self	Child	Effective Total ( $v = 2, 0$ )
Don't Work Overtime:	10	10	20
Work 1 Week of Overtime:	8 (10)	11	21
Work 2 Weeks of Overtime:	7	12	19

And if we set  $v = \{\infty, 0\}$ , then working two weeks of overtime maximizes effective utility<sub>(c, {∞,0})</sub>:

Options	Self	Child	Effective Total ( $v = \infty, 0$ )
Don't Work Overtime:	10	10	20
Work 1 Week of Overtime:	8 (10)	11	21
Work 2 Weeks of Overtime:	7 (10)	12	22

So CDU will take all three options to be permissible.

Now consider Two Parental Sacrifices. In this case there's a third subject, your partner, who will also consent<sub>m</sub> to whatever act you perform. If we set  $v = \{0, 0, 0\}$  (i.e., set  $v_{\text{self}} = 0$ ,  $v_{\text{partner}} = 0$ , and  $v_{\text{child}} = 0$ ), then not working overtime maximizes effective utility<sub>(c, {0,0,0})</sub>:

Options	Self	Partner, Child	Effective Total ( $v = 0, 0, 0$ )
Don't Work Overtime:	10	10, 10	30
Work Overtime:	8	8, 11	27

And if we set  $v = \{\infty, \infty, 0\}$ , then working overtime maximizes effective utility<sub>(c, {∞,∞,0})</sub>:

Options	Self	Partner, Child	Effective Total ( $v = \infty, \infty, 0$ )
Don't Work Overtime:	10	10, 10	30
Work Overtime:	8 (10)	8 (10), 11	31

So CDU will take both options to be permissible.

## 4.2 Morally Relevant Consent

### 4.2.1 Informed, Competent, and Uncoerced Consent

Now let's consider what kinds of conditions we want an account of consent<sub>m</sub> to take into consideration.<sup>18</sup>

First, we want consent<sub>m</sub> to require that the subject is informed. If a subject would consent to having their money taken, but only because they've been deceived into

<sup>18</sup>In addition to the informed, competent, and uncoerced conditions described below, some have suggested to me that one might add a fourth condition that requires the act to line up with the subject's desires or deep commitments in order for them to consent<sub>m</sub>. (For a discussion of a view along these lines, see Kiloren (2019).) This clause would rule out consent that was merely made because the subject was trying to be nice, or felt obligated to consent, even though they didn't really want to. This is an interesting suggestion, and I take a version of CDU incorporating such a clause to be a viable option.

thinking that they'd get it back, we should not take them to be consenting<sub>m</sub>. Likewise, if a subject would consent to giving up their salary bonus, but only because they falsely believe that if they do, it will go to charity, we should not take them to be consenting<sub>m</sub>.<sup>19</sup>

Second, we want consent<sub>m</sub> to require that the subject is competent. If a subject would consent to being lit on fire, but only because they're mentally ill, we should not take them to be consenting<sub>m</sub>. Likewise, if a subject would consent to having all of their property taken, but only because they're very young or on powerful drugs, we should not take them to be consenting<sub>m</sub>.<sup>20</sup>

Adding this competence clause raises a potential worry. Suppose one identifies competence with rationality. On some conceptions of rationality, a rational agent would never consent to a decrease in utility, and never dissent to an increase in utility. If so, and we require competence for consent<sub>m</sub>, then no one will ever consent<sub>m</sub> to acts which decrease their utility. But then CDU becomes equivalent to standard Utilitarianism, since the deviations from maximizing utility that CDU allows will never come into play.

This worry requires an understanding of competence that equates "competent" with "ideally prudentially rational", or something of that kind. This strong understanding of the competence condition is not what I have in mind. The notion of competence I have in mind here is something weaker, something like the kind of competence required in order for one's actions to be subject to moral evaluation. For example, normal human adults might be thought to be subject to such evaluation, while dogs are not. Thus one might take normal human adults, but not dogs, to be competent.<sup>21</sup> And if we assess CDU using this weaker notion of competence, this worry will not arise.

Third, we want consent<sub>m</sub> to require that the subject is uncoerced. Suppose you live under the rule of a sadistic dictator, who can scan your brain to detect whether you've consented to various things. And suppose the dictator will scan your brain and react as follows: if you consent to their cutting off your hand then they will do so, and if you do not consent to their cutting off your hand then they will kill you. If you know all of this then you might reasonably consent to them cutting off your hand. But we

---

<sup>19</sup>For some discussions of how to spell out the notion of informed consent, see Manson and O'Neill (2007), Beauchamp and Childress (2008) and Eyal (2012).

<sup>20</sup>For a discussion of some of the issues that arise with respect to assessing competence, see Grisso and Appelbaum (1998), Beauchamp and Childress (2008) and Charland (2011).

<sup>21</sup>Granted, this notion of competence seems to come in degrees. Likewise, the extent to which a subject's actions are subject to moral evaluation seems to come in degrees (see Charland (2011)). If one wanted to take this into consideration, one might consider modifying CDU to take *degrees* of competent consent into account. Alternatively, one might simply take the strength of deontic obligations to be something which admits of variable degree – a highly competent agent can do something **wrong**, while a barely competent agent can only do something **wrong**. These are interesting issues, but not ones I'll take up here. For simplicity, I'll simply proceed on the assumption that a binary notion of competence will suffice.

should not take you to be consenting<sub>m</sub> to their cutting off your hand.<sup>22</sup>

#### 4.2.2 Actual versus Ideal Consent

We want consent<sub>m</sub> to be informed, competent, and uncoerced. But there are a couple different ways in which one might incorporate these kinds of conditions into an account of consent<sub>m</sub>. First, one might take a subject to consent<sub>m</sub> to *a* iff they consent to *a* and satisfy the relevant conditions. We might call this the *actual* way of incorporating these conditions, since it requires the subject to actually satisfy these conditions in order to count as consenting<sub>m</sub>. Second, one might take a subject to consent<sub>m</sub> to *a* iff the following counterfactual is true: if the agent were to perform this act, and if *s* were informed, competent and uncoerced, then *s* would consent to the act.<sup>23</sup> We might call this the *ideal* way of incorporating these conditions, since it only requires an idealized version of the subject who satisfies these conditions to consent in order for the subject to count as consenting<sub>m</sub>.

I take both the actual and ideal approaches to have intuitive appeal.<sup>24</sup> And in some cases our intuitions regarding these approaches conflict. For example, suppose a member of an isolated tribe with no experience of modern medicine is bleeding to death. And suppose a medic has found the subject, and must insert a needle into their arm in order to give them the blood transfusion needed to save their life. In light of the subject's lack of familiarity with modern medicine, and their lack of understanding of what the needle insertion is for, they don't consent to having the needle inserted into their arm. But suppose that if the subject were informed of what the needle was for, and understood that this was necessary to save their life, they would consent.

Now, in this case should we treat the subject as consenting<sub>m</sub> to have the needle inserted into their arm? That is, with respect to the morally relevant notion of consent, should we treat them as consenting? I think most feel torn about this case. On the one hand, they don't actually consent to having the needle inserted into their arm, and that seems morally important. On the other hand, they clearly would consent if they were appropriately informed, and that seems morally important as well.

I think both ways of understanding consent<sub>m</sub> are viable. So in what follows, I'll leave it open which notion of consent<sub>m</sub> we're working with. And in the few places

---

<sup>22</sup>Although I'm not taking a stand here on how to understand coercion, those looking for a more concrete account could do worse than assuming something like Nozick's (1969) account.

<sup>23</sup>There are further important details here regarding how we evaluate this counterfactual, of course. In particular, we want to ensure that when evaluating this counterfactual we prioritize holding fixed facts regarding the subject's utility (so that their utility given each act remains unchanged) and other facts regarding their personality that bear on whether they'd consent to the act. Pete Graham has suggested that we might also consider a different notion of ideal consent<sub>m</sub> linked to a counterfactual that doesn't have the performance of the act in the antecedent. This alternative notion of ideal consent would largely yield the same results, though it would yield a different response to Slote's first argument discussed in section 5.3 (cf. footnote 33).

<sup>24</sup>For a discussion of some of the merits and demerits of appealing to something like hypothetical or ideal consent, see Thomson (1990), Stark (2000), and Enoch (2017).

where our choice between actual and ideal consent<sub>m</sub> makes a difference, I'll note how these two approaches diverge.

## 5 Objections

### 5.1 Objections to Consent Discounting Utilitarianism

Now let's turn to consider some potential objections to CDU.

One objection to CDU is that in cases where  $w = 0$ , it will permit choosing dominated acts – i.e., permit choosing an act even though there's some other act available which is strictly better for disregarded subjects, and just as good for privileged ones. And one might object that this is implausible.

CDU will have this consequence, but this is a consequence we should accept. To see this, consider a version of Parental Sacrifice in which your child would end up being equally happy with or without the gift. (While your child would greatly enjoy the gift, if you didn't buy it then they would end up finding something else to enjoy.) In this case, I take it that it's still permissible for you to work overtime, even though the option of working overtime is dominated by the option of not working overtime. And CDU yields this verdict. I.e., if we set  $v = \infty$ , then both options maximize effective utility<sub>(c,∞)</sub>:

Options	Self	Child	Effective Total ( $v = \infty$ )
Don't Work Overtime:	10	10	20
Work Overtime:	8 (10)	10	20

And so CDU will take both options to be permissible, as desired.

Likewise, consider a version of Two Parental Sacrifices in which your child would end up being equally happy with or without the gift. Again, I take it to be permissible for you to work overtime if both you and your partner are willing to make the sacrifice, even though working overtime is dominated by not working overtime. And again, CDU yields this verdict. I.e., if we set  $v = \infty$ , then both options maximize effective utility<sub>(c,∞)</sub>:

Options	Self	Partner, Child	Effective Total ( $v = \infty$ )
Don't Work Overtime:	10	10, 10	30
Work Overtime:	8 (10)	8 (10), 10	30

As before, CDU will yield the desired verdict that both options are permissible.<sup>25,26</sup>

Here is a second objection to CDU. Consider the following variant of Variable Parental Sacrifice:

Options	Self	Child	Total
Don't Work Overtime:	10	10	20
Work 1 Week of Overtime:	8	11	19
Work 2 Weeks of Overtime:	7.9	11.9	19.8

CDU will still take working one week of overtime to be permissible, since working one week of overtime maximizes effective utility<sub>(c,{2,0})</sub>. But one might worry that in this case it should be impermissible to work one week of overtime in this case, since working two weeks of overtime will only cost you an additional 0.1 utility while granting your child an additional 0.9 utility.

I take the intuition behind this worry to be this: working one week of overtime is impermissible because, given the other options available, it's an inefficient sacrifice. If you work one week of overtime you sacrifice 2 utility to give 1 utility for your child, while if you work two weeks of overtime you sacrifice 2.1 utility to give 1.9 utility for you child. And one might hold that if one is going to make a utility-decreasing sacrifice, then one's obligated to do so in as efficient manner as possible.

This worry is similar to one I addressed in section 2.3 (in footnote 9). But since it's been repeatedly raised in conversation as a worry for CDU, I'll address it again here. I think this worry is not compelling because I think it's false that permissible utility-decreasing sacrifices must be maximally efficient. That would entail that if we could work a week of overtime to buy our child a present, or could make a more efficient sacrifice by selling our leg in order to allow our child to go to college, then only selling our leg (or making no sacrifice at all) would be permissible. That's implausible.

---

<sup>25</sup>Here is another dominance-violating case that one might take to be more damning. (Thanks to Pete Graham here.) Suppose someone is drowning, and you can either rescue them (at no cost to yourself) or let them drown. If the person consents to your letting them drown, CDU would hold that it's permissible to do so. But (the objection goes) surely that's not right.

When we fill in the details regarding this case, we naturally imagine the subject as someone who is severely depressed, or is suffering from a mental illness of some kind. But such a subject would not be competent, and so could not consent<sub>m</sub> to your letting them drown. Thus CDU would not say it's permissible to let them drown. Could we fill in the details of this case in a way that makes it plausible that the subject is informed and competent? I find this very hard to do. But I find that the more I do to make it plausible that a subject could provide informed and competent consent to your letting them drown, the more plausible it becomes that it's permissible to let them drown. (In this respect this case is similar to the brainwashing case discussed below (see especially footnote 29).)

<sup>26</sup>CDU's treatment of these dominance cases entails that it will sometimes violate the kind of "Bang for your Buck" principle defended by Graham (2019), which (roughly) requires you to "make the most" out of losses inflicted on others. Thus if we consider a case where I can choose between lowering your utility by 1 to increase my utility by 1, lowering your utility by 1 to increase my utility by 2, or doing nothing, the first option would be impermissible according to the Bang for your Buck principle. But if you consent to all three options, then CDU will entail that the first option is permissible.

I think the right stance to take here is the one I described at the end of section 2.3. Namely, you should be free to determine how much of your utility you're willing to sacrifice when performing utility-decreasing altruistic acts. You can choose not to sacrifice any of your utility and not work overtime, you can choose to sacrifice a lot of utility and work two weeks of overtime, or you can choose to sacrifice no more than 2 utility, and work one week of overtime.

Here is a third objection to CDU. Consider an evil dictator who has brainwashed everyone else, so that they now consent to anything they might do. Then it seems like CDU would make it permissible for the dictator to torture everyone, since they all consent to this decrease in utility. This seems like the wrong result – it shouldn't be permissible for a dictator to do whatever they want just because they've brainwashed everyone.

Of course, everyone wouldn't consent to being brainwashed in the first place, so presumably the act of brainwashing itself wouldn't be permissible. So the evil dictator would still be doing something wrong at some point in the process. But one might think that, past events aside, the dictator is also doing something wrong *now* by torturing everyone.

How the proponent of CDU will reply to this worry depends on the notion of "brainwashing" in play. On the one hand, everyone may be "brainwashed" in the sense that they only consent to (say) being tortured by the dictator because they're not informed, competent or uncoerced. If they were informed, competent and uncoerced, they would not consent to being tortured. On this understanding of the case, the populace consents to being tortured but does not consent<sub>m</sub> to being tortured. Thus torturing everyone would be impermissible.<sup>27</sup>

On the other hand, one might construct a scenario in which everyone is "brainwashed" to consent in a manner that leaves them informed, competent and uncoerced. In that case, everyone would consent<sub>m</sub> to being tortured. This case is hard to envision – it's difficult to think of a scenario in which a populace full of informed, competent and uncoerced subjects would still consent to being tortured. But if they consent despite being informed, competent and uncoerced, then it sounds like they're not so much brainwashed as convinced.<sup>28</sup> And if this is how we understand the case, then it's not clear it's wrong for the dictator to torture everyone after all.<sup>29</sup>

---

<sup>27</sup>I take it that this is the way brainwashing is usually understood, since brainwashing is generally taken to undermine competence.

<sup>28</sup>On this understanding, we're employing the term "brainwashed" to mean something like "convinced of something that radically alters your beliefs or consenting attitudes" or "convinced to adopt beliefs or consenting attitudes that you (initially) strongly disagreed with". This does seem to be one way in which we use the term "brainwashed". Those who are convinced to enter a cult, for example, are often called brainwashed, even though in at least some cases it seems like those who enter the cult are competent, relatively informed, and so on.

<sup>29</sup>It's important here not to confuse the plausibility of there being a case in which everyone would consent to being tortured despite being informed, competent and uncoerced, with the plausibility of its being permissible for the dictator to torture them in such a case. To appropriately evaluate the latter question,

## 5.2 Objections to Variable-Disregarding Theories

Now let's consider an objection to variable-disregarding theories more generally, as a means of incorporating permissible disregarding into a Utilitarian framework.

The first objection to CDU described in section 5.1 concerned the fact that it's permissible to choose an option that's dominated by another option which is strictly better for disregarded subjects and just as good for privileged ones. I argued that this is a desirable feature of CDU – if a subject consents<sub>m</sub> to an option which leaves them with a lower utility, then *ceteris paribus* that option is permissible.

But one might worry that while this reply is plausible in some cases of disregarding, such as those involving consent, it's not plausible in others. For example, if we want to modify Utilitarianism to permit disregarding those who you're not close to, we surely don't want to permit acts that do nothing to benefit your friends and family, but make other people worse off! That is, suppose we set up a choice between acts like the following:

Options	Others	Friends and Family	Total
$a_1$ :	10	10	20
$a_2$ :	9	10	19

Surely  $a_2$  should be impermissible. If that's right, then one might worry that while variable-disregarding theories may correctly model how to permit some kinds of disregarding (e.g., oneself, those who consent), they don't correctly model how to permit other kinds of disregarding (e.g., others, those one's not close to).

I agree that in the above case  $a_2$  is impermissible. But this is not an objection to variable-disregarding theories, it's an objection to assigning disregarded subjects no weight in such cases. What distinguishes the two kinds of disregarding just described (disregarding oneself or those who consent, versus disregarding others or those one's not close to) are what weight assignments  $w$  to disregarded subjects are plausible. In the first kind of disregarding – disregarding oneself or those who consent – it's plausible to permit completely discounting the utility of the disregarded subjects and taking  $w = 0$ . In the second kind of disregarding – disregarding others or those one's not close to – it's plausible that the utility of disregarded subjects must be given at least some weight and  $w > 0$ . For while morality might allow one to *partially* disregard others or those one's not close to, it surely doesn't allow us to *completely* disregard others or those one's not close to. And as long as we give some weight to the utility of the disregarded subjects, then the kinds of dominated options described above will not be permissible.<sup>30</sup>

---

we need to first come up with a case where it seems plausible to us that everyone would consent to being tortured despite being informed, competent and uncoerced. It's only when we've managed to flesh out such a case that we can intuitively evaluate whether it's plausible in such a case for the dictator to permissibly torture them. And in my experience, the more we do to make it plausible that everyone would consent to being tortured despite being informed, competent and uncoerced, the more plausible it becomes that it's permissible for the dictator to torture them.

<sup>30</sup>For example, consider a variable-disregarding view on which the condition  $\phi$  is not being close to the

### 5.3 Slote's Objections

I've suggested that one can accommodate the intuitions behind these "altruistic" objections by modifying Utilitarianism in order to accommodate the role of consent. But Slote (1984) has argued that considerations involving consent cannot accommodate these intuitions. And the proponents of theories like SOU, such as Sider (1993) and Portmore (2008), have endorsed these arguments for why appealing to consent won't work. So let us turn to assess these arguments.

Slote (1984) offers two arguments for why consent won't suffice to explain these intuitions. He presents the first argument in the following passage:

"If someone irrationally asks me to harm or kill him, it will presumably be irrational and wrong of me to kill him, more wrong at any rate than if I irrationally choose to kill myself; yet the consent seems equal in the two cases."<sup>31</sup>

The argument seems to be this. If someone irrationally consented to my killing them, it would be wrong for me to kill them. But if I irrationally want to kill myself (and so consent to doing so), it would not be wrong (or not as wrong) for me to kill myself. So even though we both (irrationally) consent to an act that kills us, and so the status of the act with respect to consent is the same, the moral status of these acts is still different. So consent cannot explain this moral difference.

By "irrational", I take Slote to mean irrational to a degree which undermines competence. For if we understand Slote to be talking about a notion of irrationality that doesn't undermine competence – e.g., merely failing to maximize the subject's expected utility – then the claims he makes are not plausible.<sup>32</sup>

---

agent (call this condition *nc*), and suppose the weight assigned to *nc*-subjects is  $w = 0.5$ . And consider the case described above. If  $v = \{0, 0\}$  then the disregarding utilities<sub>(*nc*,0.5,{0,0})</sub> will be 20 for  $a_1$  and 19 for  $a_2$ . If  $v = \{\infty, \infty\}$  then the disregarding utilities<sub>(*nc*,0.5,{ $\infty$ , $\infty$ })</sub> will be 20 for  $a_1$  and 19.5 for  $a_2$ . But for all  $v$ ,  $a_1$  will have a higher effective utility than  $a_2$ . Thus  $a_2$  will be impermissible, as desired.

<sup>31</sup>Slote (1984), p.191.

<sup>32</sup>For example, consider a terminal cancer patient facing an extremely painful death. At times close enough to their death, it will be rational for them to choose death, since the expected utility of continuing to live will be negative. At times far enough from their death, it will be irrational to choose death, since the expected utility of continuing to live will be positive. Now consider a time just before the expected utility of death flips from positive to negative. And suppose the terminal cancer patient competently requests euthanasia at this time. This request will count as irrational, but it doesn't seem wrong to carry out their request.

Taking a step back, it's worth highlighting the strangeness of Slote's claim if we understand his talk of irrationality to just mean something like failing to maximize expected utility. We're trying to work out an account of *objective* obligation, tracking the actual consequences of our acts. Given this, it seems bizarre to maintain that the *rationality* of one's request – a subjective notion mediated by one's *beliefs* about the consequences – determines the wrongness of one's acts. To drive this home, consider two terminal cancer patients who are in identical situations, facing identical futures, who both request euthanasia. Suppose both patients are, according to the usual standards, informed and competent – perhaps they're both doctors, and have a reasonably good grasp of what their future experiences will be like. And suppose that one patient

Now let us see whether Slote's argument is compelling. First, let's assume an ideal notion of consent<sub>m</sub>. Suppose that in both cases killing the person would lower their utility, and that no one else's utility would be affected. What will CDU say about my killing someone who incompetently consents to my killing them? To work this out, we need to assess whether, if they were competent, they would have consented to my killing them. I take it that in Slote's example the subject only consents because they're not competent – if they were competent, then they wouldn't consent. If so, then CDU would say that it's impermissible for me to kill them, just as Slote says.

What will CDU say about the act of incompetently killing myself? Again, there is the question of whether, if I were competent, I would consent to my decision to commit suicide. Here the answer is trivially yes – if I were competent and decided to kill myself, I would consent to killing myself. So CDU will take it to be permissible for me to kill myself, just as Slote says. Thus a consent-based account can explain the moral difference between these two acts.<sup>33</sup>

Suppose instead that we adopt an actual, instead of ideal, notion of consent<sub>m</sub>. Given the way I've suggested we understand the notion of competence in section 4.2.1, it follows that if I'm not acting competently, then my act is not subject to moral evaluation. So if I'm not competent and I kill myself, I can't be acting wrongly, any more than a dog who killed itself could be acting wrongly. On the other hand, if someone else is not competent, and I kill them, I certainly can be acting wrongly! Again, we find that a consent-based account can explain the moral difference between these two acts.

So Slote is incorrect to think that consent-based views cannot yield the desired verdicts in this case. The moral difference he raises can straightforwardly be cashed out in terms of consent.

Let's turn to Slote's second argument. He presents his second argument in the following passage:

“If I can avoid either an enduring pain to myself or a short-lived pain to you, you and I might both agree that it would be foolish of me to prevent the shorter pain to you; judging the matter objectively, you might not consent to my taking the longer pain upon myself in order to save you from

---

is slightly more optimistic than the other, so that while living and dying have the same *expected* utility for the first patient, living has a slightly higher expected utility than dying for the second patient. According to Slote (so understood), while it could be objectively permissible to carry out the first patient's request for euthanasia, it would be objectively impermissible to carry out the second patient's request. And this is so even though they're in the exact same situation, and facing the exact same futures.

<sup>33</sup>That said, I'll confess that my intuitions here are equivocal. If you only kill yourself because you're not competent, I find (contra Slote) that the verdict that your act was wrong to be as plausible as the verdict that your act was permissible. Given this, it's interesting to note that if we adopt the alternative notion of ideal consent<sub>m</sub> suggested by Pete Graham (cf. footnote 23), CDU will hold that killing yourself in these circumstances would be impermissible after all. (That's because, given this alternative notion, you won't ideally consent<sub>m</sub> to your suicidal act, and so your utility won't be disregarded. Thus the act will be judged according to the usual Utilitarian calculus, and the act of killing yourself will be judged impermissible.)

the shorter pain. Yet there would be nothing morally wrong... in such a sacrifice. But when the positions are reversed and I can avoid a short-lived pain to myself or a longer-lived one to you and it is morally right that I should do the latter, you will presumably not consent to my doing the former and it will be wrong if I do so. Again, consent or lack of consent seems not to make the relevant common-sense moral difference."<sup>34</sup>

The argument seems to be this. If you don't consent to my performing an act  $a_1$  which would impose a large decrease in my utility to prevent a small decrease in your utility, it would still be permissible for me to do  $a_1$ . But if you don't consent to my performing an act  $b_1$  which would impose a large decrease in your utility to prevent a small decrease in my utility, it would not be permissible for me to do  $b_1$ . Since you don't consent to either, consent cannot explain the difference in moral status of these two acts.

Let us see whether this argument is compelling. Assume that these acts will have no effect on the utility of anyone else. And assume that both I and the other person in question are informed, competent and uncoerced.

What will CDU say about my performing act  $a_1$ , which brings about a large decrease in my utility to prevent a small decrease in your utility, versus act  $a_2$  in which I do nothing? From the description of the case, we can assume that I consent<sub>m</sub> to whichever act I perform, but the other person only consents<sub>m</sub> to my performing  $a_2$ . If we set  $v = \{\infty, \infty\}$  (i.e., set  $v_{\text{self}} = \infty$  and  $v_{\text{other}} = \infty$ ) this yields the result that both acts maximize effective utility<sub>(c, {∞, ∞})</sub>:

Options	Self	Other	Effective Total ( $v = \infty, \infty$ )
$a_1$ :	5 (10)	11	21
$a_2$ :	10	10 (11)	21

Thus CDU will take  $a_1$  to be permissible.

What will CDU say about my performing act  $b_1$ , which brings about a large decrease in your utility to prevent a small decrease in my utility, versus act  $b_2$  in which I do nothing? From the description of the case, we can assume that I consent<sub>m</sub> to whichever act I perform, but the other person only consents<sub>m</sub> to my performing  $b_2$ . And inverting the utilities above gives us the following values:

Options	Self	Other	Total
$b_1$ :	11	5	16
$b_2$ :	10	10	20

With a little thought we can see that  $b_2$  will maximize effective utility given any assignment  $v$ . For the disregarding utilities for  $b_2$  can range from 20-21, while the disregarding utilities for  $b_1$  will always be 16. So CDU will take  $b_1$  to be impermissible.

---

<sup>34</sup>Slote (1984), p.191.

So again, Slote is incorrect to think that consent-based views cannot yield the desired verdicts in this case. As before, the difference he raises can be straightforwardly cashed out in terms of consent.<sup>35</sup>

## 6 Conclusion

One common complaint about Utilitarianism is that it doesn't allow agents to disregard anyone. For example, as altruistic objections like Parental Sacrifice show, it doesn't allow agents to disregard their own well-being. In response to these complaints, a number of authors have endorsed replacing Utilitarianism with dual-maximizing (or multi-maximizing) views, and have endorsed handling altruistic objections by adopting something like Sider's (1993) Self-Other Utilitarianism.

I've argued that these suggestions are off-track in two respects. First, I've argued that if we want to modify Utilitarianism to permit disregarding, we should adopt variable-disregarding views, not dual-maximizing views. Second, I've argued that upon closer examination, the cases that motivate positing this kind of self-other asymmetry reveal that the morally relevant distinction is not between self and other, it's between those who do and do not consent. And by combining these two insights – adopting a variable-disregarding view which disregards those who consent – we end up with an attractive theory that yields the desired verdicts in a broad range of cases.<sup>36</sup>

## References

- Beauchamp, Tom L. and James F. Childress. 2008. *Principles of Biomedical Ethics*. Oxford University Press.
- Charland, Louis. 2011. Decision-Making Capacity. In *Stanford Encyclopedia of Philosophy*.
- Enoch, David. 2017. "Hypothetical Consent and the Value of Autonomy." *Ethics* 128(1):6–36.

---

<sup>35</sup>Portmore (2008) cites Gert (2004) as another source of arguments for why consent cannot accommodate the intuitions motivating these "altruistic" objections. Gert's arguments have a slightly different flavor, since unlike Slote, Gert's arguments have to do with subjective (not objective) obligation. But once we shift to (say) evaluating CDU with respect to consenting increases and dissenting decreases in *expected* utility, Gert's arguments take the same form as Slote's first argument, and can be handled in the same way.

<sup>36</sup>I'd like to thank Eric Chwang, Maya Eddon, Linda Eggert, Fred Feldman, Brad Hooker, David Kiloren, Hallie Liberto, Brad Skow, and the audience of the 2019 Formal Ethics conference and 2019 RoME conference, and an anonymous referee, for helpful comments and discussion. I owe a special thanks to Pete Graham for detailed comments and nearly a decade of discussions on these issues.

- Eyal, Nir. 2012. Informed Consent. In *The Stanford Encyclopedia of Philosophy*.
- Gert, Joshua. 2004. *Brute Rationality: Normativity and Human Action*. Cambridge University Press.
- Graham, Peter. 2019. "The Bang For Your Buck Principle and the Subjectivization of Moral Theories." unpublished manuscript.
- Graham, Peter A. 2010. "In Defense of Objectivism About Moral Obligation." *Ethics* 121(1).
- Grisso, Thomas and Paul S. Appelbaum. 1998. *Assessing Competence to Consent to Treatment: A Guide for Physicians and Other Health Professionals*. Oxford University Press.
- Killoren, David. 2019. "Relationship-Sensitive Utilitarianism About Animals." unpublished manuscript.
- Liberto, Hallie. 2017. "Intention and Sexual Consent." *Philosophical Explorations* 20(sup2):127–141.
- Manson, Neil C. and Onara O'Neill. 2007. *Rethinking Informed Consent in Bioethics*. Cambridge University Press.
- Meacham, Christopher J. G. 2012. "Person-Affecting Views and Saturating Counterpart Relations." *Philosophical Studies* 158:257–287.
- Nozick, Robert. 1969. Coercion. In *Philosophy, Science, and Method: Essays in Honor of Ernest Nagel*, ed. M. P. S. Suppes White Morgenbesser. St Martin's Press pp. 440–72.
- Portmore, Douglas W. 2008. "Dual-Ranking Act-Consequentialism." *Philosophical Studies* 138(3):409–427.
- Preda, Adina. 2011. "Rights Enforcement, Trade-Offs, and Pluralism." *Res Publica* 17(3):227–243.
- Roberts, Melinda. 1998. *Child versus Childmaker: Future Persons and Present Duties in Ethics and the Law*. Rowman and Littlefield.
- Scheffler, Samuel. 1994. *The Rejection of Consequentialism: A Philosophical Investigation of the Considerations Underlying Rival Moral Conceptions*. Oxford University Press.
- Sider, Theodore. 1993. "Asymmetry and Self-Sacrifice." *Philosophical Studies* 70(2):117–132.
- Slote, Michael. 1984. "Morality and Self-Other Asymmetry." *Journal of Philosophy* 81(4):179–192.

- Smart, J. J. C. and Bernard Williams. 1973. *Utilitarianism For and Against*. Cambridge University Press.
- Sobel, David. 2009. "Subjectivism and Idealization." *Ethics* 119(2):336–352.
- Splawn, Clay. 2001. "The Self-Other Asymmetry and Act Utilitarianism." *Utilitas* 13(3):323–333.
- Stark, Cynthia A. 2000. "Hypothetical Consent and Justification." *Journal of Philosophy* 97(6):313–334.
- Thomson, Judith Jarvis. 1990. *The Realm of Rights*. Harvard University Press.
- Vessel, Jean-Paul. 2010. "Supererogation for Utilitarianism." *American Philosophical Quarterly* 47(4):299–319.
- Zimmerman, Michael J. 2006. "Is Moral Obligation Objective or Subjective?" *Utilitas* 18(4):329–361.