

## SATISFICING CONSEQUENTIALISM AND SCALAR CONSEQUENTIALISM

### I. Satisficing Consequentialism: The General Idea

**SC** An act is morally right (i.e., morally permissible) if and only if its consequences are “good enough.”

### II. The Motivation for Adopting Satisficing Consequentialism

It is often thought that by adopting satisficing consequentialism we can bring consequentialism more in line with commonsense morality. Specifically, it seems that, unlike maximizing act consequentialism, satisficing act consequentialism can accommodate both of the following:

#### A. Agent-Centered Options

An *agent-centered option* is an option to either do what’s better in terms of one’s own interests or do what’s better in terms of the impersonal good. These options provide agents with the freedom to permissibly act so as to further their own interests out of proportion to their weight in the impersonal calculus.

#### B. Supererogatory Acts

*Supererogatory acts* are acts that go above and beyond the call of duty. An act  $x$  is supererogatory if and only if both (i)  $x$  is morally optional and (ii)  $x$  is, in some sense, morally better than some permissible alternative.

### III. Five Versions of Satisficing Consequentialism

#### A. Absolute Level Satisficing Consequentialism

**ALSC** “There is a number,  $n$ , such that: An act is morally right iff either (i) it has a utility of at least  $n$ , or (ii) it maximizes utility” (Bradley 2006, 101).\*

\*Note that Bradley defines absolute level satisficing *consequentialism* in terms of utility. Thus, it might, more appropriately, be called absolute level satisficing *utilitarianism*.

An Illustration:

Let  $U(x)$  = the overall utility that is produced if S does  $x$ , and let’s suppose that  $n = 100$ .

<i>act</i>	<u><i>U(x)</i></u>	<i>moral status</i>
a <sub>1</sub>	<b>+180</b>	supererogatory and morally best*
a <sub>2</sub>	<b>+90</b>	impermissible
a <sub>3</sub>	<b>+125</b>	supererogatory
a <sub>4</sub>	<b>+110</b>	merely permissible

\*Actually, ALSC does not by itself entail that a<sub>1</sub> is supererogatory. To get that entailment, we need to supplement ALSC with the following two plausible assumptions: (1) an act is supererogatory if and only if both (a) it is morally optional and (b) it is, in some sense, morally better than some permissible alternative, and (2) on consequentialism, an act is morally better than some other act in the sense relevant to (b) if and only if it produces more good/utility than the other act does.

#### B. Comparative Level Satisficing Consequentialism

**CLSC** “There is a number,  $n$  ( $n > 0$ ), such that: An act is morally right iff its utility plus  $n$  is greater than or equal to the utility of a utility-maximizing alternative” (Bradley 2006, 101).

An Illustration:

Let  $U(x)$  = the overall utility that is produced if S does  $x$ , and let’s suppose that  $n = 100$ .

<i>act</i>	<u><i>U(x)</i></u>	<i>moral status</i>
a <sub>1</sub>	<b>+500</b>	supererogatory and tied for morally best
a <sub>2</sub>	<b>+400</b>	merely permissible
a <sub>3</sub>	<b>+390</b>	impermissible
a <sub>4</sub>	<b>+450</b>	supererogatory
a <sub>5</sub>	<b>+500</b>	supererogatory and tied for morally best

#### C. Double Satisficing Consequentialism

**DSC** There is a number,  $m$  ( $m > 0$ ), as well as a number,  $n$  ( $n > 0$ ), such that: An act is morally right iff either (i) it has a utility of at least  $m$ , or (ii) its utility is less than  $m$ , but its utility plus  $n$  is greater than or equal to the utility of a utility-maximizing alternative (Bradley 2006, 101).

An Illustration:

Let  $U(x)$  = the overall utility that is produced if S does  $x$ , and let's suppose that  $m = 200$  and that  $n = 50$ .

<i>act</i>	<u><math>U(x)</math></u>	<i>moral status</i>
a <sub>1</sub>	<b>+100</b>	impermissible
a <sub>2</sub>	<b>+170</b>	merely permissible
a <sub>3</sub>	<b>+150</b>	impermissible
a <sub>4</sub>	<b>+220</b>	supererogatory and morally best
a <sub>5</sub>	<b>+200</b>	supererogatory

D. Situational Absolute Level Satisficing Consequentialism

**SALSC** "There is a number,  $n$ , such that: An act is morally right iff either (i) the situation that would obtain after the act has value of at least  $n$ , or (ii) the act maximizes utility" (Bradley 2006, 101).

Two Illustrations:

Let  $U(Sx)$  = the overall utility (in millions) that will obtain after S does  $x$ , and let's suppose that  $n = 1$  million.

<i>act</i>	<u><math>U(Sx)</math></u>	<i>moral status</i>
a <sub>1</sub>	<b>.9</b>	impermissible
a <sub>2</sub>	<b>1.2</b>	supererogatory and morally best
a <sub>3</sub>	<b>.95</b>	impermissible
a <sub>4</sub>	<b>1.1</b>	merely permissible
a <sub>5</sub>	<b>1.15</b>	supererogatory

Let  $U(Sx)$  = the overall utility (in millions) that will obtain after S does  $x$ , and let's suppose that  $n = 1$  million.

<i>act</i>	<u><math>U(Sx)</math></u>	<i>moral status</i>
a <sub>1</sub>	<b>.9</b>	impermissible
a <sub>2</sub>	<b>.8</b>	impermissible
a <sub>3</sub>	<b>.95</b>	merely permissible
a <sub>4</sub>	<b>.95</b>	merely permissible
a <sub>5</sub>	<b>.93</b>	impermissible

E. Individualist Situational Absolute Level Satisficing Consequentialism

**ISALSC** “There is a number,  $n$ , such that: An act is morally right iff either (i) in the situation after the act, each person’s welfare level is at least  $n$ , or (ii) the act maximizes utility” (Bradley 2006, 101).

Two Illustrations:

Let  $U(ISx)$  = the overall utility of the person who is the least well off after S does  $x$ , and let’s suppose that  $n = 100$ .

<i>act</i>	<u><math>U(ISx)</math></u>	<i>moral status</i>
a <sub>1</sub>	<b>99</b>	impermissible
a <sub>2</sub>	<b>120</b>	supererogatory and morally best
a <sub>3</sub>	<b>95</b>	impermissible
a <sub>4</sub>	<b>111</b>	merely permissible
a <sub>5</sub>	<b>115</b>	supererogatory

Let  $U(ISx)$  = the overall utility of the person who is the least well off after S does  $x$ , and let’s suppose that  $n = 100$ .

<i>act</i>	<u><math>U(ISx)</math></u>	<i>moral status</i>
a <sub>1</sub>	<b>90</b>	impermissible
a <sub>2</sub>	<b>80</b>	impermissible
a <sub>3</sub>	<b>95</b>	merely permissible
a <sub>4</sub>	<b>95</b>	merely permissible
a <sub>5</sub>	<b>93</b>	impermissible

#### IV. Objections to These Five Versions of Satisficing Consequentialism

##### A. Permitting Gratuitous Murder

SALSC permits committing murder for no reason at all provided that one’s current situation is so far above the threshold,  $n$ , that the situation after one commits murder will still be above  $n$ . Suppose, for instance, that  $n$  equals 1 million hedons. And suppose that, at present, there are 1.5 million hedons, and that murdering Smith will reduce the overall value in the world by only 100,000 dolars.

What’s even worse is that SALSC implies, as we’ll see below, that committing murder for no reason at all is sometimes supererogatory.

An Illustration:

Let  $U(Sx)$  = the overall utility (in millions) that will obtain after  $S$  does  $x$ , and let's suppose that  $n = 1$  million. Assume that  $a_1$  is the act of gratuitously murdering six people, that  $a_2$  is the act of minding one's own business while sitting on the couch watching TV, that  $a_3$  is the act of gratuitously murdering Smith, and that  $a_4$  is the act of gratuitously murdering two people.

<i>act</i>	<u><math>U(Sx)</math></u>	<i>moral status</i>
$a_1$	.9	impermissible
$a_2$	1.5	supererogatory and morally best
$a_3$	1.4	supererogatory
$a_4$	1.3	merely permissible

#### B. Permitting Gratuitous Harm

ISALSC is problematic for two reasons. First, assuming that  $n$  is a reasonably high number (and it needs to be for the view to be plausible), it will turn out that ISALSC will, in the actual world (where it is impossible for any one person to do anything to ensure that everyone will be over the threshold,  $n$ ), be extensionally equivalent to maximizing consequentialism. Thus, ISALSC will turn out to be just as demanding as maximizing consequentialism and, thus, too demanding. Second, ISALSC permits causing gratuitous harm in the world in which everyone is well over the threshold,  $n$ . Suppose, for instance, that  $n$  equals 100 hedons. And suppose that, at present, everyone has 150 hedons, and that punching someone in the nose will reduce his or her overall utility by only 10 dolors.

What's even worse is that ISALSC implies, as we'll see below, that causing harm for no reason at all is sometimes supererogatory.

An Illustration:

Let  $U(ISx)$  = the overall utility of the person who is the least well off after  $S$  does  $x$ , and let's suppose that  $n = 100$ . Assume that  $a_1$  is the act of gratuitously punching Smith in the nose six times, that  $a_2$  is the act of minding one's own business while sitting on the couch watching TV, that  $a_3$  is the act of gratuitously punching everyone other than Smith with 150 hedons in the nose once, and that  $a_4$  is the act of gratuitously punching Smith in the nose twice.

<i>act</i>	<u><math>U(ISx)</math></u>	<i>moral status</i>
$a_1$	90	impermissible
$a_2$	150	supererogatory and morally best
$a_3$	140	supererogatory

a<sub>4</sub>            130            merely permissible

### C. Permitting the Gratuitous Prevention of Goodness

ALSC, CLSC, DSC, SALSC, and ISALSC all permit going out of one's way to prevent some good state of affairs from coming about for absolutely no reason at all. I'll illustrate this using ALSC, but the objection applies, *mutatis mutandis*, to the other four versions.

An Illustration:

Let  $U(x)$  = the overall utility that is produced if S does  $x$ , and let's suppose that  $n = 100$ . Assume that  $a_1$  is the act of minding one's own business while sitting on the couch watching TV, that  $a_2$  is the act of dissuading nine others from donating money to Oxfam, that  $a_3$  is the act of the act of dissuading six others from donating money to Oxfam, and that  $a_4$  is the act of dissuading seven others from donating money to Oxfam.

<i>act</i>	<u><i>U(x)</i></u>	<i>moral status</i>
a <sub>1</sub>	<b>+180</b>	supererogatory and morally best
a <sub>2</sub>	<b>+90</b>	impermissible
a <sub>3</sub>	<b>+120</b>	supererogatory
a <sub>4</sub>	<b>+110</b>	merely permissible

### V. Taking Self-Sacrifice into Account: Self-Sacrificing Satisficing Consequentialism

The problem with all of the above forms of satisficing consequentialism is that they permit an agent to go out of her way, making some self-sacrifice even, so as to prevent some good state of affairs from coming about. The idea that it is morally permissible to go out of one's way to prevent some good state of affairs from obtaining is very implausible. To avoid the problem, we need a version of satisficing consequentialism that allows an agent to perform sub-optimal acts only when the sub-optimal act has 'good enough' consequences and only when bringing about better consequences would involve making some self-sacrifice. Garrett Cullity has come up with just such a version of satisficing consequentialism.

#### A. Cullity's Self-Sacrificing Absolute-Level Satisficing Consequentialism

**CSSALSC** "There is a number,  $n$ , such that: An act,  $a$ , performed by agent S, is morally right iff either (i)  $a$  has a utility of at least  $n$ , and any better alternative is worse for S than  $a$ ; or (ii)  $a$  maximizes utility" (Bradley 2006, 107).

## Two Illustrations

Let  $U_s(x)$  = the utility that accrues to S if S does  $x$ ,  $U_{-s}(x)$  = the utility that accrues to others if S does  $x$ , and  $U(x)$  = the overall utility that is produced if S does  $x$ . And let's suppose that  $n = 100$ .

<i>act</i>	$U_s(x)$	$U_{-s}(x)$	$U(x)$	<i>moral status</i>
a <sub>1</sub>	<b>+70</b>	+110	<b>+180</b>	supererogatory and morally best
a <sub>2</sub>	<b>+90</b>	+35	<b>+125</b>	permissible
a <sub>3</sub>	<b>+80</b>	+55	<b>+135</b>	supererogatory
a <sub>4</sub>	<b>+60</b>	+60	<b>+120</b>	impermissible

Let  $U_s(x)$  = the utility that accrues to S if S does  $x$ ,  $U_{-s}(x)$  = the utility that accrues to others if S does  $x$ , and  $U(x)$  = the overall utility that is produced if S does  $x$ . And let's suppose that  $n = 100$ .

<i>act</i>	$U_s(x)$	$U_{-s}(x)$	$U(x)$	<i>moral status</i>
a <sub>1</sub>	<b>+30</b>	+10	<b>+40</b>	impermissible
a <sub>2</sub>	<b>+30</b>	+35	<b>+65</b>	permissible
a <sub>3</sub>	<b>+10</b>	+5	<b>+15</b>	impermissible
a <sub>4</sub>	<b>+60</b>	+5	<b>+65</b>	permissible

### C. Why Self-Sacrificing Versions of Satisficing Consequentialism Are Unmotivated and Implausible

CSSALSC is perhaps the most plausible version of satisficing consequentialism, but we might wonder whether it's motivated. If we're willing to introduce self-sacrifice into the moral equation, we don't need to introduce satisficing in order to incorporate agent-favoring options and supererogatory acts. As we saw in the last lecture, Schefflerian Utilitarianism (SU) is a maximizing version of act consequentialism that incorporates both agent-favoring options and supererogatory acts. Moreover, SU seems superior to CSSALSC, as the example below illustrates. But, first, let's recall what SU says.

*Schefflerian Utilitarianism (SU):* S's performing  $x$  is morally permissible if and only if there is no available act alternative that would produce both (i) more utility for others (i.e., those other than S) than  $x$  would and (ii) at least as much egoistically-adjusted utility, where we include everyone's utility but adjust the overall total by giving S's utility ten times the weight of anyone else's.

### An Example

Let  $U_s(x)$  = the utility that accrues to S if S does  $x$ ,  $U_{-s}(x)$  = the utility that accrues to others if S does  $x$ ,  $U_{+s}(x) = U_{-s}(x) + [10 \times U_s(x)]$ , and  $U(x)$  = the overall utility that is produced if S does  $x$ . And let's suppose that  $n = 100$ .

<i>act</i>	$U(x)$	$U_s(x)$	$U_{-s}(x)$	$U_{+s}(x)$	<b>SU</b>	<b>CSSALSC</b>
a <sub>1</sub>	+10,010	+10	+10,000	+10,100	perm.	perm.
a <sub>2</sub>	+100	+11	+89	+199	imperm.	perm.
a <sub>3</sub>	+90	+10	+80	+180	imperm.	imperm.

CSSALSC allows agents to be far too selfish. It is one thing to say, as SU does, that agents can give some priority to themselves, but it is quite another to say, as CSSALSC does, that agents can give absolute priority to themselves so long as the utility they produce meets some threshold,  $n$ . CSSALSC implies that an agent needn't sacrifice even one hedon so as to ensure that others get thousands of hedons. For instance, in the above example, S is permitted to perform a<sub>2</sub> even though S could ensure that others gets thousands of hedons more by performing a<sub>1</sub> and the cost to S would be merely 1 less hedon.

Of course, it is probably possible to come up with a more plausible version of self-sacrificing satisficing consequentialism than CSSALSC, but what would motivate us to look for such a version? The motivation for adopting satisficing over maximizing consequentialism in the first place was to bring consequentialism closer in line with commonsense morality. But, again, once we introduce self-sacrificing into the moral equation, satisficing becomes superfluous. We can accommodate for both agent-favoring options and supererogatory acts on a maximizing version of act consequentialism.

## VI. Scalar Consequentialism

### A. Reconceiving Consequentialism: Consequentialism without Demands

As Norcross sees it, we should understand consequentialism, not as a theory of rightness, but as a theory of the comparative moral value of alternative acts. So conceived, consequentialism doesn't issue in requirements or permissions.

### B. Norcross's Arguments for Reconceiving Consequentialism

Norcross argues that we should reconceive consequentialism in the way suggested above, for the consequentialist should reject any theory of rightness where right and wrong is an all-or-nothing affair. The consequentialist should reject any such theory for the following two reasons:

THE FIRST REASON: A theory that held both that there was a duty of beneficence and that rightness and wrongness was an all-or-nothing affair “would have to say that there was a *threshold*, e.g., at 10 percent, such that if one chose to give 9 percent one would be wrong, whereas if one chose to give 10 percent one would be right. If this distinction is to be interesting, it must say that there is a *big* difference between right and wrong, between giving 9 percent and giving 10 percent, and a small difference between pairs of right actions [e.g., giving 10 versus 11 percent], or pairs of wrong actions [e.g., giving 8 versus 9 percent]” (Norcross 2006, 41).

But, as Norcross argues, the consequentialist should deny that the difference between giving 8 percent and giving 9 percent is any less significant than the difference between giving 9 percent and giving 10 percent. In each case, the difference in terms of how much good is done is the same.

THE SECOND REASON: Norcross says that a “related reason to reject an all-or-nothing line between right and wrong is that the choice of any point on the scale of possible options as a threshold for rightness will be *arbitrary*” (Norcross 2006, 41).

MY RESPONSE: I don’t think that either objection against the all-or-nothing theory of right and wrong succeeds. To see why, consider that the moral status of action might be a function of both moral reasons and non-moral reasons. That is, it might be that although one has more moral reason to give more of one’s income to charity, one has more non-moral reason to give less of one’s income to charity. And, perhaps, what determines just how much one is morally required to give is the following meta-criterion, which I defended in the previous lecture.

**MP**      S’s performing *x* is *morally permissible* if and only if there is no available alternative that S has better requiring reason to perform and no worse reason, all things considered, to perform.

If this is right, then S is morally required to give up to the point that giving more would tip the balance such that S would then have more reason, all things considered, to give less.

Thus, ‘right’ and ‘wrong’ might be like ‘worth it’ and ‘not worth it’. It might be worth it to trade some service (say, an hour of babysitting) for \$10, but not worth it to trade that same service for \$9. It’s not that there is a bigger difference between \$9 and \$10 than there is between \$8 and \$9, but a \$1 increase in the offered payment from \$9 to \$10 can make all the difference as to whether one ought to trade one’s service for that payment. This is true even if a \$1

increase in the offered payment from \$8 to \$9 makes no difference as to whether one ought to trade one's service for that payment.

And note that there is nothing arbitrary about the \$10 threshold; the threshold is set by the point at which it becomes worth it to trade one's service for that quantity payment. Likewise, there needn't be anything arbitrary about the 10 percent threshold; the threshold is set by the point at which it becomes objectively irrational to give more to charity. In both cases, the threshold lies where an increase in one of the two competing factors tips the balance.

### C. How Scalar Utilitarianism Compares to Traditional Utilitarianism

Traditional utilitarianism is objectionable for, at least, the following three reasons: (i) it holds that we are required to sacrifice our own interests so as to promote the good of others whenever doing so will maximize utility, (ii) it holds that we have very little moral freedom and that we are always morally required to maximize utility, and (iii) it denies that there are supererogatory acts.

Norcross argues that scalar utilitarianism can avoid objections (i) and (ii) and capture the intuition behind objection (iii). Scalar utilitarianism is not too demanding, for it makes no demands at all. It never requires an agent to sacrifice her own interests for the sake of promoting the good of others. (I wonder why Norcross sees this as an improvement; it seems that whereas traditional utilitarianism was too demanding, scalar utilitarianism isn't demanding enough, for, on scalar consequentialism, there's no requirement to even take a minute to save some kid drowning in a shallow pond.) Scalar utilitarianism avoids (ii), for, again, scalar utilitarianism makes no demands. Scalar utilitarianism cannot accommodate the idea that there are supererogatory acts—that is, acts that are morally superior to some permissible alternative, for scalar utilitarianism not only fails to issue any demands, but also fails to issue any permissions. There can be no going beyond duty, when there are no duties to go beyond. Nevertheless, Norcross argues that the intuition behind the idea of supererogation is the belief that some actions are morally better than what can be expected of a reasonably decent person in the circumstances, and scalar utilitarianism certainly leaves room for such a notion.

### D. Is scalar consequentialism a genuine rival, say, to deontology?

Since scalar consequentialism isn't a criterion of rightness, we might wonder if it is even a genuine rival to other moral theories that do specify some criterion of rightness. Norcross argues that while scalar consequentialism cannot provide a rival account of the substance of our moral obligations, we can take scalar

consequentialism to be a far more radical alternative to a theory like deontology, rule consequentialism, or traditional act consequentialism.

Why not reconceive rightness as a scalar concept? Norcross offer two reasons. First, there would be no point to doing so, for in that case there would be no real distinction between the rightness of an action and the moral goodness of an action. Second, as Norcross has argued elsewhere, there is no satisfactory account of good and bad *actions*, as opposed to good and bad *states of affairs*, with which to equate right and wrong actions.

E. Will consequentialism still be action-guiding? Reasons without Demands

At this point, we might wonder whether scalar consequentialism is even action-guiding and whether, if it isn't, it makes sense to even call it a moral theory. But Norcross argues that a moral theory doesn't have to issue in requirements to be action-guiding, it is sufficient for a moral theory provide an account of our moral reasons for action to be action-guiding. And scalar consequentialism gives us the following account of moral reasons. The fact that an action would produce a good state of affairs (i.e., one with positive utility) provides the agent with a moral reason to perform, whereas the fact that that an action would produce a bad state of affairs (i.e., one with negative utility) provides the agent with a moral reason to refrain from performing it. And the better the state affairs that act would produce the more moral reason there is to perform it.

I wonder, though, both whether it will be possible to distinguish moral reasons from non-moral reasons and whether it will be possible to talk about the strengths of reasons without appealing to requirements and permissions. For instance, it seems to me that the most natural way to distinguish between a moral reason and a non-moral reason is in terms of whether that reason is capable of making an act morally required/supererogatory or not. And it seems that the most natural way to understand the strength of a reason is in terms of its ability to require and justify acts. But if on scalar consequentialism there are requirements or permissions, I lose sight of what it means to call a reason moral rather than non-moral and I lose sight of what it means to call one moral reason stronger than another.