AXIOLOGICAL ABSOLUTISM AND RISK[1]

Seth Lazar and Chad Lee-Stronach

Australian National University

Abstract

Consider the following claim: given the choice between saving a life and preventing any number of people from temporarily experiencing a mild headache, you should always save the life. Many moral theorists accept this claim. In doing so, they commit themselves to some form of 'moral absolutism': the view that there are some moral considerations (like being able to save a life) that cannot be outweighed by any number of lesser moral considerations (like being able to avert a mild headache). In contexts of certainty, it is clear what moral absolutism requires of you. However, what does it require of you when deciding under risk? What ought you to do when there is a chance that, say, you will not succeed in saving the life? In recent years, various critics have argued that moral absolutism cannot satisfactorily deal with risk and should, therefore, be abandoned. In this paper, we show that moral absolutism can answer its critics by drawing on—of all things—orthodox expected utility theory.

1. INTRODUCTION

Some think that in any choice between saving a life and averting headaches, one should always save the life, no matter how many people whose temporary, minor pain one could otherwise alleviate.[2] They are the heirs of a more extreme form of absolutism, for which some moral

considerations matter so much that they could never be forsaken, regardless of what else was at stake: let justice be done though the heavens fall.[3]

We distinguish between species of absolutism below. For now, the genus is anyone who thinks there are some higher and lower considerations, such that no amount of the lower consideration can trump some amount of the higher consideration. Their 'moral preference' for the higher consideration is absolute.[4] Absolutism sometimes seems quite plausible for decision-making under certainty. But it faces challenges for decision-making with imperfect information. If you must always save a life, rather than avert any number of headaches, what does this imply if your alternatives are chancy? We surely cannot tolerate *any* number of headaches, however large, for the sake of *any* chance of saving a life, however small. But how then should we trade-off risks of breaching the higher consideration against risks of forsaking the lower one?

Advocates of absolutism have been slow to address this question.[5] Their critics have been quicker. They deny that any plausible extension of absolutism can work for decision-making under risk.[6] They argue that absolutists cannot use orthodox decision theory, then offer them some alternative, before revealing the flaws in that alternative. Responding, some absolutists try to vindicate those alternative standards.[7]

In this essay, we short-circuit this debate, and point a way forward. Absolutists can use orthodox decision theory as readily as anybody else,

once we model their 'moral utility function' in a charitable way.[8] Decision-theoretic absolutism does not, however, emerge with a clean bill of health. We show that while some of the objections against absolutist moral theories still have intuitive purchase, they have been misdiagnosed. Clarifying the problem for decision-theoretic absolutism raises fundamental questions about the nature of ethical theory, which everyone has to address.

2. Modelling Absolutist Theories[9]

The first step is to precisify our topic. This paper is about axiological absolutism, and about decision theory. In defining absolutism we speak of 'considerations'. These could be reasons, values, or anything else that bears on choice in a similar way (our goal is to preserve some neutrality between these alternatives). We talk about moral preferences: consideration *C* is morally preferred to *c* just in case, in a choice between serving *C* or *c* where all else is equal, one morally ought to serve *C*. Since we care only about *moral* preference orderings, we will sometimes simply write that *C* is preferred to *c*.

We need to quantify considerations. To keep things simple, we focus on cases in which this means multiplying tokens. If *c* is the consideration of averting a minor, temporary headache to a single person, then when we write 'any amount of *c*', we mean multiple tokens of that consideration

(not multiple headaches for the same person). *C* is 'absolutely morally preferred to *c*' means that some amount of *C* is morally preferred to any amount of *c*. When referring to 'amounts', we always mean total amounts. A 'class' of moral considerations is just a set unified by some shared property.

A moral theory is absolutist if and only if it posits that there is a higher moral consideration *C*, and a lower consideration *c*, such that there is an amount of *C* that is absolutely morally preferred to any amount of *c*. If *C* is 'saving a life' and *c* is 'averting a single minor temporary headache' then a theory is absolutist if it says that for any number of tokens of *c* one should serve *C*, saving the life rather than averting the headaches.

Absolutism can vary in *scope* and *stringency*. Scope refers to the size of the class of considerations to which the higher consideration is absolutely preferred. Stringency refers to how much of the higher consideration is needed to instantiate the absolute preference.

Extreme Absolutism is maximally stringent and has maximal scope. It posits a higher consideration, *C*, any amount of which is absolutely preferred to any amount of the class, *c**, of all other considerations. *Extreme Absolutism* might say, for example, that one may never kill an innocent person, no matter the good that could thereby be achieved.

Moderate Absolutism is maximally stringent, but has narrow scope. It posits a higher consideration *C* and a lower consideration *c* such that *any amount* of *C* is absolutely preferred to any amount of *c*. 'Life for headaches'

is intended to invoke *Moderate Absolutism*. We say that, for *Moderate Absolutism*, *C* is 'strongly absolutely preferred' to *c*. Equivalently, we also write that *C* is 'strongly superior' to *c*.

Weak Absolutism also has narrow scope, identifying a relation between specific considerations, rather than classes of considerations. But it is less stringent than *Moderate Absolutism*. It posits a higher consideration *C* and a lower consideration *c* such that *there is an amount* of *C* that is absolutely preferred to any amount of *c*. We say that *C* is 'weakly absolutely preferred' to *c* in cases like this, or, equivalently, *C* is 'weakly superior' to *c*.[10]

We focus on *axiological* absolutism. This is absolutism grounded in values: one consideration is absolutely preferred to another because no amount of tokens of the latter can have as much value as some amount of the former. *Deontic* absolutists ground absolute moral preferences in something other than value. They might think that constraints are strongly superior to value considerations,[11] or that one ought not aggregate lesser considerations to overtop higher considerations,[12] or that the presence of one reason can exclude another from consideration.[13] As we note below, it will sometimes be possible to use axiological absolutism to *represent* deontic absolutism with a value function.

Decision theory is the study of decision-making with imperfect information. Rational decision theory offers an account of rational

decision-making with imperfect information. Moral decision theory offers the moral counterpart.

Very roughly, orthodox rational decision theory tells us that, faced with imperfect information, we should identify our options and the possible outcomes to which each might lead, assign utilities to those outcomes, and probabilities conditional on taking that option, then multiply the two numbers together, before choosing the option for which the sum of those products is greatest.[14]

However, many rational decision theorists deny that these utilities are 'real'. Instead, they are induced from an agent's preferences over gambles, together with a representation theorem, which shows how those preferences given those probabilities entail a particular utility function, unique up to positive affine transformation (provided the preferences obey some formal rational constraints).[15]

Moral decision theory cannot work this way.[16] We want to *guide* choice under risk (with defined probabilities) and uncertainty (without defined probabilities), not simply to represent it. If we knew what our preferences over moral gambles should be, we wouldn't need the decision theory. And obviously these preferences must do more than obey minimal axioms of rationality. One could obey all those axioms and yet be a moral monster. Moral decision theory depends on the *moral* utility function being *objective*, not merely induced from one's preferences over gambles.[17]

In this essay we ask: can absolutists do decision theory? The first step

is to model our absolutist theory's moral utility function for outcomes. While this is most natural for axiological absolutism, we could also represent deontic absolutism using a similar moral utility function, so the lessons of this essay have broader application. To represent a moral theory with a value function is not to imply anything in particular about the character of the underlying moral reasons. [18] We might look to deontic considerations for the ultimate explanation of why we should be absolutists, but model those considerations with a value function in order to apply them to decision-making under risk.

Philosophers sometimes assume that the only way to model an absolutist moral utility function is to posit an infinite value difference between higher and lower considerations.[19] If C is strongly absolutely preferred to c, then C must be infinitely weightier than c. But this cannot be right.[20]

Imagine a consideration C_1 that is very slightly weaker than C. Surely some number of tokens of C_1 can outweigh a single token of C. Now imagine a further consideration, C_2, that is very slightly weaker than C_1. Again, surely some number of tokens of C_2 can outweigh a single token of C_1. Keep going in this way until we reach C_n, which is only very slightly weightier than c, so some number of tokens of c can outweigh a single token of C_n. But if the outweighing relationship is transitive, then this entails that some number of tokens of c can outweigh a single token of C, contrary to the claim that C is infinitely more valuable than c.

Alternatively, there must be some C_m that is absolutely preferred to its near neighbour C_{m+1}, which is extremely counterintuitive given that we can make the distance between neighbours arbitrarily fine.[21]

Axiological absolutism cannot posit an infinite value difference between higher and lower considerations. But we should not therefore reject absolutism. Rather than positing an infinite value difference between *C* and *c*, we should posit that the value of *c* cannot exceed some upper bound.[22] We deny the 'simple additivity' of c.[23] The most obvious alternative is that *c* has diminishing marginal value, which decreases asymptotically towards zero, so that the total value approaches a limit beneath the value of a single token of C.[24]

If this is how to represent axiological absolutism, then the sequence objection fails. The upper bound of C_1 might be above the value of one token of *C*, and the upper bound of *c* might be above the value of one token of C_1, but the upper bound of *c* might be less than the value of one token of *C*. So we can preserve the absolute preference for *C* over *c*, without positing that any given C_m is absolutely preferred to its nearest neighbour.

The infinitist approach seems untenable; the more charitable way to model an absolutist utility function is with a bounded value function.[25] Note that this can accommodate *Extreme*, *Moderate* and *Weak Absolutism* alike.

3. Absolutism and Risk

This is good news for absolutist decision theory. At least since Pascal, it has been obvious that infinite value differences play havoc with decision-making under risk.[26] Any fraction of infinity is itself infinite, so if *C* were infinitely prior to *c*, then orthodox decision theory would prioritise any probability of *C* over any certain amount of *c*. This would yield wildly counterintuitive results.[27] Absolutists need not endorse them: bounded utility functions are quite consistent with orthodox decision theory. But nobody yet has considered precisely how one would implement this idea for absolutist decision-making under risk. That is our task in this section.

Suppose only two considerations, *C* and *c*, bear on the moral status of some choice. *C* is weakly absolutely preferred to *c*—some amount of *C* outweighs any amount of *c*. You face a risky choice, in which you could realise some amount of *C* and some of *c*. To apply this kind of *Weak Absolutism* to decision-making under risk, we can do the same thing as any other orthodox decision theorist: consider the possible outcomes and assign them a moral utility (MU), based on our objective moral theory.[28]

Say that the moral utility of *C* increases linearly without upper bound, each token adding 20 more units of moral utility. By contrast, *c* diminishes at the margin, approaching an upper bound of 99 units of moral utility. So five tokens of *C* outweigh any amount of *c*, but some amount of *c* can outweigh four or fewer tokens of C. So, *C* is weakly absolutely preferred to

c.

Now consider this decision problem, in which ϕ and ψ are actions, and A and B are possible states of the world.

	A	B
ϕ	10*C*; 0*c* (MU=200)	0*C*; 0*c* (MU=0)
ψ	0*C*; 1,000,000,000,000*c* (MU≈99)	0*C*; 1,000,000,000,000*c* (MU≈99)

Suppose that $p(A|\phi)$ and $p(B|\phi)$ are both 0.5, as are $p(A|\psi)$ and $p(B|\psi)$. Then the expected moral utility of ϕing is 100, while that of ψing is very close to 99. So you ought to ϕ, despite the astronomically larger number of tokens of *c* that you could realise by ψing.

The weak absolute preference of *C* over *c*, however, does not entail that one will tolerate *any* risk to *c* to avoid *any* risk to *C*. Suppose the probability of B given ϕ or ψ is 0.51, while that of A is 0.49. Now you should ψ, because that has just short of 99 expected moral utility, while the expected moral utility of ϕing is only 98.

This basic approach extends to other cases involving *Weak Absolutism*. We need only state the moral utility function at stake, quantifying the absolutist preference, then input those values into our decision tables, to generate results consistent with standard decision theory. *Moderate*

Absolutism is the same. Suppose that any amount of *C* is absolutely preferred to any amount of *c*. Each token of *C* has 100 moral utility; the moral utility of *c* cannot exceed a limit of 99. Now consider the following decision table:

	A	B
ϕ	2*C*; 0*c* (MU=200)	0*C*; 0*c* (MU=0)
ψ	0*C*; 1,000,000,000,000*c* (MU≈99)	0*C*; 1,000,000,000,000*c* (MU≈99)

If the probabilities are again 0.5 across the board, then you should ϕ, despite the enormous number of tokens of *c* forgone. But if the probability of B given ϕ or ψ is 0.51, while that of A is 0.49, then the expected moral utility of ψing tops that of ϕing (≈99 to 98), and you ought to ψ. This, again, even though *C* is strongly absolutely preferred to *c*.

Extreme Absolutism is structurally identical to *Moderate Absolutism*, except that where *Moderate Absolutism* posits a specific consideration *C* that is strongly absolutely preferred to a specific consideration *c*, *Extreme Absolutism* posits a single token of *C* that outweighs *c**, the class of all other considerations. So what goes for *Moderate Absolutism* also goes, mutatis mutandis, for *Extreme Absolutism*.

This yields an interesting result. *Extreme Absolutism* starts to look much

less stringent in risky cases. Even if *C* is strongly absolutely preferred to *c**, no extreme absolutist need think that *any risk of C* is strongly absolutely preferred to *c**. This might make *Extreme Absolutism* more palatable. Even if no number of certain tokens of *c** can outweigh a single token of *C*, some number of tokens of *c** can outweigh some number of tokens of *C* when the latter are discounted for their improbability of being realised. As long as *C* is not infinitely superior to *c**, as *p*(*C*) approaches zero, so does the expected moral utility of that outcome, and some probability of realising c* will be morally preferable.

Absolutists face no special challenges when using decision theory. Of course, when we consider more complex decision problems, it would quickly become obvious that actually *doing* absolutist decision theory would be hellishly difficult. Although we take this concern seriously, it is a general problem. Moral philosophy alone cannot provide a decision procedure for realistic choices. But it can offer a criterion of subjective permissibility.[29] Absolutists are no worse placed than others to do this.

4. Debunking the Objections to Absolutist Decision Theory

With this understanding of absolutist decision theory in hand, we can turn to the problems that philosophers have raised for applying absolutist moral theories to risky choice situations. Swayed by a presupposition that

absolutism must involve infinite value differences, and so the denial of continuity, they have furnished the absolutist with a variety of different decision rules focusing on the idea of a probability threshold, in this way: suppose that C is your reason to abide by a constraint, and $p(\neg C)$ is the probability that ϕing will breach the constraint. Philosophers have assumed that the absolutist's decision rule must be something like: it is permissible to ϕ only if $p(\neg C)$ is less than some threshold. For example, it is permissible to risk killing an innocent person only if the probability that the person is innocent is less than p. They have then run a series of objections to highlight the problems with this kind of threshold view. Their absolutist opponents have uniformly shared their assumption that orthodox decision theory cannot work for absolutism, and have sought to defend some version or other of a threshold approach.[30]

Our first and perhaps most important point is this: absolutists and their critics have been labouring under a mistake. They, like anyone else, can use orthodox decision theory to extend their views to decision-making under risk. To be sure, there are many problems with using rational decision theory as a model for moral decision theory. But the inability to accommodate plausible absolutism is not one of them.

Our second key observation: the debate about absolutism and risk has been too much driven by Jackson and Smith's focus on *Extreme Absolutism*. We cannot think of a single influential contemporary deontologist who, in published work, endorses *Extreme Absolutism*. Even

Robert Nozick (who stretches the term 'contemporary') countenanced letting side-constraints be outweighed, if the stakes were high enough.[31] This misunderstanding is particularly important, since all of the most persuasive cases used by Jackson and Smith, and Huemer, involve high stakes whatever one does, often involving risking killing the innocent in order to save many more lives, or more realistic but structurally similar cases from war, or from criminal justice.[32]

These examples are, to our minds, unhelpful for understanding the implications of risk for absolutist moral theories, because only the most stringent absolutist would argue that we may never harm the innocent, no matter how great the good we can realise thereby, especially given that there is some probability that one's target is not in fact innocent.[33] What's more, as we have seen, even if one's objective moral theory *was* committed to *Extreme Absolutism* in cases like these, that would not entail that it would be equally rigorist for decision-making under risk.

Third: the most telling criticisms levelled by Huemer, and Jackson and Smith, have to do with problems of agglomeration that beset threshold-based views of subjective permissibility. Focusing on the kind of case introduced at the start of this section, these objections presuppose that absolutists must commit to a single value for $p(\neg C)$ (the probability that your action breaches a constraint), such that if it exceeds that threshold, the action is subjectively impermissible. The objection starts by claiming that this threshold must be arbitrary, then proceeds to consider cases in

which one must perform multiple acts, each with some probability of breaching the constraint.

Each of these assumptions can be questioned. Absolutists need not endorse any contextually invariant thresholds. For a given decision problem involving risking breaching a constraint *C*, if matters are simple enough then we can probably extract a threshold *t* such that if $p(\neg C)$ exceeds *t*, then it is subjectively impermissible to proceed. But this would be misleading; we could do the same for any maximising decision rule. Placing a $1 bet on a coin that will pay $2 if it lands heads, nothing otherwise, is rational if and only if $p(\text{heads}) \geq 0.5$. But obviously we wouldn't describe 'maximise expected utility' as a threshold rule. Insofar as one insists on extracting a threshold from the absolutist theory, it is obviously no more arbitrary than this threshold extracted from rational decision theory. Where it sits depends on the stakes, as it obviously should. If the probability of heads is 0.499999, then the bet is irrational, regardless of how close you are to 0.5.

A further corollary: since the 'threshold' emerges from an expected moral utility calculation given what is at stake, if we do extract one, it will depend on what is at stake in that particular decision problem. Several of the cases put forward by absolutism's critics miss this point, positing a single value for the threshold regardless of whether one's action risks (for example) killing one innocent person or two. And if you are contemplating saving two groups of innocent people, rather than just one,

then that might also affect the stakes, and so the implicit threshold at which it becomes permissible to proceed. The same will be true if more than one higher-order consideration is at stake. So all the examples used so far are ill-framed.

Ultimately we agree that these critics, albeit through a glass darkly, have identified a significant challenge for absolutism. However, they have misdiagnosed the problem. The real problem does not derive from a commitment to thresholds that absolutists need not make. Instead it derives from a set of fundamental questions that go far beyond absolutism, and must be addressed by *any* moral theory: namely, which units of action are the proper focus of moral enquiry: acts or campaigns? We concede, though, that this problem has particular purchase for those who reject the simple additivity of value. We come back to this in Section 5.

We have shown that *Extreme Absolutism* and *Moderate Absolutism* reduce to wide- and narrow-scope versions of *Weak Absolutism* in risky decisions. This must be so as long as the higher-order considerations are not infinitely superior to the lesser considerations. But might one think this an ersatz kind of absolutism?[34] In response, it is worth making one feature of our approach clear. If the moral utility of the lower consideration is bounded, and diminishes at the margin, then, as the risk of breaching the higher consideration increases, it will take more and more of the lower consideration to counterbalance that risk.

Imagine, for example, that not killing an innocent person was strongly absolutely preferred to saving innocent lives, so that it would be impermissible to kill an innocent person no matter how many lives you could save thereby.[35] (Few people hold such an extreme view, but let that slide.) But suppose that the upper bound of the value of saving innocent lives is equivalent to nine-tenths of the moral disutility of knowingly killing an innocent person. So, if the probability that this action will kill an innocent person is 0.1, say, then there is some number of lives certainly saved that could make that action permissible. Note, though, that as the probability that this action will kill an innocent person increases, the number of people whom you would have to be sure of saving in order to permissibly run this risk increases by a disproportionate amount, as each additional life saved adds less to the action's overall utility. Once you reach a 0.9 probability of killing an innocent person, no number of lives certainly saved would make it permissible to proceed. To repeat, we think *this* kind of absolutism is deeply implausible. But it is clearly not ersatz.

5. Additivity and the Act/Campaign Challenge

Critics of absolutism have argued that it is uniquely ill-equipped to address decision-making under risk. They are wrong. Absolutists can use orthodox decision theory as well as anyone else, if their absolutism is

either grounded in or adequately represented by a denial of simple additivity, rather than the assertion of infinite value differences. Since the latter generate implausible results even for decision-making under certainty, absolutists are forced to this path anyway.

However, absolutism does not emerge unscathed. Even once cleaned up and presented in its most charitable form, it generates arguably counterintuitive results. Interestingly, though, the reason it does so is quite different from what its critics have assumed—and it really has nothing to do with risk.

The basic problem is this. Some moral (and rational) decision theories include elements that generate divergent judgements of what seem to be morally identical phenomena. Specifically, they enjoin us to differently assess a sequence of acts, depending on whether we consider each act in isolation from the others, or the whole sequence as a campaign. In decision theory this is sometimes described as a conflict between 'local' and 'global' rationality. It is relatively common: philosophers have shown how vagueness, infinite utilities, rejection of the reflection principle, and risk-averse attitudes can all lead to clashes of this sort.[36] And in ethics, the phenomenon is not confined to absolutist decision theory. Indeed, Frank Jackson himself (with Robert Pargetter) has argued that their Professor Procrastinate is one such case.[37] Jackson thinks that Procrastinate ought to [accept the request to review the paper and write the report on time]. But given that he will not write the report on time if he accepts, he ought not

to [accept the request to review the paper]. If we consider the sequence of actions, then he acts impermissibly by refusing to referee the paper; if we consider the individual action, he acts permissibly. Many other such cases have been put forward that make the same point.

To convey the specific form of this problem for absolutist decision theory, we need a cleaner case than has been offered so far, by critics focused on debunking *Extreme Absolutism*, a view that nobody holds.[38]

In this case, *C* is letting an innocent person die, who would otherwise live a happy and long life; *c* is a unit of financial benefit to a person, equivalent to the value to an already wealthy person of $100.[39] Each token of *c* is one (different) person benefited in this way. Plausibly, no amount of *c* could outweigh a single token of *C*. No matter how many people we could benefit by $100, if we could instead save an innocent person's life then we should do so. This looks like a plausible kind of *Moderate Absolutism*.

Assuming that *C* is not infinitely superior to *c*, for some probability *p* and some number of tokens *n*, one must be indifferent between *p*(*C*) and *nc*. To fix things, let's stipulate that the moral disutility of one token of *C* is -9,999, while that of *c* diminishes at the margin, approaching 10 at the limit as *n* approaches infinity. On these numbers, it would be permissible to take a 0.001 risk of failing to save a life in order to give *n* people $100, for a high enough value of *n*.

Now suppose that on Monday morning, if you press button 1, you can

achieve *nc* for sure, but only by running a 0.0009 risk of killing an innocent person, A. Since *nc* is greater than 0.0009*C*, you proceed. Then on Monday afternoon, you face the same decision problem: pressing button 2 will realise *nc* for a different set of beneficiaries, but subject a different possible victim, B, to a 0.0009 risk of death. The decision problem being the same, you proceed again.

On Tuesday you face two identical decisions at the same time. The probabilities are the same as before. If you push button 1, you expose only A to a 0.0009 risk of death, achieving *nc* (that is, benefiting *n* people by *c* amount). If you push button 2, you expose only B to a 0.0009 risk of death, achieving *nc*. And if you push both buttons, you expose both A and B to a 0.0009 risk of death, achieving *2nc*. If you push neither button, A and B are not exposed to risk, and nobody gets $100. In this case, you would be permitted to push button 1 or 2, but not both. If you were unable to push only one of the buttons, you would have to do nothing.

But now suppose that you have pushed button 1 on Tuesday morning, but in the afternoon discover that button 2 can be pressed. Then, considering only the difference you can now make to the world, it would be permissible to push button 2, having pushed button 1.

This result is supposed to be counterintuitive: how can it make a difference, the critic asks, whether one pushes both buttons simultaneously, or does one in the morning, and one in the afternoon? It seems that by simply repackaging or disaggregating our options, we can

generate different deontic verdicts.

In properly cleaned-up cases, these results may not be so counterintuitive. But suppose that they are. The task, then, is to diagnose what is causing these results, to ask whether this problem is unique to absolutism, and then to sketch a response. We will proceed in that order.

Absolutism's critics think that this problem of divergent verdicts on individual acts and campaigns composed of just the same acts derives from a problem of agglomeration, itself grounded in the absolutists' need to invoke a threshold in their moral decision theory. But we have shown that the problem arises even when we do not adopt a threshold approach to decision-making under risk. For our approach, the source of the problem is the rejection of simple additivity—and it has nothing to do with risk.

If *c* has a diminishing marginal moral utility function, then the weight of any token of *c* depends on how many other tokens of *c* are also in play. Say that a single token of *c*, on its own, has *X* utility. If the *c* moral utility function is concave and increasing, then *X* is the most a token of *c* can ever weigh. If there are multiple tokens of *c*, then each additional token must weigh less than the last. Now, imagine a sequence of acts ϕ_1, ϕ_2..., ϕ_{10}, in each of which a single token of *c* is at stake. The aggregate moral utility of those ten instances of *c* will be 10*X*. But if ten tokens of *c* are at stake in a single act, ψ, they must together weigh less than 10*X*, because of the

diminishing marginal moral utility function. Suppose that both the sequence, ϕ_1..., ϕ_{10} and the act ψ involve the same costs, which are equivalent to just less than 10*X*. Then it would be permissible to carry out the sequence, but not the single act, despite their involving the same expected costs, and the same expected benefits. And this seems counterintuitive: if the costs and benefits are what really matters, then why should we care whether they are realised serially or at once?[40]

One absolutist response: this is not news! It's a basic feature of endorsing any kinds of holistic value judgements. The objections simply thump the table and insist on additivity. But if the case for rejecting additivity is strong, as examples like 'life for headaches' suggest it is, then perhaps we should tolerate these intuitive costs. None of the competing moral theories come for free—endorsing additivity means endorsing the counterintuitive verdict in life for headaches; deontic absolutist alternatives have their own problems.

A second response: these objections raise a fundamental question in both normative ethics and decision theory, as to what the proper units of ethical evaluation are. For if we should judge the sequence of acts (ϕ_1..., ϕ_{10}) as a whole, or if we should disaggregate the compound act into its constituent parts, then these inconsistencies would not arise. As already noted, precisely the same question arises in rational decision theory.

Some defenders of absolutism have implicitly noticed this possibility, and advocated for either the 'disaggregating' or the 'bundling' approach.

Much of the subsequent discussion has shown that neither of these seems to be uniformly successful. In what remains of this section, we will suggest a way forward for absolutists—and everyone else too, since everyone needs an account of when campaigns are the proper object of ethical evaluation, and when the individual acts that constitute them are.

Our basic proposal: one size does not fit all. Sometimes we should disaggregate; sometimes bundle.[41] The trick is to find a principle that distinguishes these cases.

Suppose that 100 people are under threat of a terrorist attack. There are ten suspected terrorists. For each individual, the probability that he is part of the attack is 0.9. The terrorists have built in redundancy: if any one of them lives, the 100 will surely die. In the first scenario, the apparent terrorists are collocated—you can kill them all in one go. In the second, they are separate, and must be killed separately. What is the proper unit of evaluation here? The individual act of killing? Or the campaign to stop the terrorist attack?

Here it seems clear that the campaign is morally fundamental. If we were to kill the first target, but leave the others to carry out the attack, then we would run a 0.1 risk of killing an innocent person for no good at all. The good that justifies the risk is achieved only if we impose the same risk nine more times.[42] In this case, whatever your method of killing the targets—in one go or sequentially—the campaign as a whole is the proper object of moral assessment. Each individual act of killing can bring about

the good only if you kill at least eight of the other targets.

But now suppose that the 100 victims are in ten groups of ten, and that each of the presumed terrorists will attack a different group. Then the good that justifies each 0.1 risk of killing an innocent person is independent from the goods that justify running the other risks. Whether we proceed with a single strike or hit the presumed terrorists one after the other, it seems clear that we may consider each act (or portion of the compound act) independently.

The general idea is this: when a given risk of a bad outcome is causally sufficient to realise enough expected good to justify that risk, then that risk can be assessed separately from whatever else is at stake. We can weigh the expected moral disutility of running that risk against the expected good thereby realised. Suppose killing target A is causally sufficient to realise *m* expected moral disutility, and *n* expected moral utility. If *n* is greater than *m*, then killing A can in general be assessed in isolation from any subsequent acts. But if killing A is causally sufficient to realise *n*, then that expected utility cannot count in favour of any subsequent risks, for fear of double-counting. When considering killing target B, we must set *m* and *n* aside.

Now suppose that killing A is not causally sufficient to realise *n* expected moral utility (either it realises no expected good on its own, or it realises too little to justify *m*). In this case, if killing A is permissible it must be because it is necessary to some broader sequence of actions that

does realise enough expected moral utility to justify the expected moral disutility of all the risks involved.

How does this (roughly sketched) approach fare with the standard counterexamples to absolutism? Huemer and Jackson and Smith argue that if absolutists think that killing the innocent is absolutely prohibited, then they must be pacifists, and that if they think it is absolutely prohibited to punish the innocent, then they could not defend a workable criminal justice system. It bears repeating that these are not plausible instances of absolutism. The goods achieved by just wars are weighty enough to justify risking killing innocent people; something similar is true for punishment. But forget that point. When thinking about wars, or criminal punishment, should we evaluate the practice as a whole? Or the individual acts of which it is composed?

Criminal punishment seems more like our second example. The expected goods achieved by running a particular risk of punishing the innocent are independent of the expected goods achieved by running further risks in other cases. So it seems right to judge individual acts of punishment in isolation: they are causally sufficient to realise enough expected moral utility to justify the expected moral disutility that they involve.

Wars look more like the first case. Although the permissibility of some acts in war is 'free-standing', in general fighting a war is permissible only if the many risks involved are justified by the overarching good for which

they are necessary and sufficient. At the very least, we should *also* evaluate wars as whole campaigns. Many acts in war *would not* be justified if we focused only on the risks and benefits for which they are causally sufficient.

This is just a first step in a solution for absolutists.[43] We are confident in the general prescription: absolutists (and everyone else) need to think about when we should assess closely-connected acts together, and when we should disaggregate compound acts into their constituent parts. Our specific proposal is not exhaustive. It says simply that when a risky act is causally sufficient to realise some expected good it can be considered in isolation from the campaign of which it is part, and assessed as permissible or impermissible. When one risky act depends on others to realise its expected good, then they must be assessed together.

But absolutists should take heart. First, everybody needs, and nobody yet has, an adequate account of whether acts or sequences of acts are the proper object of moral evaluation. So they have many companions in guilt. Second, non-additive, and otherwise holistic, value functions are far from outlandish. Indeed, the more controversial assumption might be that moral utility is straightforwardly additive. Rational decision theory may have to presuppose that individuals have bounded utility functions in order to address problems like the St Petersburg Game.[44] Most people already think that money has diminishing marginal utility. Holistic and contextual interactions between values are well-established.[45] One might

even argue that presupposing additivity of value involves reifying and fetishising the mathematical notation with which we are representing our moral theories.

6. Conclusion

We are cautious absolutists. We find *Moderate Absolutism* compelling at the level of objective moral theory. One of us is an axiological absolutist; the other favours deontic absolutism, but thinks that it might be adequately represented by axiological absolutism for the purposes of doing decision theory. We have argued that critics of absolutist decision theory have lumbered it with arbitrarily chosen decision rules. Once we represent absolutism with a non-additive value function, rather than with infinite value differences, there is no special problem doing regular absolutist decision theory. In particular, the central objection to absolutism under risk would then rest on a misdiagnosis. The problem is not with the use of thresholds, but instead derives from rejecting the simple additivity of value. The commitment to a non-additive value function makes matters trickier for the absolutist, but many others are in the same boat; and the biggest problem that she faces is a problem for everyone: are acts or sequences of acts (which sequences?) the proper object of moral evaluation? Many non-standard versions of rational decision theory generate just the same kind of problem.

We have sketched a solution to that problem, but addressing it at length is a matter for further research. Our main aim has been to show that, for moral decision theorists who want to learn from rational decision theory, absolutism does not introduce any special new problems. Of course, some absolutists are dispositionally opposed to the whole value-theoretic project. And yet every moral theory has to have some way to cater for imperfect information. Absolutists who reject decision theory must propose a more plausible alternative.

Of course, just as there are objections to the infinitist approach to absolutism, the rejection of simple additivity will face challenges. To use absolutist decision theory as we propose, one has to show that diminishing marginal value functions can faithfully represent a plausible absolutist moral theory. Some think that there are general reasons to believe they cannot do so.[46] We take heart from the fact that a value function can faithfully *represent* a moral theory without being committed to a particular view of axiology, but defending this claim at length is, again, a matter for further research.

References

Aboodi, R., A. Borer, and D. Enoch 2008. 'Deontology, Individualism, and Uncertainty,' *Journal of Philosophy* 105/5: 259-72.

Anscombe, G. E. M. 1979. 'War and Murder,' in *Moral Problems: A Collection of Philosophical Essays,* ed. James Rachels, New York: Harper & Row: 393-407.

Arntzenius, F., A. Elga, and J. Hawthorne 2004. 'Bayesianism, Infinite Decisions, and Binding,' *Mind* 113/450: 251-83.

Arrhenius, G. 2005. 'Superiority in Value,' *Philosophical Studies* 123/1-2: 97-114.

Arrhenius, G. and W. Rabinowicz 2005a. 'Millian Superiorities,' *Utilitas* 17/2: 127-46.

Arrhenius, G. and W. Rabinowicz 2005b. 'Value and Unacceptable Risk,' *Economics and Philosophy* 21/2: 177-97.

Arrhenius, G. and W. Rabinowicz 2015. 'Value Superiority,' in *The Oxford Handbook of Value Theory,* ed. Iwao Hirose and Jonas Olson, Oxford: Oxford University Press: 225-43.

Bartha, P. 2007. 'Taking Stock of Infinite Value: Pascal's Wager and Relative Utilities,' *Synthese* 154/1: 5-52.

Binmore, K. and A. Voorhoeve 2003. 'Defending Transitivity against Zeno's Paradox,' *Philosophy and Public Affairs* 31/3: 272–79.

Bostrom, N. 2011. 'Infinite Ethics,' *Analysis and Metaphysics* 10: 9-59.

Briggs, R. 2014. 'Normative Theories of Rational Choice: Expected Utility,' in *Stanford Encyclopaedia of Philosophy,* ed. Edward Zalta.

Broome, J. 1991. *Weighing Goods: Equality, Uncertainty and Time*, Oxford: Blackwell.

Broome, J. 2010. 'No Argument against the Continuity of Value: Reply to Dorsey,' *Utilitas* 22/04: 494-96.

Brown, C. 2017. 'Maximalism and the Structure of Acts,' *Noûs* Online First.

Buchak, L. 2013. *Risk and Rationality*, Oxford: Oxford University Press.

Buchak, L. 2016. 'Decision Theory,' in *Oxford Handbook of the Philosophy of Probability,* ed. Alan Hájek and Christopher Hitchcock, Oxford: Oxford University Press.

Carlson, E. 2000. 'Aggregating Harms — Should We Kill to Avoid Headaches?,' *Theoria* 66/3: 246-55.

Carlson, E. 2001. 'Organic Unities, Non-Trade-Off, and the Additivity of Intrinsic Value,' *Journal of Ethics* 5: 335-60.

Colyvan, M., D. Cox, and K. Steele 2010. 'Modelling the Moral Dimension of Decisions,' *Noûs* 44/3: 503-29.

Dancy, J. 1993. *Moral Reasons*, Oxford: Blackwell.

Dorsey, D. 2009. 'Headaches, Lives and Value,' *Utilitas* 21: 36-36.

Easwaran, K. 2014. 'Decision Theory without Representation Theorems,' *Philosophers' Imprint* 14/27.

Elga, A. 2010. 'Subjective Probabilities Should Be Sharp,' *Philosophers' Imprint* 10/05.

Finnis, J. 1991. *Moral Absolutes : Tradition, Revision, and Truth*, Washington, D.C: Catholic University of America Press.

Goldman, H. S. 1976. 'Dated Rightness and Moral Imperfection,' *Philosophical Review* 85/4: 449-87.

Griffin, J. 1986. *Well-Being: Its Meaning, Measurement and Moral Importance*, New York: Oxford University Press.

Hájek, A. 2003. 'Waging War on Pascal's Wager,' *The Philosophical Review* 112/1: 27-56.

Hawley, P. 2008. 'Moral Absolutism Defended,' *Journal of Philosophy* 105/5: 273-75.

Hedden, B. 2013. 'Options and Diachronic Tragedy,' *Philosophy and Phenomenological Research* 87/1: 1-35.

Huemer, M. 2010. 'Lexical Priority and the Problem of Risk,' *Pacific Philosophical Quarterly* 91/3: 332-51.

Isaacs, Y. 2014. 'Duty and Knowledge,' *Philosophical Perspectives* 28/1: 95-110.

Jackson, F. and R. Pargetter 1986. 'Oughts, Options, and Actualism,' *Philosophical Review* 95/2: 233-55.

Jackson, F. and M. Smith 2006. 'Absolutist Moral Theories and Uncertainty,' *Journal of Philosophy* 103/6: 267-83.

Jackson, F. and M. Smith 2015. 'The Implementation Problem for Deontology,' in *Weighing Reasons,* ed. Barry Maguire and Errol Lord, Oxford: Oxford University Press: 279-91.

Jeffrey, R. C. 1983. *The Logic of Decision*, Chicago: University of Chicago

Press.

Jensen, K. K. 2008. 'Millian Superiorities and the Repugnant Conclusion,' *Utilitas* 20/3: 279-300.

Jensen, K. K. 2012. 'Unacceptable Risks and the Continuity Axiom,' *Economics and Philosophy* 28/1: 31-42.

Kagan, S. 1988. 'The Additive Fallacy,' *Ethics* 99/1: 5-31.

Kamm, F. M. 2007. *Intricate Ethics: Rights, Responsibilities, and Permissible Harm*, New York: Oxford University Press.

Kirkpatrick, J. R. 2017. 'Permissibility and the Aggregation of Risks,' *Utilitas*: 1-13.

Lazar, S. 2015. 'Risky Killing and the Ethics of War,' *Ethics* 126/1: 91-117.

Lazar, S. 2017a. 'Risky Killing: How Risks Worsen Objective Rights-Violations,' *Journal of Moral Philosophy* Online First.

Lazar, S. 2017b. 'In Dubious Battle: Uncertainty and the Ethics of Killing,' *Philosophical Studies* Online First.

Lazar, S. 2017c. 'Deontological Decision Theory and Agent-Centred Options,' *Ethics* 127/3: 579-609.

Lazar, S. 2017d. 'Aggregation and Risk,' *Unpublished MS.*

Lee-Stronach, C. 2016. 'Lexical Priorities under Risk,' *Unpublished MS.*

Norcross, A. 1997. 'Comparing Harms: Headaches and Human Lives,' *Philosophy and Public Affairs* 26/2: 135-67.

Norcross, A. 1998. 'Great Harms from Small Benefits Grow: How Death Can Be Outweighed by Headaches,' *Analysis* 58/2: 152-58.

Norcross, A. 2009. 'Two Dogmas of Deontology: Aggregation, Rights, and the Separateness of Persons,' *Social Philosophy and Policy* 26/1: 76-95.

Nover, H. and A. Hájek 2004. 'Vexing Expectations,' *Mind* 113/450: 237-49.

Nozick, R. 1974. *Anarchy, State and Utopia*, Oxford: Basil Blackwell.

Portmore, D. 2016a. 'Maximalism Versus Omnism About Permissibility,' *Pacific Philosophical Quarterly*: n/a-n/a.

Portmore, D. W. 2016b. 'Maximalism and Moral Harmony,' *Philosophy and Phenomenological Research*: n/a-n/a.

Raz, J. 1986. *The Morality of Freedom*, Oxford: Clarendon Press.

Raz, J. 1999. *Practical Reason and Norms*, Oxford: Oxford University Press.

Ridge, M. 1998. 'How to Avoid Being Driven to Consequentialism: A Comment on Norcross,' *Philosophy and Public Affairs* 27/1: 50-58.

Ryberg, J. 2002. 'Higher and Lower Pleasures – Doubts on Justification,' *Ethical Theory and Moral Practice* 5/4: 415-29.

Savage, L. 1954. *The Foundations of Statistics*,

Scanlon, T. M. 1998. *What We Owe to Each Other*, London: Belknap Press.

Taurek, J. M. 1977. 'Should the Numbers Count?,' *Philosophy and Public Affairs* 6/4: 293-316.

Temkin, L. S. 1996. 'A Continuum Argument for Intransitivity,' *Philosophy and Public Affairs* 25/3: 175-210.

Temkin, L. S. 2012. *Rethinking the Good : Moral Ideals and the Nature of Practical Reasoning*, Oxford: Oxford University Press.

Tenenbaum, S. 2017. 'Action, Deontology, and Risk: Against the

Multiplicative Model,' *Ethics* 127/3: 674-707.

Thomson, J. J. 1990. *The Realm of Rights*, Cambridge, Mass.: Harvard University Press.

Voorhoeve, A. 2014. 'How Should We Aggregate Competing Claims?,' *Ethics* 125/1: 64-87.

Zollman, K. J. S., A. Bjorndahl, and A. J. London Forthcoming. 'Kantian Decision Making under Uncertainty: Dignity, Price, and Consistency,' *Philosophers' Imprint*.

---

[1] This paper was presented at the ANU's graduate workshop in Moral, Social and Political Theory. For comments, thanks to John Broome, Anne Gelling, Matt Kopec, Ten Herng-Lai, Kirsten Mann, Jeremy Strasser, and Shang Long Yeo. Seth Lazar's work on this paper was supported by ARC Discovery grant DP170101394.

[2] Dorsey [2009]; Kamm [2007]; Ridge [1998]; Scanlon [1998]; Temkin [1996]; Voorhoeve [2014].

[3] Anscombe [1979]; Finnis [1991]; Thomson [1990].

[4] Some might argue that we should reserve the idea of an 'absolute value' for a value no amount of which can be sacrificed, whatever the consequences. We think this confines discussion of absolutism to focusing on its least plausible variant. Our taxonomy includes that species of absolutism (we call it *Extreme Absolutism*), but also includes others kinds of absolutism, some of which are actually defended by practising philosophers.

[5] Though see Arrhenius and Rabinowicz [2005b]; Jensen [2012].

[6] Huemer [2010]; Isaacs [2014]; Jackson and Smith [2006]; Jackson and Smith [2015].

[7] Aboodi et al. [2008]; Hawley [2008].

[8] For that literature, see for example Arrhenius [2005]; Arrhenius and Rabinowicz [2005a, 2015]; Carlson [2000, 2001]; Dorsey [2009]; Griffin [1986]; Jensen [2008].

[9] Different philosophers have used different terminologies to refer to similar but not always identical concepts. In particular, the concept of 'lexical priority' is sometimes used to describe relations similar to those that we describe in terms of absolute preference orderings and superiority relations. Lexical priority is most commonly used to describe positions close to what we call *Moderate Absolutism*. Interpreted literally, lexical priority is committed to an infinitist representation (or something similar), and so we think doomed to failure.

[10] *Weak Absolutism* is clearly a species of absolutism according to our definition, since it posits a higher consideration, some amount of which is absolutely preferred to a lower consideration. The fact that *Weak Absolutism* is a mundane feature of most people's theories of value should encourage deontologists who wish to apply it in their ethical theory.

[11] Even Nozick, the absolutist par excellence, actually punted on this point. See Nozick [1974: 30].

[12] Kamm [2007]; Lazar [2017d]; Scanlon [1998]; Taurek [1977]; Voorhoeve [2014].

[13] Raz [1986, 1999]

[14] Decision theory is no more settled than ethics. This is a simplified statement of the orthodox view, and it punts on difficult questions like whether evidential or causal decision theory is correct. For two overviews, see Briggs [2014]; Buchak [2016].

[15] Jeffrey [1983]; Savage [1954].

[16] Here we reprise points made in Lazar [2017c]. It's worth emphasising, of course, that many decision theorists take the same view. There is growing opposition within decision theory to the representation theorem-based approach. See for example Easwaran [2014].

[17] We intend to be more or less metaethically neutral here; we are ruling out only the notion that morality is whatever you want it to be. We need not take a stand on which

broadly objectivist metaethical theory is the right one.

[18] Here we draw on Colyvan et al. [2010]; Lazar [2017b, 2017c].

[19] This assumption underpins Norcross's early objections to absolutism (Norcross [1997, 1998]), as well as discussions by Dorsey [2009]; Ryberg [2002]; Temkin [1996]. It is also a background assumption in Huemer [2010], and in parts of Isaacs [2014]; Jackson and Smith [2006]. Notice that this is particularly prevalent in attempts to model the relation of lexical priority. One could even treat it as definitional of that relation; if so, then we think that is all the more reason to do away with lexical priority, and focus instead on *Extreme*, *Moderate*, and *Weak Absolutism*.

[20] This kind of objection is discussed in most of the papers in the preceding footnote.

[21] Arrhenius [2005].

[22] We think the first to do so was Carlson [2000]. But see also Arrhenius and Rabinowicz [2005a]; Binmore and Voorhoeve [2003]; Broome [2010]; Jensen [2008].

[23] Simple additivity is the property whereby, holding the scale of value constant, every token of *r* adds as much to the total value as every other token of *r*.

[24] Some might baulk at the idea of using diminishing marginal values to understand cases like life for headaches. We cannot address those worries here, except to recall that axiological absolutism can be interpreted either as a substantive moral theory in its own right, or as a representation of deontic absolutism. In the latter case, many of the objections to diminishing marginal value functions would be side-stepped.

[25] We do not mean to imply that the infinitist approach has *no* defence against the sequence objection—Larry Temkin [1996] uses it to argue that we should reject transitivity; Dale Dorsey [2009] argues that sudden cliffs in value make sense on a multi-factorial account of value. As we note in the next section, however, the infinitist approach also raises serious problems for absolutist decision theory.

[26] Hájek [2003]; Bartha [2007].

[27] Of course, we can avoid some of these results by shifting to an unorthodox decision

theory. There are also mathematical techniques that might help. But these solutions would still leave the objective absolutist value function vulnerable to the sequence objection identified above. Bartha [2007]; Bostrom [2011]; Lee-Stronach [2016].

[28] This involves some contentious modelling choices; our aim is for a proof of concept, not to defend a specific set of numbers.

[29] For further discussion on this point, see Lazar [2017b, 2017c].

[30] After writing this paper, we became aware of a new defence of something like the threshold approach: see Zollman et al. [Forthcoming].

[31] Nozick [1974: 30].

[32] Jackson and Smith [2006: 271]

[33] Lazar [2017b].

[34] Huemer [2010] implies something like this.

[35] As one of us has argued at length elsewhere, as the probability that one's target is innocent increases, the moral disutility of killing her, if she is innocent, also increases (Lazar [2015, 2017a]). But we set that point aside here.

[36] Arntzenius et al. [2004]; Buchak [2013]; Elga [2010]; Hedden [2013].

[37] Jackson and Pargetter [1986]. See also Goldman [1976].

[38] Nonetheless, our case is indebted to those cases, especially as advanced by Huemer [2010]; Jackson and Smith [2006]; Norcross [2009].

[39] That is, we stipulate that the $100 won't make the difference to the satisfaction of any important interests.

[40] Norcross [2009] recognised this problem with objectivist absolutism.

[41] This approach is inspired by discussion with Sergio Tenenbaum, whose approach in Tenenbaum [2017] is to argue that *both* levels of permissibility matter (simultaneously). Douglas Portmore has also developed work in this vein (Portmore [2016a]; Portmore [2016b]), in which he has argued for 'maximalism', in our terms roughly the view that campaigns, not individual acts, are the proper object of moral evaluation. Two new

papers, which came out while this paper was under review, also address this question—one of them in the context of risk. Brown [2017] offers a distinct defence of maximalism from Portmore's. Kirkpatrick [2017] takes a similar position to Tenenbaum, arguing that the language of permissibility is ambiguous over whether campaigns or acts are the proper object of moral evaluation, and concluding that both levels matter. Interestingly, while moral philosophers seem to lean ether towards maximalism or to saying that both kinds of permissibility matter, decision theorists seem to heavily favour focusing on very narrowly described acts—see Elga [2010]; Hedden [2013]. This is *prima facie* evidence in favour of our preferred stance, that one size does not fit all.

[42] This case is similar to Huemer's example, in which we must choose between a 0.11 risk of killing one innocent person and a 0.1 risk of killing 1,000 innocent people. He thinks that individualists such as Aboodi et al. will separate out the risks to each of the 1,000, and so think it preferable to subject each individual to a 0.1 risk of innocent death than to subject one person to a 0.11 risk of innocent death. Set aside for now the fact that in one case the risk is that the party is innocent (death is assured), while in the other case we know the 1,000 are innocent, but death is uncertain (on this difference, see Lazar [2017a]). In this case the risks to the 1,000 are obviously not independent: you don't get any of the good unless you subject all 1,000 to that risk. So obviously the expected disutility will be much greater than running a 0.11 risk of killing one innocent person.

[43] A particularly pressing problem, for which we have no solution: what should we do when we are uncertain whether our act should be assessed as part of a campaign, or on its own?

[44]Jeffrey [1983: 154]; see also Nover and Hájek [2004: 247-48].

[45] Dancy [1993]; Kagan [1988].

[46] See, for example, John Broome's discussion of 'strong separability', which implies that the well-being of multiple individuals cannot have a diminishing marginal value function, because the value of each individual's well-being is independent of the value of

anyone else's well-being (Broome [1991: Chapter 4]). As a counterpoint, consider the deontic absolutism of philosophers like Kamm [2007], Scanlon [1998], Temkin [2012] and Voorhoeve [2014], who all argue that the moral significance of a person's interests in a choice can depend precisely on which other interests are at stake. None of these philosophers defends the diminishing marginal significance of interests across persons, but their views could arguably be represented with a value function of this form.