

Supererogation for Utilitarianism

Jean-Paul Vessel
Forthcoming in *American Philosophical Quarterly*

I. Introduction

Consequentialism—the family of ethical theories sharing the characteristic that the moral status of any bit of behavior is determined by the values of the consequences of the alternatives available to a moral agent—has enjoyed tremendous support since the pioneering efforts of classical utilitarians like Jeremy Bentham and John Stuart Mill. Continually subjected to theoretical and practical attacks, consequentialism has withstood centuries of serious criticism. Recent critics, however—especially those influenced by contemporary schools of virtue ethics and those attracted to “satisficing” accounts of morality—have lodged a new class of objections against the venerable position. These new objections seem to be motivated in the most part by “maximizing” properties of the paradigm consequentialist theory: Classical Utilitarianism. Classical utilitarianism requires that for any situation in which we find ourselves, we perform a best possible alternative available to us at the time, where a best possible alternative is one than which no other alternative is better.¹

According to classical utilitarianism we are morally *obligated* to engage in a best possible plan of action accessible to us, regardless of how much personal sacrifice such a plan would entail.² Critics claim that this maximizing feature of classical utilitarianism—the moral requirement to maximize the amount of value in the world—is unacceptable. Classical utilitarianism has been charged with being too demanding, too difficult for moral agents with naturally selfish inclinations to adopt.³ More specifically, critics charge that utilitarianism is incapable of exhibiting the allegedly important phenomenon of supererogation: In virtue of the

fact that utilitarianism requires that we always bring about the best, it appears impossible to do anything that is “above and beyond” the call of moral duty.⁴ Supererogatory action is impossible under a utilitarian scheme—or at least so say such critics; call them ‘supererogation critics’. Many believe that supererogation critics have dealt a death blow to traditional formulations of consequentialism.

This essay is a response to the objection from supererogation against classically-spirited consequentialist moral theories. First, it is argued that—despite the popularity of the supererogation objections—classical act utilitarianism does exhibit supererogatory properties. It is argued that the “ties at the top” phenomenon entails that many acts in many different kinds of cases are deemed supererogatory on the classical utilitarian scheme before recognizing that supererogation critics will be unsatisfied with the illustration: They will continue to clamor for a richer account of supererogation than can be generated by the classical theory.

Next, a “satisficing” form of consequentialism introduced to absorb the supererogation objection is considered. Anything requiring more personal sacrifice than that which is “good enough” (or satisfactory) for the world is designated as supererogatory on satisficing schemes. Drawing on recent work, it is argued that many satisficing theorists have failed in their attempts to make the notion of “good enough” clear enough for serious theoretical uses. Furthermore, Ben Bradley’s objection that in many cases satisficing theories implausibly permit us to prevent a significant amount of good from obtaining in the world is investigated.

Next, satisficers who wish to remain consequentially principled and consistent are offered a cleaner consequentialist view called ‘Egoistically-Adjusted Act Utilitarianism’ (EAU), a view that exhibits the very supererogatory features that satisficing theorists find so attractive. After the construction of EAU and the elucidation its properties, an impartiality objection is lodged against it.

In response to supererogation critics who take seriously impartiality objections, the alleged “self-other asymmetry” in moral theorizing is introduced. A version of utilitarianism formulated by Ted Sider—Self-Other Utilitarianism (SOU)—that incorporates this phenomenon is presented before its supererogatory properties are exhibited. Next, another novel moral view is presented: a version of Egoistically-Adjusted Act Utilitarianism that incorporates the self-other asymmetry. Finally, a form of social proximity utilitarianism is constructed in response to the “nearest and dearest” objection against classical utilitarianism. These latter two views—it is argued—exhibit very rich supererogatory properties, properties that should certainly satisfy the bulk of supererogation critics.

Now onto a simple formulation of classical act utilitarianism, some assumptions, and a formal introduction of the charges that the classical view is both too demanding and fails to exhibit supererogatory properties.

II. Classical Act Utilitarianism and the Supererogation Objection

The first assumption is that every morally relevant alternative (or act token) has a certain hedonic utility. Let the *hedonic utility of an alternative, A*, be the result of subtracting the total amount of pain that *A* would cause from the total amount of pleasure that *A* would cause, were *A* performed.⁵ A simple version of the classical view can then be formulated in this way:

AU: An alternative, *A*, is morally right iff no alternative to *A* has a higher hedonic utility than *A* has.

AU—the paradigm, classical version of act utilitarianism—is clearly a maximizing theory: AU requires that we maximize the amount of hedonic value in the world, that we refrain from ever performing an act token such that some alternative to it has a higher hedonic utility than it has. In this simple case, we can see the kinds of implications AU generates:

al HU
 100

<i>a2</i>	90
<i>a3</i>	70
<i>a4</i>	55

Here four alternatives are available, each of which would benefit the world if performed. AU, however, requires that *a1* be performed; it alone maximizes hedonic utility. It is impossible to do more for the world—from a hedono-doloric perspective—than to perform *a1*. Cases like this one have generated discontent in the minds of many contemporary theorists. The elucidation of two further concepts will enable us to understand clearly how this case sparks the “Too Demandingness” and “Lack of Supererogation” objections against AU.

There is a certain hedono-doloric impact that any morally relevant alternative would have upon the agent of the act, were it performed. Let the *hedonic utility of an alternative, A, for agent* be the result of subtracting the total amount of pain that *A* would cause *for the agent of A* from the total amount of pleasure that *A* would cause *for the agent of A*, were *A* performed.

There is also a certain hedono-doloric impact that any morally relevant alternative would have upon everyone other than the agent of the act, were it performed. Let the *hedonic utility of an alternative, A, for others* be the result of subtracting the total amount of pain that *A* would cause *for those other than the agent of A* from the total amount of pleasure that *A* would cause *for those other than the agent of A*, were *A* performed.

With these three concepts of hedonic utility in hand, the example can be fleshed out in more detail. Consider this chart:⁶

	<u>HU</u>	<u>HU-for-Agent</u>	<u>HU-for-Others</u>
<i>a1</i>	100	-10	110
<i>a2</i>	90	20	70
<i>a3</i>	70	0	70
<i>a4</i>	55	25	30

AU, of course, is concerned only with the first column, igniting the critics’ responses. Despite the fact that *a1* maximizes hedonic utility—that it ushers into the world the greatest possible

balance of pleasure over pain—critics have charged that no one is morally obligated to perform such an act: It is simply too demanding, especially considering the possibility of performing *a2*. If anything—critics claim—*a1* should be considered supererogatory. The agent incurs a serious personal sacrifice with the performance of *a1* so that others will prosper. It has been suggested that in doing *a1*, the agent goes above and beyond the call of moral duty. *a2*—critics clamor—is also morally permissible. Both the agent and others prosper to a serious extent as a result of the performance of *a2*. Because AU requires that a maximizing alternative be performed (*a1*), critics conclude that AU is too demanding and fails to exhibit rich supererogatory properties.

AU is a very demanding theory, requiring that we do the best that we can for the world from a hedono-doloric perspective. Whether or not it is, in fact, *too* demanding has yet to be conclusively argued in this author’s opinion, but a return to this topic will transpire soon enough. Note, however, that while AU fails to generate the result that any alternative in this current case is morally supererogatory, it does have the resources to generate such implications in a host of other somewhat similar types of case.

In many cases, the “ties at the top” phenomenon emerges. Consider this case:

	<u>HU</u>
<i>a5</i>	100
<i>a6</i>	100
<i>a7</i>	70
<i>a8</i>	55

This case differs from our original case in that it contains two maximizing alternatives: *a5* and *a6*, each of which is morally permissible on the classical scheme. (Note that in this situation, no alternative is morally obligatory according to AU.) With the insertion of the additional utility categories, how an alternative might be deemed supererogatory under AU can be easily understood.

	<u>HU</u>	<u>HU-for-Agent</u>	<u>HU-for-Others</u>
<i>a5</i>	100	-10	110

<i>a6</i>	100	20	80
<i>a7</i>	70	0	70
<i>a8</i>	55	25	30

a5 requires substantial personal sacrifice from the agent so that others will be much better off. *a6*, on the other hand, benefits the agent to a considerable degree but also benefits others a great deal. Again, each is morally permissible according to AU. Thus, it can be seen how “ties at the top” cases generate supererogatory action for AU. *a5* is supererogatory under AU. In doing *a5*, the agent goes above and beyond the call of moral duty in efforts to benefit others, for she is permitted by AU to perform the self-benefiting *a6*.⁷ Whether extra praiseworthiness should be bestowed upon an agent who knowingly performs alternatives like *a5* is a matter for another day. What seems clear, though, is this: AU generates the result that in many “ties at the top” cases supererogatory action is possible. At least in these cases, AU does not appear to be quite as demanding as many critics suggest.

Some might find this evaluation of the case a bit too cavalier, suggesting that despite the argumentation given, AU remains incapable of generating the implication that some alternatives are supererogatory. This sort of objection rests upon the controversial conceptions of supererogatory action itself. Most concur that supererogatory action must be morally optional: neither morally obligatory nor morally forbidden.⁸ Furthermore, supererogatory action must be in some way especially valuable, more valuable in this way than some competing morally permissible alternative. The sense in which supererogatory action must be more valuable than a competing morally permissible alternative, however, is a matter of rich controversy. Some believe that supererogatory action must be morally better than a competing permissible alternative.⁹ Some believe that the performance of supererogatory action confers more moral praiseworthiness upon its agent than would the performance of a morally permissible competitor.¹⁰ Still others believe that supererogatory action must benefit others to a greater

extent than a morally permissible alternative.¹¹ This author is inclined to endorse this latter analytic requirement on the term ‘supererogatory’.¹²

Classical act utilitarianism cannot generate the implication that *a5* is morally better than *a6*: They are identical in moral value under AU. But *a5* possesses interesting properties, leading to the author’s endorsement of it as supererogatory on a classical utilitarian scheme. In performing the morally optional *a5*, the agent of *a5* incurs a serious personal sacrifice so that others will prosper. There’s a sense in which such action is especially valuable.

Consequentialist moral theories frequently require that we help (or benefit) others to a certain extent. When an agent engages in a permissible line of action that exceeds the extent to which others must be benefited (perhaps at some cost to herself), one should be inclined to label such action ‘supererogatory’. Thankfulness is the appropriate attitude of recipients who benefit through the graciousness of supererogatory action (provided that such beneficiaries are in the know). Supererogatory action generates a certain kind of praiseworthiness: Those who engage in such action are wholly worthy of the praises of those whom they are benefiting. Supererogation seems possible on the classical scheme.¹³

Critics will respond, however, that while AU might exhibit the phenomenon of supererogation, a richer notion of supererogation is required, one that will generate the result that *a2* in our first case is morally permissible. It is only in “ties at the top” cases that supererogation is possible on the classical scheme. Supererogation (and “demandingness”) critics claim that there are many more cases of supererogation than those containing ties at the top. Some have chosen to abandon the consequentialist family of moral theories for this reason; others, however, have developed mutations of the classical theory capable of accommodating these supererogation intuitions.

III. Satisficing and Egoistic Adjustments

	<u>HU</u>	<u>HU-for-Agent</u>	<u>HU-for-Others</u>
<i>a1</i>	100	-10	110
<i>a2</i>	90	20	70
<i>a3</i>	70	0	70
<i>a4</i>	55	25	30

Satisficers and egoistic adjusters have put forth consequentialist theories designed to generate results consonant with the intuitions of supererogation critics. Satisficing versions of consequentialism do not require that we maximize value in the world; rather, we are merely obligated to do that which is “good enough” for the world.¹⁴ Anything requiring more personal sacrifice than that which is good enough for the world is designated as supererogatory on satisficing schemes. A generic form of satisficing act utilitarianism can be formulated in this way:

SatU: An alternative, *A*, is morally right iff *A* has a hedonic utility that is “good enough”.

How SatU might assign normative values to the alternatives in our case—and generate the desired supererogation results—can now be explored. Without a clean explication of the phrase “good enough,” it is not exactly clear how the normative statuses of the alternatives in our case are generated by SatU. Intuitively, it might be thought that SatU has these implications: *a1* should come out as morally permissible (but not morally obligatory!) according to SatU. It is impossible to do more for the world than to perform *a1*; *a1* is the sole maximizing alternative. Because it is the best thing one can do for the world, it must have a hedonic utility that is good enough for the world. *a2* is the interesting alternative. It seems that SatU implies that *a2*, too, is morally permissible; heck, is that not the reason why theories like SatU were introduced? From a certain perspective, *a2* does seem to qualify as “good enough”: It is almost as good for the

world as the maximizing *a1* and is substantially better for the world than the remaining alternatives.¹⁵

On the assumption that the hedonic utility of *a2* satisfies the “good enough” requirement, the supererogatory properties of SatU become apparent: *a1* is deemed supererogatory in light of the facts that (i) it is morally permissible, (ii) it requires more personal sacrifice than another morally permissible alternative, and (iii) it benefits others to a greater degree than another morally permissible alternative.¹⁶ SatU exhibits a very rich notion of supererogation. It generates the same “ties at the top” implications as AU. Furthermore, it designates a slew of alternatives as supererogatory in many cases where a non-maximizing alternative meets the “good enough” requirement.

Satisficing forms of consequentialism might satiate the supererogatory cravings of some, but satisficing views face serious theoretical difficulties. First, there has yet to be explicated a concept of “good enough” that is crisp enough for interesting yet plausible theoretical use.¹⁷ Perhaps more importantly, extant satisficing views generate repugnant implications in a variety of cases. Assume—for the sake of discussion—that some interesting explication of “good enough” is available. Now consider this case:

	<u>HU</u>	<u>HU-for-Agent</u>	<u>HU-for-Others</u>
<i>a9</i>	100	-10	110
<i>a10</i>	90	20	70
<i>a11</i>	100	5	95
<i>a12</i>	110	5	105

Assume, just as in the last example, that the utility of *a10* is good enough, generating the result that *a10* is morally permissible according to SatU.¹⁸ *a9*—just as before—is deemed morally permissible but not obligatory by SatU. In this case, the problems are generated by SatU’s implications regarding *a11* and *a12*. If *a10* is good enough, then clearly *a11* and *a12* are good enough, and thus permissible on this satisficing scheme. But the designation of *a11* as

permissible is anathema to the consequentialist spirit. It costs the agent *nothing* to bring about the greater benefits for others contained in the consequence of *a12*. SatU nonetheless permits the agent to engage in this “gratuitous prevention of goodness” for the world.¹⁹ This appears to be a devastating objection to theories like SatU. Moreover, Ben Bradley has illustrated how this objection and correlating analogues of it can be successful lodged against the gamut of satisficing accounts of consequentialism. Despite its supererogatory properties and contemporary popularity, satisficing versions of consequentialism simply do not seem well suited to do the requisite consequentialist work.

This author suggests that satisficing theorists who wish to remain consequentially principled and consistent should adopt a cleaner consequentialist view called ‘Egoistically-Adjusted Utilitarianism’ (EAU), one that utilizes a novel ranking system rather than simply a maximizing scheme and generates the supererogation results that satisficers crave.²⁰ The key element distinguishing EAU from classical theories like AU is its egoistic weighting. Similar to motivations underlying satisficing theories, EAU will not require the maximizing of value in all cases; instead, it will merely require that agents perform personal sacrifices that qualify as “good enough” for the world in many cases.²¹ This notion of “good enough” can be articulated in a fairly clear fashion. Reconsider the case used to motivate satisficing maneuvering:

	<u>HU</u>	<u>HU-for-Agent</u>	<u>HU-for-Others</u>
<i>a1</i>	100	-10	110
<i>a2</i>	90	20	70
<i>a3</i>	70	0	70
<i>a4</i>	55	25	30

The egoistic intuition here is that we can count our own possible future goods and evils as more important—more weighty—from a moral perspective than those of others. Suppose that we can count our own possible future goods and evils as twice as important as those of others when

considering what morality requires of us.²² Given this supposition, a new utility category is added to the mix: The *egoistically-adjusted hedonic utility for the agent* of an alternative.

	<u>HU</u>	<u>HU-for-Agent</u>	<u>E-Adj HU-for-Agent</u>	<u>HU-for-Others</u>
<i>a1</i>	100	-10	-20	110
<i>a2</i>	90	20	40	70
<i>a3</i>	70	0	0	70
<i>a4</i>	55	25	50	30

With this egoistic adjustment in play, the concept of the overall egoistically-adjusted hedonic utility of an alternative can be constructed, providing the framework for this new egoistically injected mutation of classical utilitarianism. Let the *egoistically-adjusted hedonic utility of an alternative, A*, be the result of adding the egoistically-adjusted hedonic utility for the agent of *A* to the hedonic utility for others of *A*. Here is the chart resulting from the addition of the new category:

	<u>HU</u>	<u>HU-for-Agent</u>	<u>E-Adj HU-for-Agent</u>	<u>HU-for-Others</u>	<u>E-Adj HU</u>
<i>a1</i>	100	-10	-20	110	90
<i>a2</i>	90	20	40	70	110
<i>a3</i>	70	0	0	70	70
<i>a4</i>	55	25	50	30	80

A theory can be constructed that requires maximizing egoistically-adjusted hedonic utility—but such a theory would veer us farther from the classical theory than need be: That kind of view would prohibit the performance of *a1*, rather than merely elevate *a2* into a position of moral permissibility. Consider, instead, this non-maximizing method for ranking any two alternatives: Let it be the case that *alternative A E-outranks alternative B* iff (i) $HU(A) > HU(B)$ and (ii) $E-Adj HU(A) > E-Adj HU(B)$.²³

According to this method of comparing alternatives, one alternative is E-outranked by another iff it has both a lower hedonic utility and a lower egoistically-adjusted hedonic utility

than the other. With this method of ranking at our disposal, a full blown version of egoistically-adjusted act utilitarianism can be clearly stated.

EAU: An alternative, *A*, is morally right iff no alternative to *A* E-outranks *A*.²⁴

Abandoning a maximizing scheme for this kind of ranking (or outranking) system preserves a conceptual closeness to the classical theory. The implications of EAU in our case are obvious. *a1* transforms into a supererogatory act—just as SatU implied. (Note that according to EAU, one is always permitted to perform an alternative that maximizes hedonic utility: Any alternative that is morally permitted by the classical view is permitted by EAU.) *a1* is deemed supererogatory by the new view because it demands more personal sacrifice for others than *a2*, which is deemed morally permissible on this scheme. Moreover, EAU does not fall prey to the objection that it permits the gratuitous prevention of goodness in the world in the same way that satisficing views do.²⁵ But despite its many attractive features (clarity, supererogatory properties, conceptual closeness to AU, conceptual superiority to SatU), serious objections can be lodged against EAU.

Perhaps the most serious is the classical objection from impartiality forcefully presented by J.S. Mill in his *Utilitarianism*:

I must again repeat, what the assailants of utilitarianism seldom have the justice to acknowledge, that the happiness which forms the utilitarian standard of what is right in conduct, is not the agent's own happiness, but that of all concerned. As between his own happiness and that of others utilitarianism requires him to be as strictly impartial as a disinterested and benevolent spectator.

Mill suggests here that agent impartiality is an essential property of any plausible moral theory. Mill and others moved by the impartiality objection are likely to deem impartiality as a much more attractive property than the kind of supererogation generated by a theory like EAU.

Theorists attracted to EAU, however, have some resources with which to respond to the impartiality objection. They can argue that it's universalizability—not necessarily impartiality—that is required by acceptable normative theories. Perhaps it can be argued that EAU is universalizable on a Kantian scheme, that rational agents can consistently will that everyone always act in accordance with EAU. Some might argue that the principle is one that rational planners would unanimously select from behind the veil of ignorance, satisfying Rawlsian universalizability requirements. (Note also that similar resources might be utilized to argue for a particular non-arbitrary egoistic weighting.) Regardless of who stands the higher ground in this impartiality debate, the following remains clear: While some satisficing theorists might find EAU attractive, no consequentialist moved by this impartiality objection will, leaving some still dissatisfied with the fact that impartial consequentialist theories seem to lack a rich notion of supererogation. More supererogation, though, is conceptually nearby.

IV. The Objection from Selfishness, Altruism, and the Self-Other Asymmetry

That many are prepared to reject classical utilitarianism on the ground that it is too demanding has already been shown. Others, however, have argued that the classical view is simply too selfish—and should be rejected for that reason. Consider this case:

	<u>HU</u>	<u>HU-for-Agent</u>	<u>HU-for-Others</u>
<i>a13</i>	100	100	0
<i>a14</i>	90	20	70
<i>a15</i>	70	0	70
<i>a16</i>	55	25	30

Classical AU requires that *a13* be performed; it is the sole maximizer of hedonic utility.

Nonetheless, many believe that *a14* should be deemed morally permissible²⁶—or perhaps even morally obligatory. Sure, *a13* brings about the best for the world, but all of the benefits are incurred by the agent of the act. Shouldn't the agent be permitted to forgo a portion of those

benefits (80) so that others will prosper (70)? Many theorists are inclined to answer this question in the affirmative.

An altruistic mutation of classical utilitarianism might be considered an interesting theoretical response to this objection from selfishness. In contrast to the egoistic adjustment strategy considered in the previous section, the altruistic adjustment is intended to capture the intuition that the possible future goods and evils experienced by people other than the agent of some act are considered more weighty from a moral perspective than those experienced by the agent herself.²⁷ Suppose—at least on this view—that the possible future goods and evils of those other than the agent of some act are considered twice as important as those of the agent when pondering what morality requires. Given this supposition, a new utility category can be considered: The *altruistically-adjusted hedonic utility for others* of an alternative.

	<u>HU</u>	<u>HU-for-Agent</u>	<u>HU-for-Others</u>	<u>A-Adj HU-for-Others</u>
<i>a13</i>	100	100	0	0
<i>a14</i>	90	20	70	140
<i>a15</i>	70	0	70	140
<i>a16</i>	55	25	30	60

With this idea of an altruistic adjustment for others in hand, the concept of the overall altruistically-adjusted utility of an alternative can be explicated. Let the *altruistically-adjusted hedonic utility of an alternative, A*, be the result of adding the hedonic utility for the agent of *A* to the altruistically-adjusted hedonic utility for others of *A*. Consider this chart containing the new category:

	<u>HU</u>	<u>HU-for-Agent</u>	<u>HU-for-Others</u>	<u>A-Adj HU-for-Others</u>	<u>A-Adj HU</u>
<i>a13</i>	100	100	0	0	100
<i>a14</i>	90	20	70	140	160
<i>a15</i>	70	0	70	140	140
<i>a16</i>	55	25	30	60	85

Now consider *Altruistically-Adjusted Act Utilitarianism (AAU)*:

AAU: An alternative, *A*, is morally right iff no alternative to *A* has a higher altruistically-adjusted hedonic utility than *A* has.

a13 is forbidden by AAU, constituting a consequentialist response to the objection from selfishness. *a14* is declared morally obligatory by the view, marking the altruistic nature of the theory, something that is deemed integral to many Christian conceptions of morality. But regardless of its response to the selfishness objection, it is unlikely that many (former) classical utilitarians will be moved to embrace AAU.

AAU violates Mill's egalitarian impartiality requirement: AAU requires us to sacrifice ourselves to others in a more fundamental way than does classical utilitarianism, which in itself is considered far too altruistic by many theorists, motivating the "too demandingness" objection considered in the previous sections. Note just how demanding AAU is.²⁸ AAU doesn't merely elevate the status of *a14* to moral permissibility; it requires that *a14* be performed despite the fact that *a13* ushers the greatest possible amount of value into the world. According to AAU, others are morally justified in demanding that we sacrifice ourselves to them solely in virtue of the fact that they are others. C. D. Broad—while perhaps embracing the selfishness objection against classical utilitarianism—intimates that AAU simply goes too far: "[C]ommon-sense holds that it may be right and praiseworthy for a person voluntarily to make sacrifices which it would be wrong for anyone else to impose on him."²⁹ Jean Hampton echoes these sentiments in her "Selflessness and the Loss of Self":

Hence, I do not agree that all community-benefiting, other-regarding actions are morally required, and I believe that the advocates of the concept of supererogation are right to maintain that we are sometimes morally permitted not to choose an altruistic, self-sacrificing act, but to act, instead, to benefit ourselves....To be "impartial" from a moral point of view does not always mean excluding oneself and one's own needs from moral deliberation. To treat all people equally does not mean giving everyone but oneself equal

concern. Moral people do not put themselves to one side; they include themselves in the calculation and give themselves weight in the determination of the right action to take.

Might there be some other theoretical motivation to which a classically-minded consequentialist can appeal in constructing a response to this objection from selfishness? Some philosophers have suggested that an appeal to the “self-other” asymmetry in common sense morality is capable of doing the requisite work. C. D. Broad suggests such a route in his “Self and Others.”

Suppose that an act will affect a certain person B and him alone. Then there will be a characteristic dissimilarity in the act according to whether it is done by B himself or by any other person. If it is done by B, it will be a self-affecting act; if it is done by any other person, it will be an other-affecting act. Now this kind of dissimilarity between acts, though it depends merely on the numerical identity or the numerical otherness of the agent-self and the patient-self, may be ethically relevant. If the agent-self and the patient-self be the same, the act may be right; if they are different, it may be indifferent or positively wrong. And the converse may be equally true. It is misleading to compare an act which is only self-affecting with one which would be other-affecting, however alike they may be in their consequences and in all other respects. For this dissimilarity may be ethically relevant. Undoubtedly common sense thinks that it is often highly relevant. To give to oneself an innocent pleasure is generally regarded as morally indifferent. To give to another a similar pleasure may be regarded as praiseworthy or even as obligatory.³⁰

Michael Slote also defends the (alleged) self-other asymmetry in moral thinking:

This further, and ultimately, I believe, highly perplexing element of common-sense morality consists in the permissibility, according to ordinary thinking, of *not* benefiting

oneself and of favoring *other people* even when this leads to less than optimal results. If I have a choice between conferring a great benefit on myself or a lesser benefit on someone else, and these are the only relevant factors in the situation, common-sense morality tells us that it is quite permissible to sacrifice one's own greater benefit to the lesser benefit of another. In the absence of some special relation or obligation to that other, common sense might concede that it was irrational, stupid, or gratuitous to do so, but surely not that it was morally wrong. Similarly, in a situation where no one else is concerned (or even, if you will, where no one else exists) if I ignore an opportunity to enjoy a pleasure or do not bother to avoid a pain, then (other things being equal) I do wrong by consequentialist standards, but, again, not by ordinary standards. Thus ordinary moral thinking seems to involve an asymmetry regarding what an agent is permitted to do to himself and what he is permitted to do to others.³¹

Some—including Slote himself—have thought that traditional formulations of utilitarianism are incapable of incorporating this self-other asymmetry, spelling doom for consequentialist projects. It is true that classical AU lacks the resources to incorporate this allegedly important phenomenon, but Theodore Sider illustrates a way by which consequentialists can incorporate this self-other asymmetry into their classical, normative frameworks. The view that Sider formulates—Self-Other Utilitarianism (SOU)—is a close cousin to the classical AU. Like Egoistically-Adjusted Utilitarianism (EAU), it too is built upon a non-maximizing ranking system. Consider this method of ranking any two alternatives:

Alternative A SO-outranks alternative B iff (i) $HU(A) > HU(B)$ and (ii) $HU\text{-for-Others}(A) > HU\text{-for-Others}(B)$.

According to this method of comparing alternatives, one alternative is SO-outranked by another iff it has both a lower hedonic utility and a lower hedonic utility for others than the other. SOU can now be clearly stated:

SOU: An alternative, *A*, is morally right iff no alternative to *A* SO-outranks *A*.³²

Returning to the latest case, it can be seen how SOU answers the “too selfishness” objection and generates an interesting type of supererogation.

	<u>HU</u>	<u>HU-for-Agent</u>	<u>HU-for-Others</u>
<i>a13</i>	100	100	0
<i>a14</i>	90	20	70
<i>a15</i>	70	0	70
<i>a16</i>	55	25	30

Like EAU, SOU deems permissible any alternative that maximizes hedonic utility; *a13* comes out as morally permissible on this scheme. Note also that SOU does not require that hedonic utility be maximized in all cases, particularly those in which the maximizing alternative has a lower hedonic utility for others than some other available alternative. Such is the case in this example. While *a13* maximizes hedonic utility, it does nothing for others. SOU deems *a14* to be morally permissible; it is not SO-outranked by the maximizing *a13* or any other alternative: No available alternative has a higher hedonic utility for others than it has. The selfishness objection is answered: While it is permissible to perform *a13* according to SOU, it is not required, generating the possibility for action that will benefit others to a serious extent.

Interesting supererogatory features become evident as well. SOU bestows supererogatory status upon *a14*: By performing *a14* (rather than *a13* or the perhaps irrational *a15*), the agent sacrifices 80 units of hedonic value that she is morally permitted to acquire so that others will prosper. Clearly, this is consonant with going above and beyond the call of moral duty.

Note also that cases like the one above are extremely common. We frequently find ourselves in position both to benefit ourselves and maximize hedonic utility simultaneously. We

are so close to ourselves; we know ourselves so well; we are almost always in a position to benefit ourselves. It is no surprise that strict forms of egoism and classical forms of utilitarianism generate identical implications in a wide range of cases. Possibilities for supererogatory action are abundant according to SOU.^{33,34}

Again, it must be noted that it is unclear as to whether extra praiseworthiness should be bestowed upon agents who knowingly engage in alternatives deemed supererogatory by SOU. Many classically-minded utilitarians would probably think that it should not. Regardless, SOU exhibits rich supererogatory properties and should thus satisfy a vast number of supererogation critics, especially those inclined to reject EAU for impartiality reasons. But what can be done for those who find attractive both the supererogatory properties of EAU and SOU? A combination is in order.

V. Combining the Supererogatory Properties of EAU and SOU

Super supererogation for a view somewhat similar to the classical AU can now be trotted out. Consider this combinatory method for ranking any two alternatives:

Alternative A E&SO-outranks alternative B iff (i) E-Adj HU(A) > E-Adj HU(B) and (ii) HU-for-Others(A) > HU-for-Others(B).

The resulting egoistically-adjusted self-other version of utilitarianism can be stated in this way:

EASOU: An alternative, A, is morally right iff no alternative to A E&SO-outranks A.³⁵

The supererogatory properties of EASOU are rich indeed and very evident, combining the supererogatory properties of EAU and SOU while maintaining some affinity to classical utilitarianism: Anything morally permitted by AU is morally permitted by EASOU.³⁶ Such properties—it seems—should satisfy the bulk of critics whose supererogation cravings have yet to be satiated. Nonetheless, consideration of a consequentialist response to the “nearest and

dearest” objection might illuminate a final avenue by which supererogation might emerge in an (at least somewhat) classically-spirited form of consequentialism.

VI. The “Nearest and Dearest” Objection and Some Final Supererogation Possibilities

Classical utilitarianism—as has been emphasized—is an impartial theory. It implies that no one’s future goods and evils are more weighty from a moral perspective than anyone else’s, not the agent’s, nor those other than the agent—including those to whom the agent is most “near and dear”: her spouse, children, parents, siblings, closest friends, etc. Sidgwick emphasizes this as forcefully and clearly as does Mill.

For Utilitarianism is sometimes said to resolve all virtue into universal and impartial Benevolence: it does not, however, prescribe that we should love all men equally, but that we should aim at Happiness generally as our ultimate end, and so consider the happiness of any one individual as equally important with the equal happiness of any other, as an element of this total; and should distribute our kindness so as to make this total as great as possible, in whatever way this result may be attained.³⁷

If maximizing hedonic utility requires that one benefit distant strangers rather than one’s own children—with, for example, medical treatments, education, or exciting activities—then AU requires that those distant strangers be benefited at the expense of one’s dearly loved children.

Many theorists find such implications to be very counterintuitive. Derek Parfit writes:

Most of us believe that there are certain people to whose interests we should give extra weight. Thus each ought to give priority to his children, parents, pupils, patients, members of his own trade union, those whom he represents, or fellowcountrymen. This priority should not be absolute. It would be wrong to save my child’s toy rather than a stranger’s life. But I ought to save my child from harm rather than save a stranger’s child

from a somewhat greater harm. I have special duties to my child, which cannot be overridden simply because I could do somewhat greater good elsewhere.³⁸

Describing a scenario in which you can save either your own child or two other children in a boating accident, Diane Jeske and Richard Fumerton write:

What should you do? A great many of us will conclude that it would certainly be morally permissible for you to put your child's life first and some of us would conclude that you are *obligated* to give your child priority.

So, for example, we would have no moral qualms about helping a friend or a colleague even if by doing so we forfeit the chance to help several other people whom we don't know.³⁹

C. D. Broad expresses similar sentiments in his "Self and Others". Focussing upon the family of impartial ethical theories (which includes AU as a paradigm member), Broad writes:

And among acts which are primarily other-affecting it denies any direct ethical relevance to the difference between more and less intimate relationships between an agent and his possible beneficiaries. Yet *prima facie* the special urgency of the claims of certain others upon one's beneficence seems to be founded directly on certain special relationships of those others to oneself.

Aristotle grapples with these matters in the second section of Book IX of *Nicomachean Ethics* (roughly 1164b:20 to near the end of 1165a). Wondering "whether one should in all things give the preference to one's father and obey him" and whether "one should render a service by preference to a friend or to a good man, and should show gratitude to a benefactor or oblige a friend, if one cannot do both," Aristotle suggests that we should "help our parents before all others" (when it comes to sustenance),

[W]hile to comrades and brothers one should allow freedom of speech and common use of all things. To kinsmen, too, and fellow-tribesmen and fellow-citizens and to every other class one should always try to assign what is appropriate, and to compare the claims of each class with respect to nearness of relation and to virtue or usefulness. The comparison is easier when the persons belong to the same class, and more laborious when they are different. Yet we must not on *that* account shrink from the task, but decide the question as best we can.⁴⁰

Sidgwick responded to these objections by pointing out that it is *generally* the case that focussing our moral attention primarily upon family, friends, and those to whom we are closer in some way (from a social perspective) tends to maximize hedonic utility.⁴¹ Jeske and Fumerton articulate Sidgwick's strategy.

A parent is, after all, typically in the best position (epistemically and causally) to benefit her child. We know more about our own community than we do about others and if we are trying to increase the general happiness we might do best by concentrating our efforts where we are in the best position to judge their effects.⁴²

Despite the fact that it is probably contingently true that we (or at least most of us⁴³) are most likely to succeed (or approach success) in accomplishing utilitarian goals by focussing upon those to whom we stand in close social proximity, "nearest and dearest" objectors will unlikely be moved by Sidgwick's strategy. Such objectors push for a theoretically justified moral preference (as opposed to a merely contingently justified moral preference) for those who are most near and dear to us, something that Sidgwick's response is incapable of providing. Fred Feldman imagines a consequentialist theory perhaps capable of incorporating these objectors' intuitions.

We can imagine a form of “social proximity utilitarianism”. The view is a maximizing theory much like ordinary act utilitarianism, but the system for calculating utilities is modified. We count the pleasures and pains of those socially near to us more heavily than we do the pleasures and pains of those socially more distant. The pleasures and pains of my immediate family count very heavily; those of distant starving children in Bangladesh hardly at all. The theory is designed to reflect intuitions such as the intuition that other things being equal, I have a greater moral obligation to guard the welfare of my own child than I have to guard the welfare of some unknown, distant child.⁴⁴

This consequentialist theory Feldman imagines—Social Proximity Utilitarianism—can be formulated clearly.⁴⁵ Consider this numerical measure of social proximity between act tokens and the hedono-doloric episodes that they would cause were they performed. Let the social proximity of some act, *A*, to some hedono-doloric episode, *E*, contained within the consequence of *A* be identical to the social proximity of the agent of *A* to the person who experiences *E*. Let the social proximity between any two people be represented by a number between one and zero, with the number one indicating maximal social proximity and zero indicating minimal social proximity. The social proximity of a loving father to his loving daughter, for example, is very high, close to a rating of one, while the social proximity of any earthling to some rational extraterrestrial being is likely zero. Intermediate numbers are assigned to intermediate cases.

Before considering how hedonic utilities might be adjusted for social proximity in the construction of Social Proximity Utilitarianism, consideration of a specific “nearest and dearest” case might prove helpful. Imagine that a single mother with limited means comes across the possibility of providing a serious benefit for her young daughter, perhaps a medical treatment, a stimulating educational opportunity, an exciting adventure, or maybe even a bicycle. Imagine further that just before procuring the benefit, the mother is approached by a classically-minded

utilitarian (also of very limited means) who informs her that those very resources could be better utilized in assisting a young girl (completely unknown to her) living in a severely economically disadvantaged, rural portion of their state. Suppose that the following chart accurately reflects the nature of this case.

	<u>HU</u>
<i>a17</i> : benefiting one's daughter	35
<i>a18</i> : benefiting some rural strangers' daughter	50

It can be imagined that the hedonic utilities of *a17* and *a18* are generated in similar ways.

Suppose that there are two episodes of pleasure contained in the consequence of *a17*: the daughter enjoying a large one (30) and her mother enjoying a smaller one (5). Suppose that there are three episodes of pleasure contained in the consequence of *a18*: the daughter enjoying a large one (40), her mother enjoying a smaller one (5), and her father enjoying a smaller one (5).

Suppose that—in accordance with the chart—the strangers' daughter would benefit more (if given the relevant opportunity) than would the single mother's daughter, the rural family's resources (medical, educational, recreational) being so much more limited than the single mother's.

Act utilitarianism (AU) requires that the sole maximizer *a18* be performed: The single mother must benefit the stranger at her daughter's expense on the classical scheme. Social Proximity Utilitarianism (SPU) generates different implications in the case, for SPU is constructed so as to reflect the intuition the possible future goods and evils of those near and dear to us are more weighty from a moral perspective than those of others, or—as Feldman writes—SPU is designed so that “other things being equal, I have a greater moral obligation to guard the welfare of my own child than I have to guard the welfare of some unknown, distant child.”

Here is one way that such a weighting might be articulated. Let the social proximity-adjusted value for any hedono-doloric episode, E , contained in the consequence of some alternative, A , be the product of multiplying the social proximity of the agent of A to the experiencer of E by the hedono-doloric value of E . The *social proximity-adjusted hedonic utility of an alternative, A* , can now be defined as the sum of all of the social proximity-adjusted values of all of the hedono-doloric episodes contained in the consequence of A . Applying these concepts to the case, the social proximity-adjusted hedonic utilities for the alternatives can be articulated.

Recall that there are two episodes of pleasure contained in the consequence of $a17$: a large one for the single mother's daughter (30) and a smaller one for the single mother (5). Suppose that the social proximity of the single mother to her daughter is .9. On this assumption, the social proximity-adjusted value of her daughter's episode of pleasure is 27. Calculating the social proximity-adjusted value of the single mother's episode of pleasure is somewhat trickier.

What is the social proximity of an agent of some act to herself? Some have suggested that it must be equal to one, for no one is socially closer to a person than she is to herself.⁴⁶ While some will likely find this assumption repugnant, it might prove useful to let it stand for now, addressing it later. On this assumption, the social proximity-adjusted value of the single mother's episode of pleasure is 5. Summing the social proximity-adjusted values of the hedono-doloric episodes contained in $a17$ results in the social proximity-adjusted hedonic utility of $a17$: 32.

The social proximity-adjusted hedonic utility of $a18$ is generated in similar fashion. There are three episodes of pleasure contained in the consequence of $a18$: that enjoyed by the rural strangers' daughter (40), her father's (5), and her mother's (5). Suppose that the social proximity of the single mother to each of these strangers is .4. The social proximity-adjusted

values of these episodes of pleasure are then 16, 2, and 2, respectively. Summing these values results in the social proximity-adjusted hedonic utility of *a18*: 20.

	<u>HU</u>	<u>SP-Adjusted HU</u>
<i>a17</i> : benefiting one's daughter	35	32
<i>a18</i> : benefiting some rural strangers' daughter	50	20

Social Proximity Utilitarianism can be clearly stated in this way:

SPU: An alternative, *A*, is morally right iff no alternative to *A* has a higher social proximity-adjusted hedonic utility than *A* has.

Just as Feldman notes, SPU (like the classical AU) is a maximizing theory: It requires that social proximity-adjusted hedonic utility be maximized in all cases. SPU delivers the implications that (at least some of) the “nearest and dearest” objectors crave: *a17* is deemed morally obligatory despite the fact that the single mother could do more for the world by performing *a18* (at the expense of her dear daughter, of course). SPU also reflects the intuition articulated in Derek Parfit's quotation above that we are not permitted to save our child's toy rather than save another child's life, but we are permitted to grant our child a benefit rather than grant a greater benefit to a stranger's child.

For those who find SPU's implications a bit too extreme, a form of social proximity utilitarianism theoretically closer to classical AU can be constructed in a way by now very familiar to the reader. Consider this non-maximizing method for ranking any two alternatives: Let it be the case that *alternative A SP-outranks alternative B* iff (i) $HU(A) > HU(B)$ and (ii) $SP-A HU(A) > SP-A HU(B)$. According to this method of comparing alternatives, one alternative is SP-outranked by another iff it has both a lower hedonic utility and a lower social proximity-adjusted hedonic utility than the other. With this method of ranking in hand, a second version of social proximity-adjusted utilitarianism can be clearly formulated.

SPUr: An alternative, *A*, is morally right iff no alternative to *A* SP-outranks *A*.

SPUr preserves some conceptual closeness to AU: Anything permitted by AU is permitted by SPUr. Furthermore, SPUr permits us to consider the possible future goods and evils of those most near and dear to us as more weighty from a moral perspective than those of others, but it does not require such a weighting, for we are always permitted to act for the sake of the greater good according to the view. In the case above, SPUr deems either alternative permissible, for neither is SP-outranked by the other: The single mother is permitted to grant either young girl the benefit in question.

Strapping the self-other asymmetry onto SPUr results in a theory interesting in many respects. Supererogation emerges on such a scheme similar to the way it does on the combinatorial Egoistically-Adjusted Self-Other Utilitarianism considered above. Note that the social proximity-adjusted utility of any alternative, A , is the sum of the social proximity-adjusted utility for the agent of A and the social proximity-adjusted utility for others of A . Consider now this somewhat more complicated method for ranking any two alternatives:

Alternative A SP&SO-outranks alternative B iff (i) $HU(A) > HU(B)$, (ii) $SP-Adj HU(A) > SP-Adj HU(B)$, (iii) $SP-Adj HU-for-Others(A) > SP-Adj HU-for-Others(B)$, and (iv) $HU-for-Others(A) > HU-for-Others(B)$.

Requirement (i) for outranking in this scheme captures the classical utilitarian insight.

Requirement (ii) illustrates that this system of ranking is sensitive to matters of social proximity; in fact, (i) and (ii) are the two components of SPUr's outranking system above. Requirements (iii) and (iv) illustrate that the self-other asymmetry is in play: An alternative doesn't suffer on this outranking scheme solely in virtue of having a low hedonic utility for its agent—whether adjusted for social proximity or not.

On this method of comparatives alternatives, one alternative is SP&SO-outranked by another just in case it is outranked by its competitor in each of the four ways articulated above: It must have a lower hedonic utility, lower social proximity-adjusted hedonic utility, lower social

proximity-adjusted hedonic utility for others, and a lower hedonic utility for others than some competing alternative in order for it to be SP&SO-outranked by it. Consider this resulting form of Social Proximity Self-Other Utilitarianism:

SPSOU: An alternative, *A*, is morally right iff no alternative to *A* SP&SO-outranks *A*.

SPSOU is sensitive to a variety of phenomena: the greatest good, social proximity, and the self-other asymmetry. SPSOU is also sensitive to something similar to the egoistic weighting captured in Egoistically-Adjusted Utilitarianism (EAU): The social proximity of an agent to herself plays a serious role in determining the social proximity-adjusted utility for any alternative she performs. Unlike EAU, the agent-relative features of SPSOU are motivated through somewhat different theoretical avenues: The force of the “nearest and dearest” objection combined with the self-other asymmetry are responsible for the agent-relative properties of SPSOU. A brief return to the social-proximity weighting that these versions of social proximity utilitarianism assign to the agent herself should be investigated before an exploration of the types of supererogation possible on an SPSOU scheme transpires.

In the construction of Social Proximity Utilitarianism (SPU), it was assumed—for the sake of some preliminary theoretical simplicity—that the social proximity of an agent to herself is maximal—a rating of one—on the basis that no one is socially closer to a person than she is to herself, and this rating—and subsequent weighting—plays a substantial role in determining what morality requires from the perspective of theories like SPU, SPUR, and SPSOU. Some might find this maximal weighting of the agent’s own possible future goods and evils as objectionable, particularly those who find distasteful the egoistic weighting in EAU. In fact, assuming this maximal weighting results in theories that are—in some respects—radically more egoistic than even EAU.

A spectrum of SPSOU theories can be imagined, each assigning different weightings to the agent’s own possible future goods and evils. Too high an assignment might result in theories implausibly egoistic; too low an assignment results in theories strapped with the problems associated with Altruistically-Adjusted Utilitarianism discussed above. The maximal weighting will be maintained throughout this investigation of SPSOU — again, for reasons of theoretical simplicity— but it is recognized that the such a weighting enjoys no absolutely privileged status. Theorists are encouraged to assign whichever weighting they find most apt, adjusting values for the following “supererogation” discussion in the appropriate way.

Recall the case motivating the “too demandingness” objection against classical AU, supposing the following values for the new “social proximity” categories.

	<u>HU</u>	<u>HU-for-Agent</u>	<u>HU-for-Others</u>	<u>SP-Adj HU-for-Agent</u>	<u>SP-Adj HU-for-Others</u>	<u>SP-Adj HU</u>
<i>a1</i>	100	-10	110	-10	80	70
<i>a2</i>	90	20	70	20	55	75
<i>a3</i>	70	0	70	0	55	55
<i>a4</i>	55	25	30	25	20	45

a1 and *a2* are each deemed morally permissible on an SPSOU evaluation; neither one of them is outranked by a competing alternative. *a1* maximizes hedonic utility, hedonic utility for others, and social proximity-adjusted hedonic utility for others. *a2* maximizes social proximity-adjusted hedonic utility, serving as a sort of “nearest and dearest” response to the “too demandingness” objection.

SPSOU bestows supererogatory status upon *a1* (mirroring the implications of satisficing and egoistically-adjusted forms of utilitarianism): By performing *a1* rather than the morally permissible *a2*, the agent sacrifices 80 units of hedonic value that she is morally permitted to acquire so that others will prosper.

Now recall the case motivating the “too selfishness” objection— again— supposing the following values for the “social-proximity” categories.

<u>HU</u>	<u>HU-for-Agent</u>	<u>HU-for-Others</u>	<u>SP-Adj HU-for-Agent</u>	<u>SP-Adj HU-for-Others</u>	<u>SP-Adj HU</u>
-----------	---------------------	----------------------	----------------------------	-----------------------------	------------------

<i>a13</i>	100	100	0	100	0	100
<i>a14</i>	90	20	70	20	55	75
<i>a15</i>	70	0	70	0	55	55
<i>a16</i>	55	25	30	25	20	45

The implications of SPSOU are clear in this case: *a13*, *a14*, and the irrational *a15* are morally permissible on the scheme. *a13* maximizes hedonic utility and social proximity-adjusted hedonic utility. *a14* and *a15* maximize hedonic utility for others and social proximity-adjusted hedonic utility for others. *a14* enjoys supererogatory status on an SPSOU evaluation: In performing *a14* rather than the self-benefiting *a13* (or the irrational *a15*), the agent sacrifices 80 units of hedonic value that she is morally permitted to acquire for the sake of others.

Supererogation also emerges on this SPSOU scheme when an agent sacrifices potential future benefits she is morally permitted to acquire so that people socially close to her will prosper—even if doing so fails to maximize hedonic utility for others so long as it ranks sufficiently high in either social proximity hedonic utility or social proximity-adjusted hedonic utility for others. The possibilities for supererogation on an SPSOU scheme are abundant. This is not to suggest, however, that SPSOU is immune from serious, substantial objections. SPSOU permits—and might be utilized as theoretical motivation to promote—invidious types of partiality, including nepotism, provincialism, and parochialism. Many theorists will surely find such implications unacceptable.

VII. Conclusion

Many theorists currently reject the entire family of consequentialist ethical theories on the ground that none of these theories exhibits attractive supererogatory properties. Some consequentialists have rejected classical versions of consequentialism on the same grounds, opting for satisficing versions instead. Here theoretical options open to consequentialists and supererogation critics have been articulated that have yet to be carefully explored, possible sources of supererogation springing from a variety of phenomena: the “ties at the top”

phenomenon as well as consequentialist responses to the “too demandingness,” “too selfishness,” and “nearest and dearest” objections to classical utilitarianism. Such an articulation—this author hopes—provides both a defense of consequentialism and a service to anyone looking for supererogatory possibilities within the family of consequentialist moral theories conceptually close to classical formulations of act utilitarianism.

References

- Arneson, R. J. (2009) “Moral Limits and the Demands of Beneficence.” In D. K. Chatterjee (ed.) *Morality, Affluence, and the Distant Needy*: 33-58, Cambridge: Cambridge University Press.
- Aristotle. (1915) *Nicomachean Ethics*. Edited and translated under the supervision of W. D. Ross. Oxford: Oxford University Press.
- Baier, K. (1958) *The Moral Point of View*. Ithaca: Cornell University Press.
- Bentham, J. (1789) *An Introduction to the Principles of Morals and Legislation*. London: T. Payne.
- Bradley, B. (2006) “Against Satisficing Consequentialism.” *Utilitas* 18: 97-108.
- Broad, C. D. (1953) “Self and Others.” Herbert Spencer Lecture, Oxford University. Reprinted in D. R. Cheney (ed.) *Broad’s Critical Essays in Moral Philosophy*, New York: George Allen & Unwin/Humanities Press. Available at <http://www.ditext.com/broad/self.html>.
- Byron, M. (2003) *Satisficing and Maximizing: Moral Theorists on Practical Reason*. UK: Cambridge.
- Dreier, J. (2006) *Contemporary Debates in Moral Theory*. Malden: Blackwell.

- Feinberg, Joel (1968) "Supererogation and Rules." In J.J. Thomson and G. Dworkin (eds.) *Ethics*: 391-411, New York: Harper and Row.
- Feldman, F. (1978) *Introductory Ethics*. Englewood Cliffs: Prentice-Hall.
- Feldman, F. (1998) "Critical Study of *Living High and Letting Die: Our Illusion of Innocence* by Peter Unger." *Noûs* 32: 138-147.
- Hampton, J. (1993) "Selflessness and the Loss of Self." *Social Philosophy & Policy* 10: 135-165.
- Harwood, S. (1993) "Eleven Objections to Utilitarianism." In Louis Pojman (ed.) *Moral Philosophy: A Reader*: 141-154, Indianapolis: Hackett Publishing Company, Inc.
- Heyd, D. (1982) *Supererogation*. Cambridge: Cambridge University Press.
- Hurka, T. (1990) "Two Kinds of Satisficing." *Philosophical Studies* 59: 107-111.
- Jackson, F. (1991) "Decision-theoretic Consequentialism and the Nearest and Dearest Objection." *Ethics* 101: 461-482.
- Jeske, D. and Fumerton, R. (1997) "Relatives and Relativism." *Philosophical Studies* 87: 143-157.
- Kagan, S. (1989) *The Limits of Morality*. Oxford: Clarendon Press.
- Kawall, J. (2003) "Self-Regarding Supererogatory Actions." *Journal of Social Philosophy* 34: 487-498.
- McConnell, T. (1980) "Utilitarianism and Supererogatory Acts." *Ratio* 22: 36-38.
- McNamara, P. (1996) "Making Room for Going Beyond the Call." *Mind* 105: 415-450.
- Mellema, G. (1991) *Beyond the Call of Duty*. Albany: State University of New York Press.
- Mill, J.S. (1861) *Utilitarianism*. In J.S. Mill, *Collected Works of John Stuart Mill*: 203-59.
- Mill, J.S. (1963ff) *Collected Works of John Stuart Mill*, J. M. Robson (ed.). Toronto: University of Toronto Press.

- Montague, P. (1989) "Acts, Agents, and Supererogation." *American Philosophical Quarterly* 26: 101-111.
- Moore, G. E. (1912) *Ethics*. New York: Oxford University Press.
- Mulgan, T. (1993) "Slote's Satisficing Consequentialism." *Ratio* 6: 121-134.
- Mulgan, T. (2001) *The Demands of Consequentialism*. New York: Oxford University Press.
- New, C. (1974) "Saints, Heroes, and Utilitarians." *Philosophy* 49: 179-89.
- Parfit, D. (1979) "Is Common-Sense Morality Self-Defeating?" *Journal of Philosophy* 76: 533-545.
- Portmore, D. (2007) "Consequentializing Moral Theories." *Pacific Philosophical Quarterly* 88: 39-73.
- Portmore, D. (2008a) "Dual-Ranking Act Consequentialism." *Philosophical Studies* 138: 409-427.
- Portmore, D. (2008b) "Are Moral Reasons Morally Overriding?" forthcoming in *Ethical Theory and Moral Practice*.
- Rawls, J. (1971) *A Theory of Justice*. Cambridge: Harvard University Press.
- Scheffler, S. (1984) *The Rejection of Consequentialism: A Philosophical Investigation of the Considerations of Underlying Rival Moral Considerations*. New York: Oxford University Press.
- Sider, T. (1993) "Asymmetry and Self-Sacrifice." *Philosophical Studies* 70: 117-132.
- Sidgwick, H. (1907) *The Methods of Ethics*. 7th Ed. Cambridge: Hackett Publishing Company.
- Slote, M. (1984a) "Morality and Self-Other Asymmetry" *Journal of Philosophy* 81: 179-192.
- Slote, M. (1984b) "Satisficing Consequentialism, Part I." *Aristotelian Society* 58: 139-164.
- Slote, M. (1989) *Beyond Optimizing: A Study of Rational Choice*. Cambridge: Harvard University Press.

- Smart, J.J.C. (1973) “An Outline of a System of Utilitarian Ethics.” In J.J.C. Smart and B. Williams (eds.), *Utilitarianism: For and Against*: 3-74, Cambridge: Cambridge University Press.
- Vallentyne, P. (2006) “Against Maximizing Act-Consequentialism.” In *Contemporary Debates in Moral Theory*, ed. by James Dreier.
- Vessel, J.P. (2009) “Defending a Possibilist Insight in Consequentialist Thought.” *Philosophical Studies* 142: 183-195.
- Wolf, S. (1982) “Moral Saints.” *Journal of Philosophy* 79: 419-439.
- Zimmerman, M. J. (1996) *The Concept of Moral Obligation*. Cambridge: Cambridge University Press.
- Zimmerman, M. J. (2006) “The Relevance of Risk to Wrongdoing.” In R. Feldman, K. McDaniel, J. Raibley, and M. J. Zimmerman (eds.), *The Good, the Right, Life, and Death*: 151-170, Aldershot: Ashgate Press.

¹ Pioneering formulations of classical utilitarianism can be found in Bentham (1789), Mill (1861), Sidgwick (1907), and Moore (1912).

² Some consequentially minded ethicists reject this, claiming that certain future facts about the actual world (namely, whether or not agents will engage in future immoral behavior) can generate obligations incompatible with the best possible course of action available to us. See Zimmerman (1996), Zimmerman (2006), and Vessel (2009) for critical discussions of this dispute.

³ Michael Slote famously argued along these lines in his (1984b) and (1989), so has Samuel Scheffler in his (1994) and Tim Mulgan in his (2001). Also see Wolf (1982) and the slew of

contemporary ethicists calling themselves “satisficing consequentialists,” particularly Hurka (1990), a variety of theorists in Byron (2003), and perhaps Vallentyne (2006).

⁴ Cf. Baier (1958): 203f., Smart (1973): 55, Feldman (1978): 48ff., McConnell (1980): 36ff., Heyd (1982): 76ff., Mellema (1991): 79ff., New (1974), Kagan (1989), Portmore (2007), (2008a), and (2008b). See Zimmerman’s rich discussion of supererogation claims in Chapter 8 of his (1996). The following theorists appear to believe that supererogation is incompatible with any theory containing a maximizing requirement: Moore (1912): 78, Scheffler (1984): 21f., and Slote (1989): 3 & 171.

⁵ Hedonism will be utilized as the theory of value anchoring various consequentialist theories because it is simple and because of the role it plays in classical statements of consequentialism, but the assumption is not important. Any theory of value consonant with the thesis that all values are commensurable should work equally well.

⁶ Henry West provides this scenario for the abstract case under consideration. Suppose a young physician is considering one of two posts in the *Doctors Without Borders* program, the first (*a1*) an assignment in a war-zone that would undoubtedly result in the physician experiencing substantial suffering while saving many people, the second (*a2*) an assignment in a depressed urban area where the physician would avoid substantial personal suffering but save fewer people.

⁷ Cf. Slote (1989), Harwood (1993), §8.2.1 of Zimmerman (1996), and Portmore (2007).

⁸ Cf. some of the most extensive treatments on supererogation: Heyd (1982), Mellema (1991), and Zimmerman (1996).

⁹ Portmore (2007), (2008a), and (2008c).

¹⁰ Heyd (1982), Montague (1989), and Mellema (1991).

¹¹ Heyd (1982).

¹² Feinberg suggests in his (1968) that agents must endure serious personal sacrifice or take a serious risk in the performance of a supererogatory alternative.

¹³ Some critics will not be moved by these considerations. In virtue of the fact that AU does not imply that *a5* is morally superior to *a6*, they will continue to claim that it is impossible to perform supererogatory action under a classical utilitarian evaluation. This dispute over the analytic restrictions related to the term ‘supererogatory’ remains contested. Even if such critics are closer to the truth, the type of supererogation (or supererogation*) being argued for here is extraordinarily conceptually close to their notion of supererogation, so close that this author firmly believes that its properties should be seriously explored.

¹⁴ Slote (1984) and (1989), Hurka (1990), and perhaps Vallentyne (2006), among others.

¹⁵ Whether or not *a3* or *a4* come out as morally permissible on this satisficing scheme is unclear. The concept of being good enough is just a bit too mysterious for confident judgments regarding the normative statuses of *a3* and *a4*, especially considering that SatU is a very generic form of satisficing consequentialism.

¹⁶ Note that *a1* is also deemed to be morally better than *a2* according to SatU: It generates more hedonic value for the world than does *a2*.

¹⁷ Note, however, that there are no fewer than six explications of the concept in Bradley’s excellent (2006).

¹⁸ Bradley lodges this type of objection against the whole range of satisficing conceptions of consequentialism in his (2006).

¹⁹ Bradley (2006).

²⁰ Scheffler defends a view of this kind in his (1984). Also see Portmore (2008a).

²¹ See Mulgan (2001) for a similarly motivated but very distinct view from the one about to be formulated.

²² It might be objected that the double weighting here is somewhat arbitrary. Why not count our own possible future goods and evils as three times as important, or as one and a half times as important? First, what is being provided here is a foundation from which a plethora of egoistically-adjusted versions of utilitarianism might be constructed. Second, there might be some way to narrow the range of plausible possible weightings. This matter will be addressed shortly.

²³ Theodore Sider introduces this type of method of ranking for a different reason in his (1993), a reason to be pursued shortly.

²⁴ One might be inclined to object that the focus on the agent alone in the formulation of EAU is in some way arbitrary. Why not extend that focus towards anyone to whom the agent is partial: a spouse, child, friend, or pet? For simplicity's sake, EAU will remain the view under investigation with the understanding that it can be extended to capture agents' partialities, something that will be explored explicitly in the penultimate section.

²⁵ It is left for the reader to inspect EAU's implications regarding the alternative set *a9-a12*. EAU does not permit the gratuitous prevention of goodness contained in the consequence of *a11*.

²⁶ Commenting on Ethical Neutralism—the family of ethical theories of which classical utilitarianism is a paradigm member—C. D. Broad writes in his (1953): “It seems to be in some directions immorally selfish....”

²⁷ Broad considers altruistic versions of consequentialism in his (1953).

²⁸ Note also that supererogation is still possible under this demanding, maximizing scheme. AAU delivers supererogatory implications only when the “ties at the top” phenomenon emerges in such a way that one of the permissible alternatives is more self-benefiting than another, similar to the way supererogation emerges under the classical scheme.

²⁹ Broad (1953).

³⁰ *Ibid.*

³¹ Slote (1984a): 180-181. Also see New (1974) and Portmore (2008a).

³² Sider (1993). New (1974) attempts to formulate a version of utilitarianism that incorporates the self-other asymmetry: Optative Altruistic Utilitarianism. The author finds SOU theoretically superior to Optative Altruistic Utilitarianism in that SOU entails that some alternatives are capable of being morally permissible even if they fail to maximize either hedonic utility or hedonic utility for others.

³³ Somewhat curiously, Portmore (2008a) suggests that *a13* is deemed supererogatory by SOU in virtue of the facts that it is morally optional and—in a sense—morally better than *a14*: It generates more overall value for the world than does *a14*. Both Portmore (2008a) and Kawall (2003) appear to endorse the controversial claim that completely self-regarding acts are capable of being morally supererogatory, entailing their rejection of the popular position that supererogatory action must benefit others to some extent.

³⁴ While attractive in many respects, SOU is not immune from substantial attacks. SOU permits agents to incur massive sacrifices so that others might enjoy negligible benefits. In a symposium at the 2008 Pacific Division Meetings of the American Philosophical Association, Allen Coates suggested that such sacrifices are not only irrational but immoral. Jean Hampton expresses similar thoughts in her “Selflessness and the Loss of Self”: “Moreover, I have also argued that in some situations in which duty to others and duty to self are opposed, the self-regarding choice is actually the morally superior and obligatory choice.”

For critics who find such sacrifices immoral, a close cousin to SOU can be imagined, one that restricts the extent to which agents are permitted to incur sacrifices for others, perhaps via some sort of weighting system restricting the range of permissible sacrifices just as the weighting system in EAU restricts the extent to which one is permitted to benefit oneself.

³⁵ Compare this with Schefflerian Utilitarianism (SU), which is formulated, motivated, and defended in Portmore (2008a).

³⁶ The curious reader can see that alternatives maximizing hedonic utility cannot be E&SO-outranked.

³⁷ Sidgwick (1907): 242.

³⁸ Parfit (1979): 536.

³⁹ Jeske and Fumerton (1997): 146.

⁴⁰ It is recognized that Aristotle's discussion here is substantially richer than a mere articulation and response to the "nearest and dearest" objection, but the objection is clearly in the forefront of Aristotle's discussion.

⁴¹ Sidgwick (1907): 241ff.

⁴² Jeske and Fumerton (1997): 146. Frank Jackson endorses Sidgwick's strategy—at least in part—in his (1991):

We have seen that the good consequentialist should focus her attentions on securing the well-being of a relatively small number of people, herself included, not because she rates their welfare more highly than the welfare of others but because she is in a better position to secure their welfare. Typically, this will involve her in settling on a relatively extended program of action which will take some resolution and strength of character to carry forward successfully. Before she starts she knows, if she is at all like most of us, that the chances of success are much greater if she makes the relatively small group those who are her family and friends, rather than those she hardly knows. (481)

⁴³ Frank Jackson cites Mother Teresa and Ralph Nader as possible counterexamples to this general rule on p. 481 of his (1991).

⁴⁴ Feldman (1998): 141.

⁴⁵ The author is incredibly grateful and indebted to Fred Feldman for the use of Feldman's unpublished lecture notes from his 1998 seminar on consequentialism where he motivates, formulates, and criticizes Social Proximity Utilitarianism (SPU). The first version of SPU formulated here is practically indistinguishable from the account under investigation in Feldman's lecture notes.

⁴⁶ Jean Hampton argues along this line in her "Selflessness and the Loss of Self."